Title: Leveraging Security Policies and Procedures for Electronic Evidence Discovery Transcript

Part 1: Challenges in Presenting Digital Information as Evidence

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 John Christiansen, an attorney specializing in information security, compliance, and risk management, and a founder of the Information Security and Compliance Risk Management Institute. Today John and I will be discussing how business leaders can be prepared when called upon to submit digital information as evidence in legal proceedings. So welcome John, glad to have you with us today.

John Christiansen: Thanks Julia. I'm really glad to be here.

Julia Allen: So from a legal perspective, what are some of the unique characteristics of information when it's in digital form?

John Christiansen: What's probably unique is the fact that it's the law trying to catch up with it right now in yet another arena. We've sort of been through a number of iterations where legal principles, legal standards try to catch up with what to do about digital information, electronic information, as opposed to good old-fashioned information of the kind the law's used to dealing with.

The law looks backwards. We deal with precedent. Whether it's in litigation, in cases, or in writing legislation or regulations, the question always is, “So how did we used to do this and how are we used to doing that?"

And that worked just fine for evidence, kind of forever, for literally centuries, in English and then American Law, because we knew, "Okay, it's on paper." And paper has certain characteristics; it doesn't degrade very fast and you can preserve it. And writing is pretty stable; you can preserve records for a very long time in writing. And you can tell whether records have been changed. We knew how to deal with evidence in very straightforward ways that were really well embedded in the legal system.

Well, as we've experienced, the characteristics of information have changed from then. And we've seen this before, as I say, in the law. Folks who've dealt with issues around things like copyright and digital rights management have already seen how the law has had to change a whole lot, and is still changing a whole lot, to try to deal with the different characteristics of digital information.

But compared to the hardcopy world, digital information doesn't actually exist in the way that it's presented to us. It exists as bits and bytes in a file that then gets presented to us through being processed in various applications. And data can be pulled from a variety of sources to get you the equivalent of the piece of paper that you used to have.

If I'm looking at say a tablet PC, the display there – if I'm in a wired network somewhere or a wireless network Somewhere – the display that I'm looking at, I guess even now on my computer screen, what I'm looking at has been pulled from a number of different data sources, and it's an image that is created right now. When I close the window or close the browser or when I turn off the computer in particular, that image is really pretty much gone, and it's going to be recreated again when I need to pull up that information.

We take for granted that this will happen in reliable ways so that what I look at tomorrow is the same thing that I looked at today. But, in fact, that depends on the functioning of all sorts of background components. And of course we know this. You know this very blatantly when you're no longer able to pull up that file, or when you pull it up and it's corrupted and you're not able to read it, then you may not know what went on. Generally I don't know what went on until somebody else investigates it for me. But you do know that you're no longer able to pull up the data you've got. Well that's a really obvious example of how information can get changed on you.

But it can be changed in subtle ways, of course. You can have people who hack the system or work from the inside and change data, in ways that maybe they're intending to do, maybe they're not intending to do. Maybe they're authorized to do it and it's part of their job. Maybe they're not authorized and they're doing it because they have some malicious or greedy motives that they want to want to make the records different.

What this means from a legal perspective is that (1) making sure that you have the right information and (2) making sure that you have all the information, is very difficult. Another point here being simply that information, as we know, proliferates and files can be copied all over the place.

So what we have is a situation where if, for example, I wanted to pull up, oh you know, a screen that displayed – let's pick on medical records because I like medical records and because we're trending much toward electronic medical records. If I'm trying to pull up a medical record today, an electronic medical record today, that's the same as the one that was pulled up yesterday on a patient, as I say, I'm pulling information from a lot of different places to make that happen. But at the same time I may also be copying that information to other places and at the same time the records of the transactions that made the screen come up are being created somewhere else – log files and things of that nature. There's all of a sudden sort of a smear of information around my network.

Julia Allen: Let's stay with medical health records – you've got a patient file, a tangible, physical patient file that you hold in your hand, and then you've got the equivalent that you've perhaps printed out or submitted as part of an e-discovery request. But clearly its genesis, as you said, comes from – is dynamically generated, comes from all different kinds of potential sources. So as you go into a legal proceeding, with those two different types of information, what are some of the challenges? What changes?

