

NIST Catalog of Security and Privacy Controls, Including Insider Threat Transcript

Part 1: Evolution of NIST Special Publication 800-53

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 principal researcher at CERT, working on operational resilience and software assurance.

Today I'm very pleased to welcome back Dr. Ron Ross. Ron is with the U.S. National Institute of Standards and Technology (NIST) and he is the project leader for the FISMA Implementation Project. FISMA is the Federal Information Security Management Act.

I'd also like to welcome one of my CERT colleagues, Joji Montelibano. And Joji and his Insider Threat Team have developed recommendations for control updates and additions to a NIST Special Publication that we'll be talking about today, 800-53; and based on our insider threat practices where some of those controls and recommendations have been actually implemented into Revision 4. This special publication is titled Security and Privacy Controls for Federal Information Systems and Organizations.

So that's plenty from me. Now it's time to get to the experts. So, welcome back, Ron. We're really glad to have you with us again today.

Ron Ross: Well thanks very much Julia. It's great to be with you and Joji today.

Julia Allen: And Joji, welcome to the podcast series.

Joji Montelibano: Thank you very much Julia; a pleasure to be here. And I'd like to say a big thank you to Ron and his team for allowing us the opportunity to comment on such a widely distributed publication as 800-53.

Julia Allen: Absolutely, it is broadly in use by all segments and sectors of our information security community. So yes, it's great to have the opportunity to discuss this work with both of you today.

So Ron, just to set the stage -- some of our listeners may not be familiar with this special publication. So can you say a little bit about its history and how it's being used today?

Ron Ross: Sure. NIST Special Publication 800-53 really goes back to one of the original publications that NIST was asked to write, based on our responsibilities under the Federal Information Security Management Act of 2003.

We had a couple of federal standards that were produced first: FIPS 199 and FIPS 200. But then very shortly, in 2005, we published our first iteration of the security controls catalog. And this catalog covers a range of safeguards and countermeasures, from management, operational, and technical based controls, that are deployed within information systems and the environments of operation where those systems operate.

And since 2005, we've gone through three full revisions, coming up to our current Revision 4. The most historic revision, prior to this one, was the Revision 3 where for the first time we entered into a partnership with the Department of Defense and the Intelligence Community under the Office of

Director of National Intelligence, and in collaboration with the Committee on National Security Systems to produce a unified security control catalog that covered both national security systems and non-national security systems. So we're always constantly upgrading the catalog through these various revisions for us to keep up with the threats and other new technologies that might come down road.

Julia Allen: Excellent, appreciate that background. So obviously today's conversation is about Revision 4, leading up with the background that you've described. So what prompted your release of Revision 4? What are the primary changes in content and areas of emphasis in this revision?

Ron Ross: Well this was another historic revision because it was probably the most extensive revision we've ever done. About a year ago, we did our first data call. This is the first time we ever went out to the public, before we updated the publication, to ask for their suggestions. And the primary drivers for the revision, I would say, go down to several different important areas.

The first is the threat landscape. Threats today are constantly evolving. Whether its nation-state level threats, organized crime, terrorists, hactivists, the threat landscape continues to change and we had to be able to react to that change.

The second area was the actual attack data that we've collected over the last several years. We have a lot of empirical data now -- knowing where the attacks are coming from; how these attacks are put together; where they're targeting. So the capabilities, the intent, and the targeting of adversaries against our critical systems -- all of that data then is collected and we can now use that as we try to update the catalog.

There were also significant gap areas. One of the major gap areas was in the insider threat, which really prompted our relationship with CERT at Carnegie Mellon. And then there were other gap areas that we noticed -- application security, advanced persistent threat. We have a new appendix for privacy controls.

And then there was this other very important area that a lot of folks have forgotten about today. It's trustworthiness and assurance of information systems; in essence making sure that we not only have a robust, continuous monitoring program, but we've built it right to begin with, so we have something very solid to monitor over time.

So those were the primary drivers for Revision 4. And we got over 1000 comments coming in, in that data call. And that prompted our work with your organization and others to really bring this catalog up to a very state-of-the-art set of safeguards and countermeasures, with the appropriate breadth and depth that we need to support the war fighters, the Intel community, and all of our civil agencies. And of course any organization in the private sector, which chooses to use the catalog on a voluntary basis, is always welcomed as well.

Julia Allen: So speaking of use, Ron, -- and obviously this special publication has a long lifetime -- can you just say a little bit of how you've seen the evolution of its use from maybe going back a couple of years and -- with Revision 3 being a first major change or paradigm shift, and now Revision 4 -- how have you seen the change of use of this special publication?

