Why Use Maturity Models to Improve Cybersecurity: Key Concepts, Principles, and Definitions Transcript

Part 1: What Are Maturity Models and Why Are They Useful?

Julia Allen: Welcome to CERT's Podcast Series: Security for Business Leaders. The CERT Program is part of the Software Engineering Institute. We're 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 principal researcher at CERT, working on operational resilience. Today, I'm very pleased to welcome back Rich Caralli. Rich is the technical director of CERT's Cyber Enterprise and Workforce Management Directorate. He is also the architect of the CERT Resilience Management Model (CERT-RMM), which we've talked about a lot on this podcast series.

Today, Rich and I will be discussing some key concepts and definitions of maturity models, some of the underlying principles and ideas, and how these can and have been applied to a wide range of disciplines and market sectors.

And just for our podcast listeners' benefit, our conversation is based upon two SEI whitepapers, which will be newly published with the show notes that we'll link to. And we'll also link to two of Rich's previous podcasts where we've discussed some of these ideas as well.

So, welcome Rich. It's really good to have you back on the podcast series.

Rich Caralli: Thanks Julia. I'm really glad to do this and get some good information about maturity models out there to the community.

Julia Allen: Great. So this is kind of interesting timing because maturity models as an idea and in application have been around for a long time. So from your point of view, why are we talking about this now? Why is it particularly time critical?

Rich Caralli: Well I think maturity models in general are getting this second wind. And you and I have been involved in this maturity model business for now 10 years, believe it or not. And I think we're starting to see that a whole new community is waking up to the benefits and the advantages of using a maturity model to transform their organizations and their businesses and the way they do things.

So I think now's the right time to have this conversation; particularly as maturity models become very prominent in the cybersecurity space, which we're starting to see with a lot of the government efforts and private industry efforts.

Julia Allen: So are there some other aspects that drive us to talking about this now?

Rich Caralli: Yes, I think it's useful and practical to maybe talk about how we got here because we went on a similar journey that the community's going to go on. Back in 2003, you and I started to work on how we could better improve security as a process. And we started to really look at the lessons that were coming from the Capability Maturity Model Integration (CMMI), the CMMI side of the house, in dealing with improvements in software and systems

engineering. And we started to apply those concepts, to the best we could, to more continuous processes like cybersecurity.
The banking and finance community in 2004 came to us and said, "A lot of what you're talking about resonates with us but we like the bigger concept; and that bigger concept is resilience." And that's how we got to our early resilience engineering framework work; and then the follow on, which was the CERT Resilience Management Model (CERT-RMM).

Since 2010, when we published the first version of RMM, there have been a lot of derivatives including a very nice piece of work that we do with our DHS friends called the Cyber Resilience Review, which helps owners and operators at the critical infrastructure level to look at how mature their resilience practices are.

And then the follow-on, and a related effort which we worked on last year, which was the ES-C2M2, which is the Electricity Subsector Cybersecurity Maturity Model; which has had a lot of impact in the power sector, and of course a lot of other sectors like oil and gas are now looking at building a similar model.

On the heels of a lot of that, working in that community, Mark Knight, who is a collaborator of ours who is the chair of the Gridwise Architecture Council, noticed as we were noticing that a lot of maturity models were being built in the operational and cybersecurity space; but they all weren't created equal. A lot of them purported to be based on CMM (Capability Maturity Model) but yet didn't demonstrate any of the architecture or characteristics of a CMM.

And so we started to think about the fact that maturity models can be a really effective tool for transformation if done properly, but if done improperly, can really have a detrimental effect. I look at it as if I've built you a roadmap to try to get from location A to location B, and if that roadmap was inaccurate or inconsistent, you likely aren't going to get to your destination. So this is why we started to talk about what does a maturity model really constitute and how can it be an effective tool? And all of this has been driven by me to measure effectiveness.

The operational and cybersecurity community wants to move away from compliance and implementation-based measurement. And measuring process maturity, if you remember back in 2003, was one way we thought we might be able to measure effectiveness—and it turns out that might be an effective way to do it.

Julia Allen: As I listen to you speak, I also think about the fact that a lot of the community thinks about process maturity as high overhead, labor and resource intensive, something that you just have to make this huge organizational commitment to.

