The Man-in-the-Middle Defence
(Transcript of Discussion)

Ross Anderson

Computer Laboratory, University of Cambridge

The man-in-the-middle defence is all about rehabilitating Charlie. For 20 years we've worried about this guy in the middle, Charlie, who's forever intercalating himself into the communications between Alice and Bob, and people have been very judgemental about poor Charlie, saying that Charlie is a wicked person. Well, we're not entirely convinced.

Once you start to consider dishonest insiders, you have to change the whole threat model, and the rather simplistic assumptions and models that people used in the 1970s and early 1980s just don't run any more. Now it might occur to you that E-commerce is perhaps in the same state that protocol development was 20 years ago in the mid 1980s. E-commerce assumes that you've got three parties, the customer, the merchant, and the bank, but in real life the threat is usually not Charlie, the long eared individual lurking in the shadows. The threat is Bob, the shop that you go into and do business with – because the shopkeeper is bent. Sometimes it's the bank who's bent. How do you deal with this?

The card you may think is your card, but it isn't; the bank's terms and conditions say very clearly that it belongs to the bank, and even if belonged to you, you've got no say over the software in it. It would cost you quite a lot of money to dig that software out of the card and disassemble it, and find out what it actually does. So you're not represented via the card, the issuing bank is. Similarly the terminal represents the merchant's interests, or perhaps – if we're being slightly more careful – the interests of the merchant's acquiring bank. No-one represents the customer. But thanks to Mike Bond, and earlier work by Jolyon Clulow, Sergei Skorobogatov, and so on[1], pulling this protocol apart, we now know that we can fix that.

Meet my attorney. This attorney is a version zero prototype. The attorney that you will eventually be able to get, if someone takes this up and runs with it as a business, will of course be a lot smaller and slicker. It will be a dinky little device about this size, and it will have on it an LCD display and a little red button, and it will have the male part of a smartcard – which you can see here – and it will have the female part – which is the socket here. You put your chip and PIN card into the attorney socket, and you put the attorney into the terminal at Texaco, or whatever, and the LCD display says £34.99, or whatever is actually being asked from your card. Never mind the display on the merchant's terminal, that can say what it likes; but we don't care any more because now

[1] See position paper.

B. Christianson et al. (Eds.): Security Protocols 2006, LNCS 5087, pp. 157–163, 2009.

we have got a trusted path and a trusted display. So I put my card, with its little card condom – which is the phrase that Markus came up with for it – into the merchant terminal; and I see the amount; and I say, yes, £34.99, that's the amount of petrol I took at the pump; and I press the little red button; and then the device, which knows my PIN, enters the PIN into the card, thereby authorising the transaction, and also writes an audit trail entry. Whatever PIN I enter into the merchant terminal doesn't matter, four zeros, whatever, those bits are just thrown on the floor.

OK, so what happens then? Well, phantoms appear in my account, and I go to the bank and I say, "Oi, I didn't make those!" and they say, "You must have done because our system's secure" and I say, "No, hang on a minute, my system's even more secure than yours because I hired this smart young man, Peter Ryan, from Newcastle University, to formally verify it all, and it's EAL6+[2], right, and compare that with your twenty million lines of spaghetti COBOL? Forget it, pal!"

Tuomas Aura: What if someone steals your wallet, the wallet contains both the smartcard and your attorney, and nothing is required to use this because it is sufficient to press a button.

Reply: You will get a premium device with a thumbprint reader attached to it, so you will have to swipe your thumbprint, and it will check that you are alive before it will do the transaction. Once you can start using the middleman, not just as an attack but as a defence, you really have an opportunity to start fixing a lot of the stuff that's wrong with protocols. That's the fundamental point here. You want biometrics? You can have it, with the middleman. You're no longer locked down by the contracts that twenty or thirty thousand banks have made with each other, with VISA, with Mastercard, and with the suppliers. That is just too much momentum, you can't do any upgrades or maintenance on that kind of security.

Kenny Paterson: This is wonderful Ross, but why would a merchant or bank put your device in his system?

Reply: Well, the merchant doesn't mind; the merchant just cares about being paid. The banks might object because it undermines their deniability, but they might have problems on the public relations front if they tried to object publicly.

Bruce Christianson: But presumably this is not something under my personal control. The attorney must be provided by some third party?

Mike Bond: Where do you find a good lawyer, is the question!

Reply: Lawyers get their money from their clients, and so the lawyers have their incentives properly aligned. As EMV fraud transaction padding, redirection, and

[2] ISO/IEC 15408 Evaluation Assurance Level.

all the rest of it begin to climb, there will quite possibly be a market for people to build devices like this. Certainly, we know the electronics industry in China would be quite capable of producing these devices for a few pounds each: all you need is a small microcontroller and a relatively small and simple piece of software. The means for verifying small and simple pieces of software have been worked on for forty years.

