How to Speak an Authentication Secret Securely from an Eavesdropper

(Transcript of Discussion)

Lawrence O'Gorman

Avaya Labs

Matt Blaze: The model is assuming only one side of the channel will be used liked that?

Reply: Yes, I should reiterate that, because that's very important. The attacker model is, the eavesdropper hears one side, Brutus can attack from the other side, and these guys can collude, but they can't hear both the challenge and the response.

Frank Stajano: What happens if the nurse gets the numbers wrong, does she as many tries as she wants?

Reply: There's a trade-off, but the answer is no, she doesn't get as many tries as she wants.

Ross Anderson: It's worth pointing that if the nurse can only memorise four colours, then for your trust scenario, you're better off giving two digit challenge response rather than four digit challenge response, because you get more entries that way.

Reply: Yes that's interesting, good point.

So we started off with merely four colour substitutions, that was great because the healthcare worker is saying four colour/number substitutions, you know, has a little bit of the load of memorisation, but it's not terrible.

James Malcolm: Couldn't you get the healthcare worker to add the randomisation? It may be not very good, but...

Reply: We could, and we can talk about that.

Matt Blaze: Sorry, maybe I missed something obvious, why do the camouflage elements have to not include the real numbers?

Bruce Christianson: Independence.

Reply: Why do they, because of the histogram attack. There's two scenarios here.

Matt Blaze: Here you never include the elements, but I could imagine a middle ground in which you statistically reduce the incidence of the camouflage elements

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 230–236, 2009.

in the random string of the ones that match the real elements, so that you flatten the histogram, but that still doesn't require you to put everything at the end.

Reply: Don't say anything more, because you're giving away the end of the talk. It took me a long time to figure this out so I'm really mad that you figured it out in a few seconds; I guess everyone else did here, are there any other questions with regards to that?

Bruce Christianson: If you say, we're going to have at least six elements at the end of the colour sequence, and Brutus gets to know you add independence in that way, is it then the case that Brutus gets no information at all?

Reply: Well it depends on the question here. If they get the entire sequence and are allowed to keep on going, then by virtue of the information coming back to them, that, no, this is not authenticated, then they know that something is wrong here. They don't know which colour is wrong, but you're right, there is information that they can get.

Bruce Christianson: Provided Brutus hears at least six numbers after the last colour, he can't tell that he's not in the best case.

Reply: Right, right, now the attacker knows exactly the method that we're using here, and so they'll know this method that, I'll talk about next. We can call this the Blaze solution here, we're going to just do both, adding the dependent camouflage elements which defends against Eve, and then we add independent camouflage elements. Picking a number out of the air, let's try to do this with two thirds probability, that I will have at least three of the four colour/number substitutions in my added camouflage elements. To do that, probabilistically speaking, I have to add ten extra camouflage elements.

Frank Stajano: It looks to me that the main problem you have with these attacks is that Brutus gets to try again and again. You have a very abstract way of framing the problem here where the input from the speaker is just numbers, as if they were punched in a keyboard, instead they are spoken. You obviously need to have a speaker independent number recognition system in there, but you could also use some kind of speaker dependent voice print matching and say, I will behave differently whether I believe that this is Sue, the nurse, or whether it's someone who doesn't sound anything like Sue, maybe because they've got bigger voice, or something like that, so in that case you are going to react differently to whether you let them try again or not.

Reply: I think you're saying two things there but let me comment on one. At the very beginning I said that I'm just throwing biometrics out the window, so we work in a noisy environment, this is a noisy application here and so speaker verification, I'm going to say, doesn't work. I think what you're saying is you can combine it, and maybe this will be a help.

Frank Stajano: Yes, you would be not reliable to do the speaker verification just from the voiceprint, but you use it as a hint to say, how am I going to react when I get a bad match on the digitised links.

Reply: Yes, that's a good comment.

Marios Andreou: How about if the mapping from number to colour is not uniform for every worker, so each person has their own, that's one thought. The other thought is, some of the online banks use a method where when you register you give them the pass phrase, and then they challenge you with three or four random characters from the pass phrase, so they give you 1, 7 and 10 and 12 from the pass phrase. I don't know if that's any more secure than what you've done.