And I think we've found with some of the applications -- Nader Mehravari talked about some of the same applications that you mentioned in our last podcast -- that you can really skinny down and make the ideas play and be scoped and tailored specific to the particular problem you're trying to solve, without having it be a burdensome overhead, right?

Rich Caralli: Absolutely. And I think ES-C2M2 is a good example of that -- still deploying some of the core principles but doing it in a way that makes the barrier to use much lower and the ability to use it over time much greater.

Julia Allen: Okay great. So let's start with the fundamentals; put some definitions and principles in place. So in your experience, in your observation, the work that you've done, what is a maturity model? And you've said a little bit about its usefulness. But can you say a little bit more about why you would want to even consider using one?

Rich Caralli: Sure. I think there are very broad and maybe academic perspectives. A maturity model is really a set of characteristics, attributes, indicators, or patterns that represent progression and achievement in a particular domain or discipline. Put more succinctly, it's the practices and the technologies and the processes that organizations use and stack in a progression, aligned to particular domains or disciplines. So these characteristics, attributes, Indicators, and patterns can be what the community uses to define their space.

It also then creates a benchmark against which an organization can assess their current level of achievement or capability and compare that with others in the domain or industry. And I think that that's actually why you see maturity models become so popular in the space because organizations, particularly critical infrastructure sectors, want to do a lot of this benchmarking.

When you start to talk about it in that perspective, you get this community view of a maturity model. And I would argue that that's what happened in the use of maturity models in the software and systems engineering domain. The community came around. The maturity model was a place to start. It reflected their experience and knowledge. It gave them a common language and shared vision that they could coalesce around; and it helped them define what improvement in maturity meant for organizations in those domains.

As such, it then became a framework for them to use to prioritize their next steps and their actions and their roadmap to drive improvement towards maturity. So all of those benefits that we saw and we have experienced as folks who have worked at the SEI and in the community with the CMMI, are all available to the cybersecurity community through the use of maturity models.

Julia Allen: Right. And what I've found through our applications of RMM is this whole idea of having a reference model, a dictionary, a thesaurus, a place where the terms are defined and where the relationships are defined. And I know when I use it, I can tie just about every aspect of a new security improvement initiative to some underlying piece of a goal or a practice or guidance in, in this case, RMM. Its utility as a reference model, as a compendium of knowledge, I've just found incredibly useful. Would you agree with that?

Rich Caralli: I agree. I think with RMM in particular, a lot of people are put off by picking up a model that's 1000 pages long, has a lot of information. But we've had a lot of reviewers who have publicly stated that what that gives you is a body of knowledge to work from. And it gives you the target to throw darts at.

No model is going to be completely correct, but it's something that everybody can gather around and use and improve over time because the models don't stay static either. So it's definitely -- to me something like RMM is almost more like a body of knowledge from which you can derive useful and practical models.

Part 2: Structure and Components

Julia Allen: Great. So let's talk about the structure of that body of knowledge. So when you talk about a maturity model -- and later on we're going to talk about a couple of different types of maturity models. But what are some of the key content and structural characteristics, some typical components of a maturity model? In other words, how would I know one if I saw one?

Rich Caralli: Yes. So let's start with the components because I think those are the building blocks. So maturity models by nature have levels. Those levels are often referred to as

maturity levels or capability levels, although they could have a lot of different names. The name of the level is typically characteristic or representative of the characteristics that exist at that level. So a pretty simple concept. And that's typically a horizontal construct in a model.

The more vertical construct is a domain -- so a group of like attributes. When we talk about models like RMM or the CMMI, those domains are really process areas or the grouping of different processes that are related. So for example in RMM, we talk about incident management as a process area. So that's a domain. And the domains are where the characteristics and attributes and indicators are then divided across levels.

The actual attributes are at the intersection of a maturity level and a domain. So for example at the domain of incident management, at maturity level 1, there would be a set of attributes that are reflective of and contain the core content of the model at that level. So that could be practices, for example. It could be technology.

It could be processes like you often see in a CMM. And then part of what glues all of that together is the ability to use that model -- its domains, its levels, and its attributes, which are the core concepts -- in an appraisal process where you can actually, in a consistent and repeatable way, examine an organization against that model and let them know where they are in their evolutionary journey. We call those appraisal methods typically.

The SCAMPI method of CMMI (Standard CMMI Appraisal Method for Improvement) is an example of one of those appraisal methods. ES-C2M2 for example has its own appraisal method that's loosely based on SCAMPI as well. And then there's a larger concept that often includes the concept of a maturity model and that's usually an improvement roadmap.

