Getting in Front of Social Engineering
Transcript

Part 1: Why We're Susceptible

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 executive outreach. Today I'm pleased to introduce Gary Hinson, an IT governance specialist with expertise in building security awareness. Gary manages two popular websites: NoticeBored.com and ISO27001security.com. Today we'll be discussing social engineering and actions that business leaders can take to mitigate risks associated with social engineering. So welcome Gary. Glad to have you with us today.

Gary Hinson: Thanks very much Julia.

Julia Allen: So for starters, for our listeners, what is social engineering, and in your experience, why is it so difficult to defend against?

Gary Hinson: I think one of the best definitions of social engineering that I've seen comes from Kevin Mitnick that hopefully some of your listeners will have heard of. He said, "Social engineering is a form of hacking that relies on influencing, deceiving, or psychologically manipulating unwitting people to comply with a request." That's quite an open-ended definition. The request could be anything really, but he's using it in the context of hacking. So people who would generally ring up a target company, try and get hold of people, try and get information out of them as part of a hacking attempt.

And it's difficult to defend against because you don't really know the nature of those kind of calls. If you're expecting calls from customers and suppliers and partners and employees, particularly say if you're on a call desk, trying to spot a social engineer calling in, it's not obvious what kind of calls would be a social engineering type call. So the people need to be pretty alert and pretty aware of the possibility that they're being milked for information. Unfortunately they're asked for information all the time from the normal callers, so it's a really tricky area to distinguish between social engineering calls and legitimate business calls.

Julia Allen: Right, because as human beings, we all tend to be quite trusting. If someone's calling us up and they appear to have legitimate purpose, or they appear to know something about us that only those that we should trust would know, then as you said, it is difficult to discern if it's a legitimate call or not, right?

Gary Hinson: Absolutely. That's part of the game they're playing. Social engineers use various tricks and various bits of information to appear more credible. Quite often, for example, they will try and appear like an insider. They will try and make out that they're an employee at a distant office, or an employee in the local office who's just arrived, so that the voice wouldn't be recognized. They use little bits of information that originally they might obtain from the web or from published sources like marketing brochures, that kind of thing, to make it appear that they know something about what they're talking about.

That's really the hook. Once they get the person on the other end to believe them, they can start to get a little bit more information out of them. Even if that call doesn't work out, they go away with a bit more information than they started with, and so they go on to the next call.

Julia Allen: Gary, to what extent do you think social engineering plays a part in successful security breaches? There are all kinds of security breaches, many of which are quite automated. What have you seen, how social engineering might be combined with other forms of attack to make a successful breach?

Gary Hinson: My feeling is that most of the directed attacks probably involve social engineering. If you're looking at automated attacks, such as worms and viruses going out to everyone in the world, those are non-specific, non-directed attacks and they're usually just automated.

But if somebody's got their eye on a particular company or a particular person, then social engineering is a more obvious way to go. The idea there is to get some inside information, get inside the company, or inside the person's computer systems to find out more about them. And so that's the kind of situation where they'll use social engineering attacks.

A classic example really is phishing, where you get an email saying, "Your bank has had a problem. Can you just log in and confirm your log in details?" And you're logging into the phisher's website. Now that's a social engineering tag. That's a social engineering hook that gets you to sign in and give your details away.

The standard phishing attack is a very generic attack. But there's something called spear phishing where they target an individual person or small group of people using information that seems specific to them. Those I would suspect stand a much greater chance of success because people fall for them. People realize there's some information there that probably isn't generally known so they think this must be a legitimate request. So that's a situation where a social engineering attack can be very effective.

Julia Allen: Who is most likely to use social engineering? What kind of people are on the prowl using this particular technique?

Gary Hinson: Starting at a very broad level, we're all social engineers pretty much. We all use social engineering techniques to get information out of people that we're conversing with or doing business with. So at one level, it's part of just being human and just having normal social relationships.

If we're talking more directly about hackers and people attacking either companies or individuals, then I would say it's the hacker groups. It's the people who are deliberately hacking information. Some of them these days are either funded or directly are criminals, and we're seeing quite a lot more interest in hacking and phishing type techniques, for example, in identity theft. So there are situations like that when the hackers have basically another tool in their toolbox, another means of getting at information they can use.

Julia Allen: People are really the primary source of a vulnerability because, as you say, sometimes they don't even realize that they've been attacked by virtue of the information that they've divulged. Other than basically being trusting souls, why do you think people are the primary source of the vulnerability?

Gary Hinson: You've touched on it entirely there. It's around trust. It's around the fact that when somebody calls you or speaks to you about something, that your natural inclination as a human is to give them the information that they're after and try and be helpful. It's a particular problem for help desks or call centers who are trained to do that. They're trained to be helpful, and to try and get the other person to feel good about the call and so on.

So in those sort of situations, a good social engineer can use techniques to get information out that probably shouldn't have been disclosed. And there are tricks and lies and various psychological techniques. I've heard Kevin Mitnick talk about, for example, using pressure. So a classic example would be saying, "I need this information right now. I've got the chief exec waiting on the phone for this information. Are you going to get me to call the chief exec and so you can't give him the information because you wouldn't give it to me?" There are some people that will respond very badly to that kind of pressure. Some will just say, "Go away. I'm not having anything to do with this." But a lot of people will just say, "Okay, it's obviously urgent. The account number is this and the password is whatever."