Reply: But it doesn't matter what the scheme is, the eavesdropper is going to hear particular characters each time, and they're going to do this histogram attack.

Bruce Christianson: I assume different employees do have different mappings?

Reply: Oh absolutely, yes.

Tuomas Aura: But in the medical model you show it doesn't help. Having different substitutions for different employees, because an attacker who is listening to only the response, doesn't get any information.

Reply: That's true.

Tuomas Aura: So then, in that attacker model, it doesn't help to have different.

Reply: From the eavesdropper point of view, yes. But if someone gives their password away, you don't want everyone's password to be given away.

Frank Stajano: Explain what memorising voice commands means in this context.

Reply: There are about seven commands that you speak to the voice agent, and you say, connect to Larry O'Gorman, or, connect to the nearest cardiologist, and you can say, get my messages, and, Weaver take a break, Weaver's the name of the voice agent.

Frank Stajano: So when you have an attack, when you discover an attack, you are sure that there is a Brutus in there who's trying things out?

Reply: Well after a certain number of erroneous things you say that, you're guessing.

Frank Stajano: When this happens, what do you do, you have to give them another mapping of colours, surely that's going to annoy the nurses no end?

Reply: Yes. This is obviously the trade-off. There's always trade-offs with security, and, yes, this is one of the things that I'm suggesting. Right now we're very liberal in what we do in the Johns Hopkins test, the security is not high. Well, you know, they turned it off, so it's zero now, but even as it was designed, we didn't ask a whole lot from them.

Frank Stajano: Could you think of a way of, for example, making the mapping not overlap with the previous one, because if you used to have like blue = 7, and green = 9, and then now it's green = 3, and blue = 4, it's certainly going to be mixed up?

Reply: Well that's like a chain hash table, or something, and you can do a chain hash table with a piece of paper.

Frank Stajano: All that I'm suggesting is instead of using colours, the next time you use animals, or shapes, or something cognitively different.

Reply: Well they have to memorise more things, you're asking a person to memorise, to do colour first time, animals second time, and planets.

Frank Stajano: Yes, that's better than colours first time, and different colours the second time, and different colours the third time.

Bruce Christianson: But at least they won't get confused between the second phrase and the first.

Reply: Well that's a good point, yes.

Bruce Christianson: The problem seems to me to be that colours are different from numbers. What if you ask people to remember that certain numbers were not themselves, but the rest were. So four numbers are permuted, all the rest are themselves, and then the challenge is the numbers 0-9 in a random order. That seems to get you both the properties you want.

Reply: Exactly, so the suggestion here is that, I'm giving away information by having a challenge which has numbers and colours in it, because Brutus immediately says, well, the things that aren't colours are numbers, and so I know now what the colour substitution is. So instead of that, having all colours there, but the person knowing that purple is not a colour that they've memorised, the onus is then on the user to say some random number.

Bruce Christianson: I wasn't advocating that, but yes, it could be done.

Ross Anderson: I think what Bruce was suggesting was that you would memorise, for example, 9 mapped to 8, so when the number's 9 you say 8, but 7 maps to nothing, so when 7 is spoken you say a random 2.

Reply: Right, but you don't have to say anything actually.

Ross Anderson: No, you must say a random number.

Bruce Christianson: People are bad at picking random numbers, Eve will get the hint. Instead say there's a permutation of four numbers, all the other numbers map to themselves. Eve will therefore hear the numbers 0 to 9 in a random order, Eve gets no information.

Reply: Yes, so it's just the Eve defence that we still have, and because we have all the numbers there in the challenge we don't have to worry about Brutus anymore.

Ross Anderson: In the threat model where there's one sided of eavesdropping only, that's true. I've another concern with this though, which is completely separate from the threat model. Passwords were great so long as people only ever logged onto the PDP11 in the corner of the lab, but once people had to log onto 101 websites, passwords break badly. Now if your mechanism ever becomes widespread, I suspect it will break badly for the same reason.

Reply: Yes, maybe true.

Bruce Christianson: Or else you have to say that your authentication is always local.

Reply: Right, right, so you're asking not use the same colour/number substitution at computer B, or my second nursing job that I have.