In CMMI, that's called IDEAL (Initiate, Diagnose, Establish, Act, Learn). So it's a plan/do/check/act cycle where you assess against the model. You find out where the gaps are, you make plans and action plans for those gaps, you implement, and then you diagnose again to make sure that you've actually improved. So this improvement roadmap is a larger concept into which the maturity model exists as one part.

One of the biggest challenges in a maturity model, as you might imagine, is the transition between the levels. So, effective maturity models that are impactful and useful and really transform communities have measurable transitions between levels that are actually based on empirical data that's been validated. So it is very clear that when I move from level 1 to level 2 there are indicators or markers that I can observe that tells me I've made that transformation. And as you know Julia, that's often one of the steps that's the most difficult in building and using a maturity model.

Julia Allen: Right. In other words, you've worked with subject matter experts and you've worked in a particular domain or a particular sector and captured what the contributors know. But as you said, the validation -- in other words, the observation: If I put this practice in place and I do this initial step -- let's say in incident management I do a reasonable job at detecting but I'm not doing any root cause analysis because I don't have that level of maturity yet. You have to actually validate that transition or that evolution or migration of practices really giving you some effective result, right?

Rich Caralli: Absolutely. And what we see is that the trend in the maturity models that we're seeing that are being put out there is to group practices or group attributes together in a way that is anecdotal. So I create step 1 and I put these practices there; then I create step 2 and I put these practices there. And there is no real logic in where we draw the line.

There certainly is no empirical data that says where to draw the line. And that can still be very useful for a Community because you might say to a community, “Take on these practices first and see how well you do with those; and then move to these set of practices.” And that may in fact be an indicator of maturity. But when we're talking about real maturity models that are very transformative, that transition between those levels has to be able to be proven.

And so this is by the way going to be the biggest challenge in the cybersecurity space because most of the models we put out, even those of us who do this for a living, those transitions are not often empirical at first blush. So we usually let those models kind of simmer with the community and then we start to learn about where those empirical transitions are. And of course we're learning that in RMM.

Julia Allen: Right, right. And I think that's why we sometimes get criticized as a community because of that lack of empirical validation and that lack of empirical evidence. But I think what you've said is very important, which is the body of knowledge, giving it a try. It does represent the expertise of a particular community.

And I think you used the word 'simmer' -- in other words actually get it out there. I know in all the derivatives that we've developed, when you take these models and apply them to a specific problem, that's where you begin to get some of this evidentiary information.

Rich Caralli: Right, if you're going to build a maturity model you need to be pretty tenacious because you have to know going in that the first model you build is going to be wrong. And iterations of that model in the future will also be wrong to some degree. But that's the important part of building it and using it in the community because that's where the community participates in defining what maturity really means for that community.

And over time, iterative changes to that model reflect more and more empirical data, more and more factual basis for the transitions, and really does reflect the community's practice.

Part 3: Three Types of Models – Progression, Capability, Hybrid

Julia Allen: Great. Well let's -- given that we're kind of laying some foundational principles and constructs here, let's talk about some of the types of maturity models that you describe in your white papers. And I know right now we've got three notional ones.

So let's talk about, first about a progression model and then later on we'll also talk about a capability model and a hybrid, just to give a preview to our listeners. So what is a progression model and how is it distinct in its features and characteristics?

Rich Caralli: Okay. So we go into pretty good detail on this in the papers that you referenced at the beginning of the podcast. So I'll go through this fairly quickly and folks can go there for the details.

But a simple progression model is just really a progression of scaling of those characteristics, indicators, attributes or patterns. So the level names or the definitions indicate progression -- meaning they focus on the domain-specific attributes. And the levels are often arbitrary; the grouping's arbitrary. So you might say -- if I'm using the example of accounting I might say the first level's pencil and paper and the next level is abacus and the next level's a calculator and the next level’s computer.

This is not measuring capability or process maturity, even though it's often confused with doing that. And the movement between the levels is not validated.
Now with that said, this can still be very useful for some communities because it at least gets the practices in a form that organizations can walk through them in a progressive way.

Julia Allen: Right, and do you have an example, a security example, of what a progression would be, just to make it a little more tangible?

