Measuring Operational Resilience
Transcript

Part 1: Answer Questions, Inform Decisions, and Affect Behavior

Pamela Curtis: 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 Pamela Curtis. I'm a senior researcher at CERT, working on operational resilience.

I'll be your podcast moderator today because Julia Allen, who normally moderates CERT podcasts, will be discussing her team's work in measuring operational resilience. In the interests of full disclosure, I am a member of Julia's team.

We've posted five previous podcasts on various aspects of CERT's operational resilience work and the CERT Resilience Management Model or CERT-RMM as we'll refer to the model from here on -- that listeners may want to check out as background for today's podcast. Links to those podcasts are in the show notes.

Welcome Julia; glad to have you as a podcast speaker.

Julia Allen: Hi Pamela. Yes, it's kind of funny to be on the opposite side of the call. But I'm looking forward to telling everyone about all our great work.

Pamela Curtis: Good. All right, let's begin. By way of background, would you please define operational resilience as we use it in our research?

Julia Allen: Of course. And there's a lot more, as you had mentioned, there's a lot more back story and foundational content on our resilience work in the previous podcasts. So I would certainly encourage listeners as well.

When we define and use the term 'operational resilience', we're really talking about a quality or an emergent property that an organization exhibits -- that is, their ability to continue to provide service, critical service, customer facing service, whatever it is that they're known for -- in the presence of or while operational stress and disruption is occurring; particularly if it does not exceed its limit. Occasionally there will be things that will happen that are beyond all bounds, and so you have to regroup and recover. But we're talking about that you can keep on ticking even when something bad is happening and provide the services that customers really count on.

Stress and disruption, as we define it in our work, comes from, most often, the four categories of operational risk. And those are deliberate or inadvertent actions of people; insider threat is an example. We typically think about systems and technology failures; such as a security incident that might compromise a critical service.

The third category is failed internal processes that might result, for example, in customer data being publicly disclosed. And the typical ones are external events: severe weather, fire, earthquake, flood. Those are the kind of things when we talk about operational stress and disruption that there we're referring to. And in the model, what we do is we identify the

functions and activities around those domains or disciplines: information security, business or service continuity, aspects of (IT) operations, and how well the organization is managing those. And the model defines those and the relationships among them. So that gives the listeners a little bit of the foundational definition.

Pamela Curtis: Okay great. So operational resilience is, as you said, what emerges from those activities.

Julia Allen: Correct.

Pamela Curtis: Defined in the model. Okay. Well we've been working on defining approaches for measuring operational resilience since January of 2010. Would you fill our listeners in on the purpose of this research and provide a bit of history?

Julia Allen: Sure. I'm going to chunk this into three responses because we have three reports that we've published that match up with this. So as you mentioned, we started at the beginning of 2010. And the foundational work that we did the first year was to define the key research questions. What questions are we trying to answer or address or inform with this measurement research? And we also came up with our key terminology and foundation of principles.

So some of the key questions, at the highest level, include how resilient is my organization? As you said, it's an emergent capability or property. So how able am I as an organization to withstand a punch or to keep on going? Am I resilient enough? How resilient do I need to be -- in other words, what's my threshold or what level of resilience am I aspiring to? Do I need to spend more money and if so, on what? And what am I getting for what I've already spent? Fundamentally I guess the real operative question is what is the business value of being more resilient? There's many things that decision makers and leaders need to invest in and does resilience and increasing resilience make the cut?

For our measurement work, the key question is: What should I be measuring to determine if I'm meeting my performance objectives for resilience? So once I've set all those thresholds and targets then how do I measure to see if I'm achieving those? Fundamentally measurement is about informing decisions and affecting behavior, so that leaders can control current operations and, as best they can predict the future.

Another key aspect of our 2010 work was to begin to categorize and define different types of measures. And I'll just mention two of those as most noteworthy: implementation measures versus effectiveness measures. So implementation measures -- and given that we're doing this against a process model, CERT-RMM —- talk about is the process being performed, and to what extent? So am I doing what I said I intended to do? But it doesn't make any judgments about how well or if the outcome or result of the process is actually helping me become more resilient. So that's what effectiveness measures are, is how well am I achieving my results and is it allowing me to become more resilient?

