Establishing Trust in Disconnected Environments

featuring Grace Lewis as Interviewed by Suzanne Miller

Suzanne Miller: Welcome to the SEI Podcast Series, a production of the Carnegie Mellon University Software Engineering Institute. The SEI is a federally funded research and development center operated by Carnegie Mellon University and funded by the U.S. Department of Defense. Today's podcast will be available on the SEI website at sei.cmu.edu/podcasts.

My name is Suzanne Miller. I am a principal researcher here at the SEI. Today, I am very pleased to introduce you, once again, to my friend Grace Lewis, who has done many kinds of research at the SEI. She has dealt with edge computing, with cloud computing, software architecture, service-oriented architecture, technology evaluation. I could go on.

Today, we are here to talk about trust, not trust between people in a normal way but a technology-based kind of trust. So, I am going to let Grace say a little bit about how she came to this research, why is this research important to you right now, and we will get into talking about it.

Grace Lewis: Thank you Susie. Like you said I have been dealing with what at the time had been emerging technologies, so cloud computing, service-oriented architecture. I like new stuff. As I was doing my research, I always felt that the tactical space was really not well served. Maybe it is not where the commercial market place is interested, but I always felt it was not well served. These people in the field need so much technology, need so much help because they operate in environments that are very difficult.

Suzanne: Impoverished.

Grace: They are impoverished. They do not have access to power. They don't have access to great bandwidth. They do not have access to just the normal things you would find back in the data center, right? In addition, there is stuff happening around them. If you look at a first responder, there could be buildings falling. In the military…Stuff happens, right? I really started to look into this area of research, and that led to what has probably been our biggest output, which is tactical cloudlets.

Tactical cloudlets are forward-deployed, discoverable computing nodes that are placed in proximity of mobile devices. What mobile devices do is that they discover them. They are able to use them either to offload computation. For example expensive computation does not execute on a mobile device but rather executes on the cloudlet, and also for data staging, meaning the cloudlet can serve as a collection point for data that is being gathered in the field.

Or, the other way around. It can store data that has either been pre-provisioned in advance to support a mission or maybe data that is being synced back with the enterprise as connectivity becomes available. So that is how I got to tactical cloudlets really trying to serve the people at the tactical edge.

Suzanne: What is the problem with tactical cloudlets and trust?

Grace: One of the key features of the tactical cloudlets, like I said, is that they are discoverable, right? So a mobile device in the field is going to say, *Are there any cloudlets around me*? From a mobile-device perspective, I want to make sure that cloudlet is a friendly one. It is a good cloudlet, right? And the other way around is also true. If I am a cloudlet and the mobile device says, *I would like to connect to you*, I need to know it is a good one.

Suzanne: In a first-responder situation one of the unfortunate things that happen is you get looters that follow around the first responders looking for riches. This is a case where the trustworthiness, a smart looter would discover whatever computing resources to find stuff and do things we do not want to happen. We want to be able to have some way of preventing that kind of situation from occurring.

Grace: Correct, and the same is true in the military.

Suzanne: So, what have you found?

Grace: What we did was a combination of things. First of all we created a threat model. We said, *What are the biggest threats in this type of environment*? Not surprisingly, the biggest threats are node impersonation and capture, exactly what you said. So what we did was, *OK, let us try to focus on what would be the extreme scenario*. So the extreme scenario is you have cloudlet that is fully disconnected from the enterprise, so there is no reach back to the enterprise.

Suzanne: I cannot verify your permissions or anything like that.

Grace: Nothing. So let us take the extreme scenario, and let's create a solution that will be very generic and not make so many assumptions about things like hardware or whether credentials were pre-provisioned in advance.

What we did was, going back to our four requirements, what we did is, *Let us create a solution that addresses four requirements*. The first one like I said before is, *I have no connectivity, so I*

cannot rely on going back to the enterprise to do any type of online credential validation or anything like that.

The second, *No hardware requirements* because in a lot of these situations a lot of organizations come together to support a mission, right? So you cannot place any requirements on hardware. Or, you should not place any requirements on hardware. So we said no…

Suzanne: *Only the IPhones over here. Only the Androids over there. You can't do that.*

Grace: No reliance on trusted platform modules, TPM, nothing like that. The third one is, *Because we cannot guarantee that cloudlets and mobile devices are going to be deployed at the same time by the same people, we said no credentials pre-provisioned on mobile devices,* so we are going really extreme. And the fourth one was whatever we created had to address the threats that we had developed in the threat model.

After a lot of investigation looking at different things, looking at, for example biometrics, behaviometrics, looking at what people do in mobile ad hoc networks. We ended up finding a solution that Stanford had created many, many years ago called Identity Based Encryption (IBE).