Rich Caralli: Yes like authentication could be an example. So you might start with passwords; and then the next level would be strong passwords; another level up would be password changes on a specific interval like 60 days, then you might move to two-factor authentication; then to three-factor.

So there is maturity in the progression in terms of the techniques or the practice. And it could be useful but it isn't measuring capability per se, it's measuring implementation of a practice.

Julia Allen: Right, so it's a progression of techniques or methods or tools or technologies. But I think the key word you used is it's not measuring capability. So let's talk about capability models. How are they different or how do they expand or add value to a progression model?

Rich Caralli: So a capability maturity model often measures organizational capability based on a collection of those characteristics, indicators, attributes, and patterns, which are typically expressed as processes. So you often hear a capability model being referred to as a process model, or the use of the CMM, the process improvement, is being characterized as model-based process improvement.

These kinds of models reflect the maturity of the culture and the degree to which the capabilities are actually institutionalized in the culture. And this is a very important distinction. When you're trying to measure effectiveness rather than implementation, you might want to measure something like maturity, because what you're looking to do is not just to measure whether the practice exists or is being done but the degree to which the practice is embedded in the culture, institutionalized, and can be retained under times of stress.

So the levels in a CMM are often indicative of the degree to which you've achieved some level of process maturity, where the base practices typically exist at level 1. So for example, in the CMMI you might pick a base practice like requirements definition; and you can raise that practice from an ad hoc level to a managed level, then to a defined level, then to a quantitatively managed level and then to an optimized level.

It's still the base practice -- base practices in requirements definition -- but the degree to which these practices are institutionalized, embedded in the culture, retained under times of stress, measured and managed, starts to go up as you measure your levels. RMM has the exact same concept in it.

Julia Allen: Right, so just to make this more tangible by example -- so in your progression model, when you talked about authentication, you talked about going from passwords all the way up to two and three factor authentication, a progression of practice.

If I was going to take that same example and put it into a capability maturity model construct, I would be wrapping around it management and planning and measurement and the kinds of activities that would need to take place in the organization -- we also use the term “making it

sticky." So as you said, in the face of disruption and stress, that process is still robust, it's still -- authentication would still happen in the way that it's intended, correct?

Rich Caralli: That's right, because in a CMM, the dimension that's actually being measured is that representation of organizational capability. And this is important because it measures more than the ability to just perform a simple task. It looks at a broader organizational capability and it reflects the maturity of the culture, and again the degree to which those capabilities are actually embedded.

So the transitionable states in a CMM actually describe states of organizational maturity, relative to the indicators of maturity that are applied at each one of those levels. So it's a much more complex undertaking. You and I know as we were writing RMM how difficult it is to write a model with that level of structure and that level of architecture. But in practice when you start to see how that kind of model can actually help you measure not only whether or not you're doing something but whether you're doing it well -- which is the big leap here -- CMM is the only way to go to do that.

Julia Allen: Got it. Okay so let's kind of ratchet this up another level and hopefully this will be clear as we go forward. But we actually have one really good example right now, and many coming, that are a hybrid or a combination of the best of a progression model and hopefully the best of a capability maturity model. So can you talk about this hybrid idea a little bit?

Rich Caralli: Right, so we sort of coined the phrase 'hybrid' to define these more agile, better able to be implemented models that reflect not only the nice transformative progression half of a progressive model but also incorporate some of these stickiness concepts from the CMM. And so in essence what a hybrid model is, it's basically a progression model that defines the characteristics, indicators, attributes and patterns. So you still have that stacking like we talked about in the authentication methods where you're going from very rudimentary practices to very sophisticated practices.

Then it overlays a capability model approach to reflect the transition between the levels. In other words, when you leave one level and go to the next level, it's because your transition is based on the degree to which you're institutionalizing those practices.

So each domain then in a hybrid model becomes a logical grouping of practices, organized by the level -- the maturity levels -- but still representing progression as you move through the model. So the best way to describe a hybrid, if you've never seen one, is take the best of breed of a progression model and overlay the characteristics of a CMM on top of it, and get a model that gives you the ability to measure both progression and maturity at the same time.

Julia Allen: I think an example would help here. And I know we have one good one with ES-C2M2. So how is that combination reflected in that model?

Rich Caralli: So in ES-C2M2, we use what we call the MIL (Maturity Indicator Level) scale, which was a new scaling that we developed for these kinds of hybrid models that reflects the maturity scaling that you see in CMMI and RMM. And for the practices that we grouped at MIL level 1, those practices are performed but considered to be done in an ad hoc way.