Ron Ross: That's a great question Julia. And I've noticed a very significant change and I think this is a natural evolution of the catalog. When we first published the catalog -- and we still do this today -- we have a set of baseline security controls. These are a starting set of controls that are applied to one of the three categories of systems that we define under our FIPS Publication 199. It's either a low, moderate, or high impact system -- where impact describes the potential adverse effects on your mission if that system is compromised or breached. Originally, I think people implemented the baselines without a lot of a change. But with this greatly expanding catalog of controls, going from

600 controls up to I think 850- some in this iteration, we're now getting to the point where we're specializing our security plans. We're taking those baselines and we're applying the tailoring guidance that we provide in the control catalog, and we're building security plans that are very specialized to mission, technology, or environment of operation.

So the military, for example, will create a military tactical overlay, which is a specific set of controls for military operations in combat environments. The Space Command may do the same thing. Our Nuclear Regulatory Agency created an overlay concept for nuclear power plants. So it's taking the basic controls and applying them to very specialized situations, which is really different than taking them right out of the box and trying to apply the controls to all situations and all types of systems.

Julia Allen: I'm really glad to hear that that kind of application and tailoring guidance is out in the community. Because as we all know, you can't secure everything and you have to make some very tough decisions about where to invest your limited resources. So I'm pleased and appreciate that explanation.

Part 2: Recommended Insider Threat Controls Based on More Than 500 Cases

Julia Allen: So let's get Joji into this conversation. Joji, can you say a little bit about the recommendations that our Insider Threat Team made for this revision -- summarize them and then maybe we can do a deeper dive on a few of your favorites.

Joji Montelibano: Absolutely. First of all, let me just follow on, on what Ron described about the evolution of 800-53. We like to think that we as an organization, CERT, have aligned very closely with NIST -- growing coordination with both the Department of Defense and the Intelligence Community in developing especially Rev 3.

And as Julia knows, and the folks who follow these podcasts, CERT is a federally funded research and development center. And we are also closely aligned with those two communities. So this work goes hand-in-hand with both our organizations' missions, both CERT's and NIST's.

And following up also on what Ron said about baselines, I like to think, having grown up with NIST 800-53, that NIST got it right the first time. And that's what Ron was talking about, baselines. The baselines pretty much remain, I think, unchanged.

So we, with the exception of one additional control that we proposed, most of our recommendations were targeted to the specialization that Ron was describing: tailored guidance. Because the baselines pretty much -- if organizations implement baselines correctly, they're in pretty good shape. And but you have to go a step further. Like Ron said, the one gap we saw was insider threat.

And along the themes of empirical data, our recommendations were 100 percent empirically based, 100 percent objective. At the time that we made our recommendations, we had over 500 cases that we looked at -- 500 cases of actual insider attacks. And we took the attack vectors from our case studies, from our database, and we looked at NIST 800-53 Rev 3 and said if an organization were to implement these controls, would they be able to address these gaps?

And so if you think about it, 800-53 has a lot of controls. And we were very hard-pressed to find gaps because we think it's an excellent publication. It gave us an opportunity to really focus in, to hone in on specialized guidance. And so on a high level, we gave recommendations to 10 control families; overall about 20 new recommendations for those families. The two main families we focused on were the access control family and the personnel security family. And then we also recommended a new control for the project management family, the PM family.

Julia Allen: Great. So why don't you keep on that thread, Joji, and tell us a little bit about either some of your favorites or the control families that you focused your recommendations on?

Joji Montelibano: Sure. Our recommendations spanned the breadth of very granular technical recommendations to, I would say, enterprise, organization-level recommendations. I'll take the three families that I just mentioned: the access control, the personnel security, and the project management families as examples.

So the access control family -- an example of a very granular control that we recommended was for AC-2, Account Management, where we thought there wasn't enough emphasis on shared accounts. And in our studies we've seen that insiders use shared accounts to launch their attacks -- mainly because shared accounts, like administrator accounts and privileged accounts, they tend to hide your tracks. A lot of people don't track who is using those shared accounts. And those shared accounts give you carte blanche and they pretty much allow you to do as you please in an organization's infrastructure.

But one of the most pertinent findings that we've had over a decade of research is that insider threat is not just a technical problem; it's not just an IT problem. And 800- 53 addresses that as well. That's why they have families that are not only focused on technical controls.

And this is where I'll move on to the personnel security part. The breadth of our recommendations for the personnel security part was inter-departmental communication, meaning human resources, legal, physical security; special emphasis on human resources. Those departments have to communicate not only with information technology and information security but with each other. So a very basic tenet that we highly recommended inclusion into this Rev was notification.

If someone gets transferred, if someone gets terminated, if someone gets sanctioned, you have to have a notification mechanism in place to make sure the organization knows about it. And even today we see that mature, relatively mature, organizations do not have a process in place to let critical personnel know when someone is let go or when someone leaves the organization. And so what you have are these accounts that are lying around that people still have remote access to an organization they left. And sadly, especially in the cases of termination, those departures do not occur in a very amicable way. And when bad feelings are involved, bad things can happen.