John Christiansen: Well we haven't seen a lot of this yet but it's coming. I have, in fact, been involved in some cases where electronic medical records have been changed. And it can really be not pretty. There were some really bad consequences, as you can imagine, when information, false information about a person's medical condition or health status is changed.

What's different is right now in history it's going to depend on how sophisticated the plaintiff is – in other words, whether the lawyers for the side that's asking for this information are sophisticated enough to make hard questions about that whether they want to. If they have a suspicion that information has been changed or is not correct and that matters to the case, a sophisticated plaintiff's lawyer is going to ask for basically all the information you've got that's relevant to the reliability of the record you produced.

So, for example, if that were me, asking for your records, I wouldn't accept your paper production. I wouldn't accept a printout of what was on the screen. I'd say, "That's fine. Now I want to know what your log files say about how that was created. I want to go back into the database and see what's been going on in that database. I want perhaps even to see how your applications function, to see whether or not – to have my experts tell us whether or not, in fact, they are producing accurate results; whether they're presenting the information in an appropriate way; whether they're changing the information.” If that matters to me, I know enough at this point to start asking some very hard questions.

Julia Allen: You're really questioning or attempting to determine all the processes, the technologies, all the different ways in which that information could have been generated, I assume to either build confidence or question the validity of the output. Correct?

John Christiansen: That's right. And actually I'm doing that for a very good legal reason, which is that the goal of all of this is ultimately to get the information that matters into court. And there are a couple of thresholds for that.

One is, is it sufficiently reliable so that it can even be brought in? And right now we're developing the standards for understanding that. We've had some challenges to breathalyzer information, which have been taken pretty seriously, and had some wins on how reliable that information actually is – wins for the plaintiff's side, that is, getting the information created by the breathalyzer thrown out of court. We've had some challenges to voting machines, which is rather interesting this year I know.

But the fact is we're beginning to see those sorts of challenges happen. And sometimes we're getting evidence actually excluded in court. Now once it's in, there's also the question of credibility. Why should I believe it? If I'm a judge, if I'm on the jury, why should I believe that information? And if it requires a detailed explanation of how it got to be there, and we can poke holes in why it is what it's supposed to be, or why it's not what it's supposed to be, maybe it doesn't get believed.

In criminal cases, as I 'm sure you know, we've got proof beyond a reasonable doubt that the person is guilty, which means that the evidence is held to a very high standard. That would be why the breathalyzer information, in particular, is subject to such challenge because if you can create a reasonable doubt about how he applications functioned, then you win, you get your client off, if you're the defense lawyer.

In civil cases, there is the preponderance of the evidence rule. Basically it says if I have slightly more weight behind my evidence, if I'm slightly more believable – if I'm 51% and you're 49; or even if I'm 50.01% and you're the balance, I win. Of course, that's up to the jury, or sometimes the judge, to decide whether there's the preponderance or not. This isn't a hard and fast, quantitative rule.

But what that means is that if I can really raise some doubts about the validity of the information out of your system, then perhaps I do get a win. Perhaps I can void that financial transaction because I can raise doubts about whether or not this was something that was created by a hacker, by some unauthorized person.

And I've actually been involved with and dealt with a couple of situations where there have been transfers out of financial institution accounts, that the account holder said, "No, wait a minute, I didn't authorize that." And the question is whether or not there are valid records for that authorization and could those records have been altered? So what that means is if it's been altered, perhaps the financial institution is liable to the account holder for the transfers. And some of these can be very large transfers as you can imagine.

Part 2: Be Prepared: Policy and Process

Julia Allen: So John, do you find that kind of in wrestling with this issue – and as you said the legal profession tends to look back and you need to establish a precedent and have some cases that really bring these questions to bear in some case law that helps establish some positions. But do you find in this particular case that if organizations have well-defined processes in place, or standards in place, or tried and true policies and procedures that describe how information is generated and collected and archived and destroyed, do you find that having some of that kind of foundational practice in place is helpful in establishing the validity and the reliability of the information?

John Christiansen: Oh I absolutely think so. One thing I really hate is throwing things together in emergency mode. And the scenario we'd be looking at is, say, you're a large organization, you've got complex systems, you're doing a lot of activities, and you get served with a summons and a complaint. Well, under the rules we've got in place today, your duty at the time that you know that you're about to be the victim of litigation, if you will – or if you're going to file suit – your duty is to preserve all relevant evidence in your system.