Bruce Christianson: Oh sure, what I'm asking about is the threat model. The attorney isn't just your agent, he's not just a part of your end-system, he really is a man in the middle.

Reply: Absolutely, and if you have bent attorneys that can cause serious problems. Then there's the question of what happens when the merchant has his attorney, and you have your attorney, and the bank has his attorney, and you've got a whole room full of lawyers, and that's something I'm going to come on to shortly.

Bruce Christianson: There comes a point where the bank is going to say, well actually the fact that my customer had the benefit of an attorney strengthens my position. The question is, why is the bank going to allow, or encourage you, to plug an attorney into the card. Now I can see certain assumptions on the attorney, under which the banks are going to be only too eager for you to do that.

Reply: Well the banks cannot, as a physical matter, stop you plugging an attorney in, because of the way they designed the protocol. The only way they can stop you using an attorney is if they move from static data authentication to dynamic data authentication, but that would cost them money.

Mike Bond: There are still ways to use an attorney if you look inside the protocols. If they change the protocols enough then they can write it out, but even with dynamic data authentication you can still, for instance, send a PIN to the card. The card would have a private key on it, so you can send the correct PIN to the card, but you cannot piggyback, you can't decrypt things coming from the terminal. So for instance, if you wanted you could use your attorney to give you a one-time PIN, which you type at the terminal that doesn't encrypt it, the attorney could check it and pass on the real PIN. So there's a range of different attorneys, some of them would work with dynamic data authentication, some of them wouldn't.

Audience: But that doesn't stop the banks putting legal things in place to stop you.

Reply: That is why this is not just a technical paper, but also it's a challenge to the banking community: guys, are you honest or not? If you are dishonest, object to this paper and say that you will not let your customers use attorneys,

if you are honest then say, that's fine. So please Mr NatWest, tell us, are you or are you not a crook? May I or may I not use an attorney when I am doing business with one of your point-of-sale terminals?

Frank Stajano: Maybe this is a silly question, but I can't understand why the terminal doesn't notice that you're not using the PIN from its own keypad.

Reply: Because that's the way the protocol is designed. The terminal sends the PIN to the card, and either gets back a MAC or it doesn't. So if my middleperson device, my attorney, takes the customer-entered PIN from the terminal and throws the bits to the wind, and then puts in its own PIN, the terminal has no way of knowing.

Matt Blaze: So you're basically adapting EMV to the Chaum Wallet model[3], where you have a device that represents the customer's interest acting as a wrapper around a device that represents the bank's interests?

Reply: We're aware of Chaum's work, but he's actually doing something rather different for his protocol.

Matt Blaze: Well fair enough, the implementation is different, I'm talking about the model, right, and essentially what you're doing is observing that EMV is susceptible to Chaum's model.

Reply: Chaum is representing only a privacy interest, and we're providing trusted path, trusted display, and audit.

Chris Mitchell: I think the banks would have a legitimate concern about you putting your device in the merchant terminal, because although you may not think it's a very secure scheme, to some extent the security relies on the merchant checking that the card is a genuine card. A malicious card – or a malicious sleeve – could generate a false MAC and give it to the terminal, and if the terminal doesn't do an on-line authorisation, it doesn't have any way of knowing whether the MAC provided for the transaction receipt is a genuine MAC or just a random string of bits.

Reply: So the local small businesses are all compelled to go and do an on-line check.

Chris Mitchell: But it does violate part of the security model the banks have.

Reply: But that's basically a niggle, that's the sort of thing that the bank would use as an excuse.

[3] Wallet Databases with Observers, Crypto '92.

Chris Mitchell: Well they would say that, but you would say that too.

Mike Bond: I can think of a compromise here which is, devices with card-retained functionality – that's where the card goes all the way in – you need quite an expensive, complicated attorney, with a short-range radio link. And even then if they choose to desire a card retained in such a way, they can make it very difficult for you to make a neat attorney with a screen and keyboard that you can still physically access whilst it's reading the card. So those things would make it difficult to use your attorney. Anything where you partially insert the card into a slot, such as shops and supermarkets, you probably could use the attorney.

Steven Murdoch: There are a lot of cases where the merchant never actually sees the card, never holds the card. If you look at the train station, the terminal PIN pad is too far away, it's behind a glass screen.

Mike Bond: The culture's changing in electronic payment now in the UK in that the merchant doesn't want to see your card any more. You do business with the box, and then the merchant just looks at the result from the box and says yes or no. You do business with something that's on the merchant's table but which belongs to the bank, and you don't really do business with the merchant any more. He just looks, says yes or no, and then gives you the goods.