And finally, the one new control we recommended, which we were happy to see included in the draft of Rev 4, was PM-12, was the creation of insider threat program. This is not novel to us by any means; in fact, the Office of Management and Budget released a memo, M-1108. It didn't mandate the creation of an insider threat program but it did inquire about the existence of an insider threat program within federal agencies.

And the new Executive Order, released in October last year, EO- 13587, went a step further. It superseded OMB 1108 and it went a step further in establishing an insider threat task force and making sure that, at least in the classified area, that agencies are creating insider threat programs to help safeguard classified data.

And the major characteristic of an insider threat program is the cross-communication, the cooperation and coordination of the different departments that I mentioned. It's not just an IT problem. You have to have HR, legal, physical security, IT, and information security involved, in order for such a program to be successful.

Julia Allen: Oh Joji, thank you so much for highlighting some of those key controls and the history from which they derived.

Part 3: Developing Revision 4 and Beyond; Monitoring Control Effectiveness

Julia Allen: So back to you Ron. In your earlier remarks you said that this was a noteworthy effort on this revision because you put out a call for comments; received over 1000 comments, obviously ours being in that mix on the insider threat recommended controls.

So how did your team decide to handle all that? And specifically how did they decide to address the recommended CERT updates?

Ron Ross: We have a very disciplined and structured process for taking the public comments. We have a matrix we set up and every comment is laid out in a table. And every one of our team members gets an opportunity on their own to take a look at the comment and make their own recommendations.

And I go through and do the exact same thing on every comment, looking, giving it my perspective. And then I take all the team input, in addition to my perspective on the problem, and would try to come up with a consensus-based solution -- whether we make a specific change or we reject a specific comment for whatever reason.

But again, this was an enormously important update, especially in the insider threat area, and all the terrific input that we got. And I just want to again give great credit to the CERT organization at Carnegie Mellon. You guys really are the world-class experts in insider threat. And we very much appreciated the opportunity to receive the input that really was prompted by some of the big events that happened over the last couple of years. I know you guys have been tracking this for over a decade. But I think the thing that really brought it to the forefront, really to the core of its importance, was the WikiLeaks episode.

And so as we were looking at all the inputs from your organization, from CERT, we're trying to decide -- you say, "Would these kinds of changes have stopped something like WikiLeaks?" And I think the answer is yes. Had we put these controls and the enhancements in place a couple of years ago and they were actually fairly rigorously enforced, I think we can absolutely stop some of the things like WikiLeaks. It was an enormously damaging type of a breach.

And so that's the process we go through. And then we come up with our final -- we call it a markup on the previous version. And then that's what really goes out to the public in our public review process. And again, I've become such a big fan of the public review process. Because this document, in draft form, goes out not just to the federal agencies but it goes out to anyone who wants to download the document from our website.

And we get comments from not just the federal side, but we get lots of comments from the private sector. We get comments worldwide. So at the end of the day, the objective is to have the best set of safeguards and countermeasures we possibly can. The catalog of controls is going to get broader and deeper over time because the threat is expanding.

And again, making sure that you pick the controls, as Joji was mentioning, that are necessary to really do the job. We'd like to be able to stop all cyber attacks at the boundary. However, we do know that some of those attacks are getting through at the high end, especially with the advanced persistent threat (APT). And there again, the APT really emphasizes the importance of an insider threat program. So if adversaries do get through your boundary -- your initial defenses -- you have the ability to detect, respond, and limit the damage that those adversaries can do once they are detected.

So again, these additions in the insider threat area I believe are one of the most critical things we can do today, most important things we can do today, for this very, very high visibility area.

Julia Allen: Well at the expense of putting you on the spot, because I put Joji on the spot with this question. We'll give it a try. So what prompted my asking this is your talking about the WikiLeaks incident.

So in your work on this controls catalog and the other related special publications, do you have an opportunity to observe the efficacy or effectiveness of some set of controls? In other words, someone picks up NIST 800-53. They implement a set of controls. Does your organization have an opportunity to actually get some data back on whether or not the set of controls that were implemented actually produced the desired effect; to use your example, actually stop something like WikiLeaks?

I know I've asked Joji how do we know that our insider threat practices are actually accomplishing the job that we claim that they will accomplish? So I'd like to throw that out to you and see what you think.

Ron Ross: Well that's another great question. I think this ties in very nicely to our continuous monitoring program. So our new philosophy is "build it right initially and then continuously monitor over time to make sure the security state of your system and your environment of operations is maintained."