And as I say, it's sort of smeared around now for a lot of purposes. So if you don't know where that evidence is or might be, you have a very difficult, basically, forensic task for your IT team. It's going to take them a long time to figure out what's going on. They may not be able to do it entirely reliably, and at the very least it's going to be a horrible burden, and you're going to no doubt have a number of painful meetings while people try to figure out who's responsible for what and how you conduct these activities.

So I absolutely think that you need to have a policy in place. I have a policy I'd be happy to make available for listeners – I think, if we can post that – which really I modeled on the idea of security incident response because we have experience with that. But it's really the same – conceptually it's the same sort of problem. All of a sudden we have an event where we need to know where the information is in our systems, how to get hold of it, how to make sure that it isn't changed in inappropriate ways, and then how to get that information out and use it. The best precedent I've found so far is security incident response.

Now obviously, when you look at what I'm talking about, I'm also implying, very strongly – and I'll state it very strongly – that if you have a good, competent security policy infrastructure and people are actually working within that infrastructure to manage your systems, you'll already have a good idea where information is. You'll already have people who are responsible for and understand how the systems work. Therefore you will already be in a position where you can intervene and respond with more efficiency. It'll still be a burden but you can respond much more efficiently and with a much better sense that you are reliably pulling out what you need to do.

And that's very helpful also if the other side starts arguing about whether or not you produced all the data, or the validity of the data, if you're able to say, "We had an organized response. This is what we do. This is how we did it. Let's explain to you the processes, or let's explain to the judge." And a judge is going to be very much more inclined to say, "You know, they've put in a good-faith effort. They know what they're doing. So we'll just let them do it."

Because one thing that happens is that if the judge gets the impression that you're not cooperating or you're not managing the process competently, then they can actually penalize you for failing to produce the information the right way, for failing to preserve the information, and things like that.

Julia Allen: So it seems to me, based on what you're saying, that when you're faced with, as a business leader, when you're faced with litigation or these types of requests, that perhaps one of the areas of recourse that you have, when perhaps the discovery request seems particularly burdensome, you don't have sufficient resources, you may not be able to respond by the due-date – is it correct that maybe one of the areas of recourse you have is to describe how you go about producing the information and using that as a basis for negotiating both what gets turned over and by when?

John Christiansen: Yes, that's always a part of what you can do. And if you're in a position where you can say, "Look, we're working on this diligently," it helps a lot.

As we all know, in large, complex systems where lots of information is moved and used and stored, you can have lots of places where information gets backed up and stored and archived, and it can perhaps be a bit difficult to pull out, in particular, older records. You do have the opportunity to say, "Hey, look, this is an unduly burdensome request. Let's do something that works a bit better for us, that burdens us less, and yet and we'll still will reliably pull out the information."

It's going to be much more effective to be able to go in and say, "Look, we need to narrow this request, or maybe we need to figure out some other way of addressing this." If you're able to say, "We understand how our systems are configured. We work with them, we know where the information is, and we can tell you specifically why we have problems about it" as opposed to, "The sky is falling. This is all horrible. We don't know what we're doing and we just need relief."

No judge is going to look at somebody who's in a position where they're saying, "We just don't know how to do this and it's really hard. Could you give us relief?" They're not going to look on that kindly. But if you say, "We're working diligently, we've been working diligently, and this is just where we are. And it's not fair to force us to do more, because we're doing something that's very reasonable and in good faith," a judge is much more likely to look at that and say, "You're right, we'll give you some relief."

Part 3: Involve Key Roles; Practice Most Likely Scenarios

Julia Allen: So in addition to some of the things that we've been discussing, how do you advise your clients on actions that they can take, both in anticipation of, and also when faced with these kind of requests? What kinds of steps, or processes, or methods, or approaches do you recommend that they put in place?

John Christiansen: Well as I say, you sort of start with policy in this area, because you need to delineate responsibilities and accountability and have a sense of who actually does the various things that need to happen because it's a team effort. Frankly, the business leaders, probably the business leadership in an organization probably has a better sense than most of where major litigation is going to come from because you know what the interactions are. Your counsel, legal counsel, is also going to know that. And legal counsel will have a better grasp than anyone of the dynamics of what all this means in litigation itself and in the courtroom.