It [IBE] is very useful for these environments because it is very lightweight. It is called IBE because what it does, is that it does not really rely on certificates and things like that, but rather it says, *OK, your identifier, which is basically your public key is going to be any arbitrary string*. For example, in the context of secure email, which is what they had been using, your ID or your public key was your email address, which was unique.

For us it was great because we could use something, for example in our case, we just used the Android device ID as the public key. Based on that, it allows for a very lightweight solution, where you do not have to do a lot of certificate processing and things like that.

Suzanne: There is not a lot going on in the background.

Grace: Correct. The second thing was that in order to address node impersonation and capture, we had to combine that with secure key agreement without a trusted third party.

Suzanne: Secure key agreement without a trusted third party. So trust between the two of us.

Grace: Without having to rely on an external party. So basically what that means is using out-of-band channels, like for example using voice or using physical proximity or using something where we can establish that first encounter, that first channel of trust …

Suzanne: Without the device.

Grace: Yes, and then we get into this. So we developed what we call a 4-step solution. In the first part of the solution, we call the bootstrapping phase. So in the Bootstrapping phase,

basically the cloudlet sets itself up as a key generation center, meaning that it is going to generate credentials for itself and for any device that wants to connect to it.

The other thing that happens during the bootstrapping phase is that mission duration is set. That will become clearer when I talk about credential revocation, which basically means, this cloudlet is going to be available for the duration of this mission, which we are going to decide now is three hours.

Suzanne: And after that, everything we have done is now void.

Grace: That is what is going to happen at the end, so right. So that is the bootstrapping stage. The second stage is what we called device pairing. Basically what would happen is if I have a mobile device, and I am going to go out into the field and do something with this cloudlet, I have to physically come up to the admin of the cloudlet and say, *I would like to pair this device with your cloudlet*. After visual confirmation, which would be the out-of-band channel and physical proximity, because we implemented for Bluetooth and USB, what happens is that the cloudlet generates credentials for this device based on the device ID, which is only valid for that device, and sends me those credentials. Why? Because we established a secure channel based on visual confirmation. That is device pairing.

So now, I go out into the field, and I want to use the cloudlet. So the third step is WiFi authentication, which means, once I have my device and it has been paired and a cloudlet says. *I am here! I am here!* I do WiFi authentication, why? Because I want to know that the cloudlet is good, and the cloudlet wants to know that I am good. Because we already exchanged credentials we can do that very easily so we use something called RADIUS, which is a pretty well known WiFi authentication protocol to do that.

Once it is done, step four is API requests, now the mobile device can start interacting with the cloudlet in a very secure way. Now, in addition to that because node impersonation in capture are such high priority threats, we have to have a way to revoke credentials. The first one is, we call it automatic, which is when the mission time expires, when the duration expires. Done.

Suzanne: *We are done!*

Grace: Done! That device can no longer use those credentials. The other one is manual. So for example so if I know that the device has been captured, has been compromised, I can manually go into the cloudlet as the admin, and say, *I do not want this, this device can no longer connect*. So that is basically what we did.

Suzanne: This is fun!

Grace: It is fun.

Suzanne: What comes next with this research?

Grace: Right now we have two ongoing threads. The first one is…Something else that is very important in these types of environment is to be able to share information. We are not talking security now but just being able to share information. Sharing information in environments where you cannot guarantee that there is going to be connectivity between everybody, and not only everybody in the field but back to the enterprise. One thing we are working on now, we are running experiments, and it is looking really good, it is something called delay-tolerant data sharing. What we are doing with that is we are taking advantage of something called broadcatching.

Suzanne: Broadcatching, not broadcasting.

Grace: Broadcatching.

Suzanne: The opposite of broadcasting maybe?

Grace: Maybe, I do not know. Basically broadcatching, what it does is it combines BitTorrent, which is a delay-tolerant protocol, it combines it with RSS, which is what people usually use to subscribe to things. Basically it is a solution based on that, that enables people in the field to say, *I have a file, and I am going to tag this file with whatever tags. It then uses RSS to make that announcement*. Then people on the other hand will say, *If there is ever any file that is published with this tag, image, map…*

Suzanne: *I want it.*

Grace: *I want it*. It uses RSS for that. Once RSS is used to do that then what it does is it uses BitTorrent, which is a delay-tolerant protocol to get that file assembled, and there is your file.

Suzanne: And so the delay tolerance is basically *I wait until there is enough connectivity, I wait until there is enough power…*

Grace: BitTorrent does that.

Suzanne: That is what I am saying, that is what that does. It gives you that ability to not know if you are going to have connectivity. It takes care of that for you in the middle.

Grace: BitTorrent is great, and the results are really good.