So a few examples might help illustrate. So an implementation example might be something like the cost or schedule to perform a process. You have those but they don't really tell you a whole lot about “am I better or worse?” So the companion effectiveness measure might be difference in planned versus actual cost and schedule over time. So am I getting, am I hitting my targets or am I improving? And is this allowing me to do more with less? Another effectiveness measure might be change in resources needed to support the process over time. Hopefully either I'm adding capability, so there might be a slight uptick in resources, which is

expected. Or I'm getting better and better at tailoring and making my process more repeatable, so my resources are declining.

Maybe something a little bit more technical -- an example: The cost and schedule to detect and respond to a security incident. Again, if I know -- or the number of security incidents. So let's say I have those count-type measures. But what does it tell me? What I really care about, on the effectiveness side, is the reduction in impact and consequences due to a security incident. For example, is it taking me less time between the time I detect and the time I respond and recover? Or do I now have fewer incidents, because of the improvements I've put in place, that require either regulatory agencies or law enforcement agencies to be involved? Because sometimes that's a fair amount of overhead. So that was -- that's what we did in 2010.

Pamela Curtis: Okay, great. And how about your work this year?

Julia Allen: Okay, so we've had two, I think, nice pieces of work; research that have moved this conversation forward. Early in this calendar year, in 2011, we've applied the foundational thinking from 2010 to help us identify a set of top ten strategic measures. And we'll talk about those a little bit further as we move on in our conversation.

There are 26 process areas in CERT-RMM and each of them have an example set of measures associated with them. So we used this as an opportunity to do a horizontal integration across all the measures, in all 26 process areas, and did a thorough scrub in consistency and clarity; we fixed some errors. And in the process of doing that, we actually identified 36 (by current count) global measures that apply across all 26 process areas. So we pulled those out of the individual process areas and defined those.

And the work that we've most recently done is to drill this down even further. Measurement occurs in a context and if you don't have a meaningful context defined, it doesn't make a whole lot of sense to measure. So again, given that CERT-RMM is a process capability model, we've defined and described what it means to have an implementation-level defined process. There are activities defined in the model. But it really doesn't -- it's more the what; it doesn't tell you how. And so we've created definitions of example processes, created templates for processes and supporting procedures, and then described how you define measures within the context of those processes.

Pamela Curtis: And that information is available in your reports, correct?

Julia Allen: Correct. So as I described it, the 2010 work is available in one report. The early work in calendar year 2011 is in a second report. And the process definition work, the implementation-level process definition work, is available in a third report. And those will all be cited in the show notes.

Part 2: Strategic Resilience Measures Compliment Existing Security Measures

Pamela Curtis: Okay. Now there appear to be many promising efforts to measure information security, which is one aspect of operational resilience. How does your research differ from or complement those efforts?

Julia Allen: Indeed, there are many communities, efforts, documents, standards, guidelines, that describe how to measure information security. And we certainly don't claim to be better or replace any of those. Our work is intended to be complementary. So by way of comparing and contrasting -- the way we see organizations use these guidelines and standards is typically to

collect what we call measures of type count: number of incidents, number of systems with patches installed, number of people trained, maybe number of compliance requirements met.

But having those type count measures, while useful -- because you can watch their trends over time -- we find they don't really inform decisions in the way that we describe in the model because you're really missing the context for how those measures can be used. So within our resilience measurement work, we have a definition of what is an operational resilience management system -- or if you prefer, operational resilience management program -- as the basis for measurement. And it describes at a very strategic level the relationship of organizational objectives; how resilience objectives and requirements derive from organizational objectives, including a consideration for risk tolerance; given those resilience objectives and requirements, the controls necessary to protect and sustain those high-value customer-facing services and assets.

And so in the ORMS, the operational resilience management system, we talk about both managing the conditions and the consequences. Conditions tend to be more on the information security side, the protect side. Consequences tend to be more on the continuity side. Even with all the protections that you've put in place, bad things happen and you have to be able to respond.