So in MIL level 2, as we progress to the next set of practices -- so the next most mature practices -- at MIL level 2, those practices are typically documented; the stakeholders of the practices have been identified; the resources have been provided to do the practice; and standards that support the practice have been identified.

So when we step from MIL level 1 to 2, we not only incremented the practice -- meaning the practice got more mature -- but the stickiness factors also got more mature. We went from doing the base set of practices in an ad hoc way to doing the next highest level set of practices in a documented, more planned way.

And then as we move to MIL level 3, yet another iteration of practice maturity -- meaning we've now moved from changing passwords every 60 days to doing three-factor Authentication -- but we've also ratcheted up the stickiness factors. So at MIL level 3, the activities that are performed are guided by policy; they're reviewed on a regular basis; people are assigned responsibility for those practices; and the people who are doing them are highly skilled.

So in this way we get the progression of the practice in a domain and we also get some of the stickiness factors that measure the degree to which you're doing these things well or doing them better.

Julia Allen: Right, and my understanding -- I've not actually been in the field working with ES-C2M2 -- but my understanding is what this makes this appealing to the community that's using it is that instead of having the practices separated from institutionalizing or capability attributes or characteristics, you've got them all interwoven. So that as they actually implement or assess or try to identify an improvement path, it's much more clear what the steps they have to take are. Is that the message I should be getting?

Rich Caralli: Absolutely, absolutely. And the thing that I really love about a hybrid model is -- and if anybody has listened to this podcast for 30 minutes, they'll understand what I'm talking about -- the maturity concepts are complex. They're difficult to understand, they're difficult to apply sometimes.

In this particular model, we embed those in a way that the user doesn't really know and isn't completely overtaken by measuring that capability or that maturity aspect. They're focused on the practices and the maturity aspect of it is sort of a byproduct. So it takes the emphasis and puts it back on the practice like a progressive model does. But it gives you the benefit of the CMM kind of scaling without really having to be an expert in that space. And that's why I think we're going to see more models like this -- much more agile to use, much more, much easier to apply, and much easier to understand.

Julia Allen: Great. Well I know that we've just explored the tip of the iceberg. And as you said, the whitepapers I think add a lot more detail. But I know you've got some exciting other sources to point our listeners to, including some planned future activities. So can you say a little bit about additional information on this topic?

Rich Caralli: Yes absolutely. So as you said, we will be posting the two papers that sort of -- these papers sort of kicked off our effort with ES-C2M2 last year, trying to help the energy community understand all of these differences so that they could build a model that could actually be useable for them. Those papers will be posted with the podcast.

I also invite people to go to the CMMI Institute website. We talked a lot about CMMI and models like the software and systems engineering models. But there's also CMMI for Services, which is a nice cousin to CERT-RMM. So you'll see some of this CMMness that's in action.

There's also the RMM website where folks can download version 1.0 CERT- RMM, so they can see how we took those concepts and brought them to the operational side. There's also the

ES-C2M2 website where folks can download the very model we're talking about and maybe have that sitting alongside them as they listen to this podcast. And they can start to see some of these concepts come alive as they look at the actual model.

We're also working on a technical note that's going to describe the MIL scale and the MIL scale stands for maturity indicator level. And basically what it does is it takes the maturity levels that you see in CMMI and RMM and breaks them down into more manageable chunks. One of the complaints about RMM early on was the ability to transition from level 1 in the model to level 2 was a huge leap.

A lot of organizational structural things had to happen for that to be possible. So the MIL scale breaks this down into measureable and much more consumable chunks. And we're going to put that technical note out in the next couple of months, so folks can start to think about using this scale as they're building hybrid models.

And then with my friend Mark Knight and company and Austin Montgomery, who also works at the SEI, we're about to begin building a paper called "Maturity Models 201." And basically what that's going to be is a selected overview of maturity models that are in use today, the management of those models, and the community benefits of using such models. So it's going to give a broader perspective on how communities can be transformed through the use of maturity models.

Julia Allen: Well Rich, this has been an excellent, excellent conversation -- putting some key foundational concepts out there for folks to understand and digest and interact with. So I think you so very much for your time and preparation today.

Rich Caralli: Thank you Julia; always a pleasure.