Ross Anderson: Because there's a Mafia controlled part.

Kenny Paterson: If I understand things correctly, users in your system have to produce two kinds of responses. Sometimes they have to repeat a digit, and sometimes they have to translate a colour into a digit? Is there any evidence that users maybe hesitate slightly more when translating?

Reply: That's a good point, and so what we do is, we just jigger the challenges a little bit, so we randomise the time of challenge. But actually the users, when they hear a number have to think, oh that's a number, and they repeat the number, when they hear the colour, they think, oh that's a substitution, and so it takes about the same time in my little experience.

Matt Blaze: So what they're authenticating into is an information system that they're doing independent of their immediate healthcare duties, that is, this isn't, turn on the defibrillator, this is, log me in so that I can get patient records subsequently.

Reply: Yes, I don't know if I'd call that secondary, I mean, they're not speaking to their mum, but, yes, it's not what they're doing with their hands right now.

Matt Blaze: I remember reading Ross' paper about healthcare security, and he pointed out the important issue in healthcare is not confidentiality, but the accountability. And you know, I didn't believe him at all when he first said it, but after thinking about it, I believed that it was obviously true.

Reply: Yes, it's a legal aspect, yes.

Matt Blaze: Finding out who accessed records, is much more important that you preventing someone from getting unauthorised access, because of the emergency access aspects of this. So how do you deal with the problem of an impatient cardiologist wants to find out something right now, and the authentication just doesn't work.

Reply: Well, the answer to that, is always that's a trade-off.

Matt Blaze: Is there a policy built into this system that allows for such breaks, there's something you can do?

Reply: Yes, axe to get through the emergency backdoor.

Matt Blaze: I imagine that's an axe that gets used quite a lot in systems that have annoying properties.

Reply: Well, annoying security properties. This headphone is logged into at the very beginning of the day when hopefully you haven't gotten into your emergency, you wear it all day, you don't have to authenticate each time.

Dan Cvrcek: You could ask them to add some numbers together, instead of repeating the same numbers they are hearing.

Reply: Yes, you could do that, and then that would be random. But then you've asked them to do three things, you've asked to memorise, you've asked them to do the authentication, and you've asked them to do a little bit of maths.

Reply: It's easy for us, you have more confidence in humankind than I do, but that could be the right thing rather than this 20 seconds.

Mike Bond: Alternative technology which might work very well, although it would probably have lower security in terms of total number of permutations, is to use Steganographic channels to do the authentication. George and I did some work on looking at authentication protocols for secret societies, how to design a funny handshake, or a funny walk, that identifies you as a member of a secret society[1]. The way you could apply it here is to have, say, some secret object in the reception, like a vase with flowers, which you put in a different colour every day, and the query to the nurses is, what's the colour, and anybody who listens to that query doesn't know what the object is that you're supposed to say the colour of, but the nurses all know that it's the vase of flowers, or they know that it's the children's toy in the toy rack that's been put on the top of a wardrobe.

Bruce Christianson: That only authenticates whether someone's a nurse or not, it doesn't tell you whether they are the correct one.

Mike Bond: No, but maybe if you start working on that, and you figure out ways to give people different objects, and you figure out ways to alternate the changing of the objects, the nice thing about that is that it uses your associative memory, everybody remembers the key is something in the world around them, and it's something that's really easy to remember, and even quite easy to change. Compared with something abstract, this key is very real, and all you've got to do is the challenge on attributes then.

Frank Stajano: So long as the receptionist doesn't go and rearrange the flowers.

Aaron Coble: I was thinking you have these authentications that are proofs of specific segments in time. If you distributed it out, throughout the day, with just one piece of information being asked, you're working against the eavesdropper

[1] Bond and Danezis, "The Dining Freemasons", 2005, LNCS 4631, pp258-265.

because they no longer know when you're authenticating and when you're not, if you choose your responses carefully.

Reply: Another nice thing about that is it gives you the ability to have different levels of security, logging on, versus going to the morphine cabinet, you know, then you answer another challenge.

Aaron Coble: And a minute and a half sounds like a long time, but if you're being asked one question an hour, it seems...

Reply: Yes, then it seems to be a lot less, even though it adds up to a minute and a half, yes.