So when it comes to the community's work in information security measurement, there's lots, as I said, there's lots of bodies and codes of practice. And we've actually developed a crosswalk, a CERT-RMM crosswalk, that maps to all of these bodies of practice; such as ITIL, COBIT, the ISO 27000 series, BS25999, which is the British Standard for business continuity, and things like the Payment Card Industry Data Security Standard, PCI DSS. So you might think of those as complementary. So if you had a strategic measure as expressed by our measurement work, you could see the natural connections to all of the measures that are called for in these other bodies of practice. So we see them as complementary.

Pamela Curtis: Okay. And so you're helping to tie the kinds of things that are typically measured related to those codes of practice with strategic objectives that are the things that are actually affecting decisions.

Julia Allen: Right. We want to make sure that anything that's being measured -- because measurement is expensive to perform and sustain -- we want to make sure that anything measured has a direct tie to some business strategy, some business critical success factor, some business objective. And so we view the, in particular, the strategic measures that we've identified as being the bridge between the more detailed technical, tactical measures and the business objectives.

Pamela Curtis: Could you give an example of some of your strategic measures and explain why you selected them?

Julia Allen: Sure. As I said, the selection has been driven very much by business objectives, in concert with the structure that the ORMS (operational resilience management system) provides. So it's probably best to illustrate by an example.

So one of the strategic measures is *in the face of realized risk* -- in other words you have a risk that you're managing and it actually happened. So it could be a disruption in continuity of a critical service; it could be a security incident. What we say in our definition is that *the ORMS ensures* -- so all the things that you've put in place to be operational resilient; that's what we mean by ORMS -- *ensures the continuity of essential operations*. So you can keep on doing

what you need to do to satisfy your customers; *of high-value services* -- those that are most critical -- *and the assets that support them*. So again, in the face of realized risk, you're able to continue to provide critical service.

Sometimes we frame this as "the probability of delivered service in the presence of a security incident." So this is important because this is what business leaders care about. If I take a hit, how confident am I, or how able am I, in some measurable way, to know that I've prepared myself, I've protected my key services, and that I can -- the four asset types are people, information, technology and facilities.

So have I made all of those sufficiently robust to provide that service? And one of the supporting measures that roll up into that strategic measure is something like the confidence factor that risks from all sources have been identified and prioritized. So am I confident that I know all my risks and that I'm prepared and are managing those? Or there will be certain risks that I'll just accept or tolerate, and know that I have to invest in some recovery action if that risk is realized.

So to contrast that particular measure, let me just give you a brief description of the other types of strategic measures. So we have strategic measures around the extent to which resilience activities support or don't support business objectives. As I said, we always want to have that strong tie. We have measures around demonstrating that high-value services and assets satisfy their resilience requirements; that controls that address those requirements are effective and adequate, or if they're deemed to be ineffective and inadequate, that they're corrected; and that we're effectively managing the risks to the assets that could adversely affect our ability to deliver service. So again in our second report, those are laid out in more detail. But those are the general categories of strategic measures.

Pamela Curtis: All right, thank you.

Julia Allen: You're welcome.

Part 3: Thoroughly Defining a Measure; Getting Started

Pamela Curtis: So you've clearly been doing a lot of work in defining measures. So can you say something about the information that is required to thoroughly define a measure?

Julia Allen: Indeed, this is the hard part. This is where you really have to roll up your sleeves and do some hard work and, I think, where we've begun to appreciate when we suggest that people measure, just how much effort is involved in doing that.

So we've defined a measurement template and we have a number of examples where we've filled out that template. A measure is thoroughly defined by the typical six questions: Who, What, Where, When, Why and How.

So let me expand on each of those briefly. So who's the measure for? If you don't know who you're measuring for and who you're reporting to and how they're going to use it, you have some more homework to do. Who are the stakeholders for that measure? And on the collection and reporting side, who's actually going to collect the data, the measurement data, and make it available for analysis?

With respect to “What,” what is being measured? And as I said earlier, we recommend that it be done within the context of a process. So as part of what process, or processes, is that

measure being taken? Where is the data and information stored? You need to develop infrastructure, some type of data repository or a database that you're going to use to house all of this information. And it too has to be protected and sustained and secured.