But then you absolutely need your IT team saying, "Okay, here's how we do things. And we are the people who are responsible for actually going forth and gathering information and identifying what we've got, and acting on preserving evidence."

And I think this is where your security folks also need to be part of the process. Because depending on how your organization does things – obviously, because everybody seems to do it differently – but the security folks have a particular expertise in knowing how to do forensics. And I think it can be very, very valuable in bringing that in and saying, "Okay, here's how we know what went on with that application, with that database, with our network. Here's how we can state what happened within this, with some degree of reliability." Or perhaps we have to say, "You know, that isn't reliable."

Of course, if they've been doing their job, and they've been listened to before then, you will have tightened your systems up considerably to your own help. But you start with policy.

What I would also strongly suggest is at the very least doing some tabletop exercises. Put together a straw man, a scenario where you say, "Okay, assuming that we had – that we were participating in a lawsuit, how would we react to that?" And in many organizations, in many industries in particular, you have the ability to say, "This is a recurring type of lawsuit that we see."

Let's pick on healthcare again. You're fairly constantly going to be sued for alleged medical malpractice, for medical errors, for – leaving aside the question of fault – because people frankly die and get injured in hospitals, one way or the other, and it usually winds up in a lawsuit. What you can and should do in a situation like that is say, "Okay, we know this kind of lawsuit happens. How do we respond to that?" And let's, at least, walk through it on the table top and see – do we understand what we're doing? Do we have gaps? Does everybody know their responsibilities?

In complex systems, and systems where it's perhaps more difficult to keep a handle on what's going on, where information might be smearing around, maybe you should actually run some exercises where you say, "Okay, let's pretend we've had an event. Let's pretend we're under the gun. Let's actually have some forensic activity in the system. Let's actually try to pull that information out. Let's see what we would get and let's see what we think of that."

I think where you know that you're going to have litigation going on, on a reasonably frequent basis, and you know you're going to have to pull information out, and you just see this coming down the pike at you – even if you haven't seen an electronic discovery request yet, or had to respond in depth because the lawyers for the other side haven't yet caught up with it, they will. And you'll get more sophisticated requests and you'll have to respond.

So it's worth (1) putting your policies in place; and (2) testing, to some degree, whether or not you actually can do what you believe you need to do.

Julia Allen: Yes, your saying about the tabletop exercises kind of reminds me that so many of these disciplines are related. But certainly in business continuity, disaster recovery, you mentioned incident management and response. There are certainly other established disciplines where this idea of an event happens that you need to respond to, where you can kind of draw knowledge and expertise and practice from those disciplines.

John Christiansen: That's exactly right. Business continuity is another great example. Because, of course, that's another one where something bad happens – and litigation is bad, as far as I'm concerned. Even though I'm a lawyer, I don't like litigation much as a way of dissolving disputes, but it's what we've got. But something bad happens. It calls into question whether you can work with your information systems appropriately, whether you can fulfill your mission, whatever/however that mission is defined.

And I guess I would have to say that one thing we're doing now is that we are adding to the mission of the IT Department the ability to identify and produce admissible evidence, in the form of electronic information from that system, and to support the organization in litigation, if it has to.

Julia Allen: Well John, this has been a very, I think, rich and helpful, valuable introduction to this complex topic that we all need to learn more about. Do you have some key resources that you use or would advise where our listeners can learn more?

John Christiansen: One of the best sets of resources, which has the great advantage of being both from a very authoritative source and free, is the Sedona Conference. This is an event that's been going on for a few years now with some of the leadership in the electronic evidence field. They get together for conferences but they also produce reports and guidelines around electronic evidence. And they've got – I think they probably have six or seven reports out there by now that, as I say, are free for download on the web.

And then talk to your lawyer, talk to your in-house counsel, talk to your principal law firm, whoever it is, and say, "What are we doing to address these issues?"

Because you don't want to get a blank stare. You want, frankly, your lawyers to be helping lead the charge, if they're up for it. If they're capable technologically, maybe they should be leading the charge though they need to do that with a strong team. But, as I say, you certainly don't want the blank stare. That's a bad sign and that means you need to really start riding herd on how that gets taken care of.

Julia Allen: Well John, I so appreciate your time and your expertise today, and the work you have done on behalf of our community with the Information Security and Compliance Risk Management Institute. So thank you so much for your time today.

John Christiansen: Thank you Julia.