In the case of any of the controls we had, to include the insider threat controls, the best thing we can do is over time observe the effect of those controls. And the way we do that is, again, we have a very close collaboration with our colleagues on the Defense Department side and the Intelligence Community.

And we can look at the types of attacks that are occurring over time, and we can use that empirical data. So if the attacks are occurring in areas where we've already deployed controls, that may tell us something very important. Either the controls are not effective and we need additional controls, stronger controls, or we don't have the right set of controls.

And so that's the way you measure over time. One of the key tenets of continuous monitoring is to determine the effectiveness of your deployed controls over time. And that can really only be done with observing the empirical data, the types of cyber attacks that are occurring.

Now the one unknown is these attacks are evolving over time. So you never know what the adversary is going to think up tomorrow or next month or next year. There are always surprises out there. But we can learn a lot from this type of empirical data. And that's the way we try to respond, to make sure we're developing the right controls and we're making the right recommendations to our customers.

Julia Allen: Excellent. Well thank you Ron. And Joji, did you want to add anything to that? Because I know we've talked about this a little bit.

Joji Montelibano: Yes absolutely. And we are on the continuous monitoring bandwagon, if you want to put it that way. We've worked closely with DHS FNS (Department of Homeland Security Federal Network Security) to help with the maturing of FISMA. And in our own work with our customers, we implement our controls. We practice what we preach and we actively solicit input from those organizations that implement our controls.

The controls, the very controls that we've recommended inclusion for in 800-53, we've actually deployed to some of our customers. And we are in regular, sometimes daily, contact with those customers to see how those controls are working. And that always feeds to us and which in turn we feed to NIST.

Julia Allen: Great. So Ron, as we come to our close, what are -- here you are, you've just gone through this major effort to get Revision 4 out and I'm going to ask you now what are your future plans for this special publication and those related to it? But looking down the road, do you have some areas that maybe didn't make it into this revision or some plans for maybe a year to two years out?

Ron Ross: Well I think the future plans are fairly well laid out at this time. Our vision continues to be carried out in this whole project with the Joint Task Force. Certainly we want to make sure that we've finalized Revision 4. Our plan right now is to lock that down sometime in the July time frame.

We may come out with a final draft depending on how many comments we get on the public review process. And we don't like to add a lot of new stuff in a final version. We want to make sure our customers always have a chance to look at whatever some of the new controls might look like.

So getting that finalized is our primary objective. And I think for the future, we're going to try to maintain a two-year update cycle. Now the next Revision 5 will not be nearly as extensive as this one because this is a zero-based look. But for the future, I think Joji picked up on a very important point, with the insider threat. It's the integration of these controls over an enterprise level.

And again with the WikiLeaks case, we have the IT department involved. We have the security folks and the HR. But the communications across the enterprise, and how these controls are implemented in a cross-cutting effect across the enterprise, that's going to be a future direction for 800-53.

And again, continuing the notion of the overlay concept where we deploy the baseline set of controls, which are best practices across all systems, and then specialize that security plan for the specific mission, environment of operation, or technology. Like cloud computing or mobile may have certain types of overlays that are a little bit unique to their types of operations. That's the direction in the future. Again, making this control set the best we can make it to serve our customers and helping them defend their systems to the greatest extent possible.

Julia Allen: Terrific. Well, do you have some places to point our listeners for additional information?

Ron Ross: Well, everything that we do at NIST is on our Computer Security Division website. And any listener out there can go to the website and they can download any of the Joint Task Force publications under the FISMA umbrella.

They can get literally hundreds of other documents and other more specific guidance in technology areas, evaluation, cryptography -- all available on the NIST Computer Security Division Resource Center website.

Julia Allen: Great. And Joji, places that you'd like to point our listeners?

Joji Montelibano: To our Insider Threat website, and the CERT website. And that'll have documented controls and guidance on how organizations can implement these controls to combat insider threat.

Julia Allen: Right. And I would be remiss if I didn't mention that we've done a number of great CERT podcasts with Joji's colleagues on insider threat. So I would encourage our listeners to give those a look as well.

Well first of all Ron, I cannot thank you enough for your time, your attention, your preparation for today's podcast, and the high value contributory work that your organization is doing. So thank you so much for your time today.

Ron Ross: Well thank you very much Julia and to Joji. And again, many thanks to the CERT organization at Carnegie Mellon, the Software Engineering Institute, for your continued support of this very important area and helping the nation defend itself against these ongoing cyber attacks. We very much appreciate that.

Julia Allen: You're most welcome. And Joji, it's great having you as a contributor to the podcast series and the foundational work that you and your team have done, resulting in this significant accomplishment of getting our insider threat controls reflected in NIST 800-53. Thank you for your time today.

Joji Montelibano: Thank you so much Julia. And let me reiterate your thanks to Ron. It was a privilege and a pleasure to work with you sir.