With respect to “When” and “How” -- When, and how frequently, are the measures collected? Monthly, quarterly, biannually, annually? Why is this measure important versus others? Making some hard choices about what to measure and trading off what measures you care about the most. And what we've found, not surprisingly, is the most meaningful information with respect to measurement is conveyed by reporting trends over time versus point in time or single measures at a particular point in time. So sometimes it will take a while, a couple of reporting periods, to actually help you determine the value of the measure.

And then lastly “How” -- how is the data collected? What is the way that the measure analysis is reported? In other words, what's the visual representation? Is it a histogram? Is it some type of curve? Is it a Pareto diagram? Is it a Kiviat or spider chart? Thinking about the visualization because that's really key for decision makers, and that ties back to who your key stakeholders are. And how will the measure be used? What decisions will it inform? So as I said, we have a template that defines all these fields with some examples so folks can get an idea of what it takes to define a measure in the way we recommend.

Pamela Curtis: Excellent. Thank you. Well it's clear that if measurement is going to be done thoroughly and thoughtfully in an organization, it's going to require integration into business processes in the organization. So could you tell us what some of the first steps that you would recommend for starting a resilience measurement program?

Julia Allen: You made an excellent point, Pamela, about integrating into the way that the organization does business. So if you're doing measurement in another part of your organization or -- take the way that you're rolling up and generating your financial reporting information. There may be opportunities to integrate resilience measurement into an existing measurement or reporting process. And that would be optimal because then you're not bringing in yet another new thing. You're making it just part of the normal course of business. So that's a nice key point.

What we recommend is you got to first, as I said, identify who the measure's going to be for. Who are the sponsors and the stakeholders for this effort? What questions are you trying to answer or what resilience objectives are you trying to inform progress against? What information do you already have and what information do you need to collect and within what processes?

And then we recommend that you really start small. Take a couple of key measures. One of everybody's primarily interests is always incident handling, incident response, incident recovery. So maybe if you don't have some measures in place, maybe you pick a small number of measures in that particular process area. Collect them, analyze, report, and refine. Put a really, if you can, a barebones measurement process in place: the role that's going to collect, the role that's going to analyze, how you're going to visualize it, where you're going to store the data; whatever small pieces of infrastructure you need to put in place to start this out at a very modest level.

And then as time goes on, you really need to quantify the value of each measure. How much is it costing you to collect it? And are you deriving sufficient benefit for the cost that you've invested, both for an individual measure and in comparing measure A with measure B with

measure C. Because there's -- you're always resource constrained. So you have to make sure that you're getting good bang for the buck.
And then obviously refine and retire measures as you go. And I really want to emphasize retire. If there's a measure that everybody's used to seeing but it's not informing any decisions or affecting any behavior, our suggestion would be that you retire it, or refine it, or update it, or maybe combine it with another measure. But don't get caught in the inertia, which can happen as you're collecting and reporting measures. So those are some of the steps we'd recommend.

Pamela Curtis: Okay excellent. Thank you.

Julia Allen: You're welcome.

Pamela Curtis: Where can our listeners learn more about the measurement work that you've been doing?

Julia Allen: Well there are a couple of places. The CERT website, and in particular the CERT Resilience website, has lots of good information on our model work in total; some of the training that we offer; as I mentioned, a bunch of different podcasts on the subject which will help listeners understand the context of our measurement work.

And as I described, we have three reports cover the work done in 2010/2011. And we have plans for this coming fiscal year. So we'll obviously have more publications. But there are three available right now that listeners can review. And each of them has quite a number of supporting resources at the end of each report. Obviously this work rests on a very foundational body of knowledge. Noteworthy is our -- the SEI's -- Process Program, Software Engineering Measurement and Analysis work. We've drawn very heavily from concepts there. So you'll see other work inside and outside the SEI cited as supporting resources for our efforts.

Pamela Curtis: All right. Well thank you so much Julia for sharing with us today.

Julia Allen: Well Pamela, you're most welcome. I've very much enjoyed my collaboration with you in this body of work and appreciate the opportunity to share it with our listeners.

Pamela Curtis: And thank you for listening. And please join us again for the CERT Podcast Series.