Alex Shafarenko: But how are you going to prove to the bank that your attorney is not a spoof itself?

Reply: I don't have to prove anything to the bank. It's none of the bank's business.

Alex Shafarenko: Well, do you participate in the transaction if you're suspicious of me?

Reply: Well, for example, I refused to participate in an EMV transaction at Tesco's because when I put my card into the chip and PIN terminal, the girl snatched it away and said, "No, you mustn't do that, I must do this." Then she swiped my card through a mag stripe reader into a chip and PIN dock. Right, that means that Mr Tesco has got my mag stripe data and my PIN, and it went in clear text through computers that he programs. I said, "No, I am not putting my PIN into this terminal," and I paid cash instead.

So the attorney is not there to prove anything to the bank. It's there to prove to the judge – a big difference. I don't care about the bank.

Kenny Paterson: Since we're all having such fun do you mind if I sketch a phishing attack on your system. So the phisher sends out a whole lot of bogus sleeves containing a dodgy program?

Reply: Oh absolutely, then you're dead.

Kenny Paterson: Yes, so that does lead to the observation that this only works if you can trust the supplier of your device.

Reply: Absolutely, so you must initiate the purchase. You must go down to RadioShack, or whatever, and say, I want a sleeve that carries a guarantee, and the people who sell this must actually sell it with a guarantee, you know, "We have got EAL3+," or whatever they reckon is enough certification for it.

Kenny Paterson: Then my recourse is to the insurer.

Reply: And presumably you sell this as part of a package of fraud assurance mechanisms, whereby you say, "If you use our device and evil happens to you, then we will provide you with ten thousand pounds' worth of legal assistance to sue the bank."

Kenny Paterson: I'm sure we'd have great fun cooking up some scare stories.

Bruce Christianson: This is not unlike the old version of a guardian. If you've got a piece of hardware that belongs to someone else, and you're required to attach it to a network in order to do a transaction, you don't attach it directly, you put a firewall on a bridge, which you control.

Reply: Well exactly, this is what we come to with this slide.

Bruce Christianson: As an extension of that principle, it's very sound.

Reply: Right, mail guards, anti-virus software, firewalls, and so on, we've all got...

Matt Blaze: When I said this, you yelled at me. [laughter]

Reply: I thought you were in effect asserting that I'm infringing David Chaum's patent, and I'm saying, it's insufficiently similar to constitute an infringement.

Matt Blaze: Ah, I actually said it's similar to the model.

Reply: So there's the dining attorneys sitting round the table objecting to each other's schemes, and this is what you can use in a firewalling environment.

Also, it's the place to put the human back into the protocol, which is the theme of this workshop. Too often stuff is designed so that the person who owns the kit is designed out of the loop. Designing the bank customer out of the loop is absolutely classical of the way that techies do stuff: the bank is paying you your consultancy money, so you think about the bank's interest. Every banker knows that people are more likely to change their spouses than to change their

bankers, so who cares about the customer? But all of this adds up; eventually it tips. We've got to have ways of putting the human back into the protocol, as well as getting the bugs out.

So what's rehabilitating Charlie about? It's about usability, and it's about maintainability, which I reckon are likely to be the two big problems in security engineering over the next five years. Voilà.

Tuomas Aura: I just want to point out this attorney that wants to give the PIN directly to your credit card, does not communicate with any other terminals, or any other middlemen?

Reply: He shouldn't, because if you have him in a WAN telling your PIN to the Mafia attorney, and the KGB attorney, and the North Korean attorney, and so on, that's bad stuff.

Tuomas Aura: So it may well be that you cannot compose this attorney, because all of them have special requirements that are not compatible.

Reply: Exactly so, I know. But my point is that the Maurer-Massey[4] insight gives us a conceptual framework in which you can explore when things should commute, and when they definitely shouldn't.

Mike Bond: I think what Tuomas is getting at is, what modifications of the EMV protocol are feasible, and which ones aren't. Is this really a protocol which can be hammered into a new shape, rather than being redesigned, and even deployed. And the answer is, with static data authentication, you would have a very hard time making multiple middlemen work usefully. Maybe with dynamic data authentication, where every player has a private key, it's a bit more feasible, but you still have the denial of service problem.

Reply: And then you've got the terrible problem of knowing which is the genuine terminal. You may be able to put your PIN in, encrypted under a terminal's public key, but one of the middlemen says, hello, I'm a terminal, and another says, hello, I'm a terminal. The Mafia own a bank so they can certify a terminal, so their certificate looks as good to the EMV system as the genuine terminals.

Bruce Christianson: In a weird kind of way things get better once the EMV protocols are redesigned with middlemen in mind.

Reply: Embrace and extend.

[4] Cascade Ciphers: The Importance of Being First, ISIT '90.