The other thing that we are doing is, we are recognizing and people are recognizing, Internet of Things [IoT]. It is here, and there is a lot going on. So we are extending that concept of cloudlets. Where we always thought of cloudlets as, *There is the enterprise. There is the cloudlets as an intermediary layer, and there's devices. Well, there is another fourth layer out here, which are sensors and actuators, IoT.*

How do you incorporate those IoT elements into this whole system considering that IoT devices are even more constrained. I mean, some of them do not even have a user interface.

Suzanne: And some of them do not have any security either.

Grace: Do not have any security, do not have any user interface, do not have a lot of processing power.

Suzanne: Yes. Our phones, which is the most common peripheral device today, actually has huge amount of processing.

Grace: We are starting to work in that area.

Suzanne: But sensors are a lot, more diminished, more impoverished. I guess that is my word of the day.

Grace: Yes, it is.

Suzanne: Excellent. Excellent stuff. So transition, what are some of the transition activities, since that is something we worried about at the SEI, is not just doing good research but actually getting it out to people who need it.

Grace: In addition to publications, which are available from my website, we have released the tactical cloudlet software as open source. It is available on GitHub, and the address is github.com/SEI-AMS/pycloud, p-y-c-l-o-u-d.

Suzanne: Ok, so people can actually experiment with this on their own and the publications will give them instructions.

Grace: No, it is actually well documented, believe it or not, it is very well documented. Our software is very well documented. The GitHub report has a wiki that comes with it and all the instructions…

Suzanne: Good, so you get all that.

Grace: And people are starting using it, and it is great the community that forms itself when you have an open source product because people have told us. *Oh, found a bug*!

Suzanne: *Oh. We got to fix that!*

Grace: Or people are like, *I want to use this way or want to use it the other way*! And so we have created, I am not going to say that we have tons of followers, but we have created a little bit of a following of people that are using it and are very interested in what we are doing.

Suzanne: What is the most unusual or unexpected use of this tactical cloudlet idea that has come up as a result of this community?

Grace: You know what? The usage, I think it is really…

Suzanne: It is consistent with…

Grace: It is consistent. I would say it is one of two things. One is kind of like the more traditional, because this tactical cloudlet comes out of cyber foraging that idea about offloading computation.

So one use case is really, *I have to do speech recognition. If I do this, or object recognition, if I do this on my mobile phone, I am going to drain the battery in two objects.* So being able to locate a cloudlet and say, *Hey, I just took this picture can you recognize the face for me or this object for me*? So computation offload, is one type of scenario, that is one use case.

Suzanne: I am envisioning a new kind of geo caching game.

Grace: Of course you are.

Suzanne: That takes advantage of this. I do not know why…

Grace: Of course you are.

Suzanne: All right gamers out there.

Grace: And the other one is more for like what I said before, like data staging. So, for example being able to say, *I am going to go out on a mission, and I am going to need maps. I am going to need this capability, that capability, this software and being able to just pack everything inside your cloudlet and just kind of like, take it with you*.

Suzanne: Or, *I am a marathon runner, and I need maps of the area, I need to know where the all the rests stops are and all the rest of it.* I am envisioning some other edge cases.

Grace: Not that I have done it, but another big use of this technology is in healthcare. There is a wonderful project out there, I do not remember the name right now, where basically what they are doing with this type of technology is using it for Dengue detection.

Suzanne: Dengue fever.

Grace: Dengue fever, exactly. Where basically what they do is that they take a blood sample, and they use these stickers that turn different colors based on the blood composition at the moment. If you take a picture of that, it can be processed by a cloudlet, and now somebody can tell if a person has Dengue in seconds. Whereas before, they used to have to take the samples, go

to a lab who knows how far away? I mean things like that, that can help you do processing just on the fly are great.

Suzanne: You can be very proud. There is going to be lot of things based on this that nobody even thinks about because they do not have a way of getting that edge power.

Grace: So definitely, healthcare, education.

Suzanne: You are going to have to come back and tell us how cool things are happing in the next couple of years.

Grace: Would love to.

Suzanne: I want to thank you as always. Every time I talk to you, it is just the coolest stuff. You do like the emerging technology. I will verify that.

Grace: I do.

Suzanne: So, that is wonderful, thanks for coming, I know you have a blog post on this topic.

Grace: Yes, I do.

Suzanne: As well as we have got the GitHub, so lots of places people can go to find that.

Grace: Yes, we just recently published the paper where all this is described.

Suzanne: This is all good. We will provide links in the podcast to all these kinds of resources that you have provided for us. So everybody will be able to have some fun with this, I might even have some of my own fun.

Grace: There you go.

Suzanne: This podcast, as we all know, will be available on the SEI website at sei.cmu.edu/podcasts and it will also be available on the SEI YouTube channel and on Carnegie Mellon University's iTunes U site. As always, if you have any questions please do not hesitate to contact us at info@sei.cmu.edu. Thanks for watching.