Julia Allen: So the obvious user ID, password, account information, perhaps information that would identify you. Are there other types of information that you've seen or that susceptible users might divulge to someone that's calling them up?

Gary Hinson: Yes. It starts with the basics really. When you're first starting a social engineering type attack, you need to gather information to build your credibility as a caller. So at that level, you're wanting general information about the company, about the people, about names of projects, names of sites. Things like company lingo, terminology for the way things are done, the names of procedures, even the types of ID. So in some companies they use first name, last name. Other companies use last name, initial. So getting that kind of structure can be quite an innocent sounding call but that information is useful for the hacker to take forwards.

The other kinds of things that they like are maps. If they're planning a physical attack, and this is probably more along the lines of industrial espionage than straightforward hacking, but if you're planning to actually visit the headquarters of a target company and get inside and get confidential documents out, then a map of the building would be really useful. So that's the kind of thing that insiders may think, "Well, it's just a map. It's not significant. It's just a telephone directory or a list of user names." Or even an email with a nice long carbon copy list, so they can see all the different people in the company and try to guess relationships. These kinds of very general information sources are very useful to the social engineers in the early stages of hacking. And also personal information, so things like names and bank account details of people who work there.

Part 2: Identifying Social Engineering Attempts and Attacks

Julia Allen: So this could be kind of a daunting threat to protect against. So given all of this, and in the awareness work that you've done, what have you found are some effective actions that business leaders can take to address social engineering? I know you talk a lot about training and awareness, but what have you found to be some effective practices?

Gary Hinson: I think social engineering is one of those things that really the only effective control is through people. It's making people aware of the possibility that they might be socially engineered and giving them the tools, the kind of techniques to identify social engineering attacks, or potential attacks, and then what to do with them.

So one of my favorite controls in this area is the idea that if somebody thinks they might be dealing with a social engineer that they should pass it to what are called frontline workers. These are

people like PAs (personal assistants) and secretaries and call center staff, who are dealing with telephone calls, dealing with unknown callers all the time. The idea there is to try and focus these calls on those people and give those people specific training in social engineering. So they're more aware of the threat and they understand, when somebody passes them a call that says, "I'm not sure about this guy. He's asking some strange questions and I'm really not confident," they immediately think, "Okay, this could be an attack. Let's try and find out what's going on." So that's quite a useful technique.

The part of making that work is getting people to understand across the company that these things can occur and that there are policies and procedures in place. And the policies and procedures for them refer to passing calls to the front desk, or passing calls to somebody who can deal with it. So that's probably, I would say, the most popular type of control for social engineering.

Julia Allen: Yeah, it seems to me that that's kind of a fine balance because if you're in a customer-facing role, you certainly want to be accommodating and helpful and solicitous to your callers. But by the same token, you want to make sure you're talking to a legitimate customer, right?

Gary Hinson: Absolutely. That's the fine balance all the way through with social engineering because they can be really tricky. They can use all sorts of lies and deceit to appear legitimate callers. In my experiences, PAs, personal assistants and secretaries, particularly for senior people, are very adept at identifying sales calls. So these are sales people who ring up using, effectively, social engineering techniques to try and get hold of the boss, to try and ask him some questions and try and get him interested in a product. So they do that kind of thing pretty much all the time and they get quite adept at identifying potential sales calls.

So that's one way, one kind of trick in the training process, is to get them to think of social engineering just as particular version of that kind of pushy sales-rep-type call. If they make a mistake every so often and identify an actual sales call as a social engineering call, then it's no big deal. They manage to deflect the call and ask for more information and basically put them off. You don't really lose much that way. The down side is not too drastic. If they identify a legitimate contact, a legitimate caller, as a social engineer, then again, the impact can be reduced depending on what they do about the call. So if they just ask for more information, and ask the person to call back later or send an email or try and basically verify that they're a legitimate caller, then again, you don't really lose very much. A real caller will take the trouble to try and provide more information. A social engineer may give you more information, and may try and be even more brazen in an attempt to get through, but a lot of the time, they will just give up and find an easier target.

You're playing that game with them, really, that kind of mind game, to try and get information from the caller without being too pushy and too abrupt and too untrustworthy yourself. But it's a case of playing the game back on them, getting information out of them.

Part 3: Detective, Corrective, and Preventive Actions

Julia Allen: Have you found, if a social engineering attack has been successful and the organization or the person is aware that they've been socially engineered, what some detective or corrective actions could be? I know this is probably a particularly tough one to detect.

Gary Hinson: It is. The whole point about it is that people just don't realize they've been fooled. But it depends really on the company culture. If you've got somebody making a serious attempt to socially engineer another company, then they're going to be making a number of these calls over a period, maybe as much as a few days or even weeks. So if you can get people to report suspicious

calls like this to a central point and then correlate the information centrally, there's a chance you might be able to recognize that it looks like an attack is underway and send out a quick alert. To be honest, even if there is no attack, there's still value in sending out that quick alert to say, "Be on your guard. We think something's happening." It sounds kind of vague but people respond to that and they think, "Oh, I suppose, yeah, I hadn't really thought of that, and that's a good idea, and I'm going to be more careful." So it actually works for preventive control, even if there is no attack.

But it is hard to get people to make those calls and to identify when they've been attacked. Some people, for example, will not admit to the fact that they've fooled. This has been a problem with the phishing attacks, where people have gone through the phisher's website, given their information away, and then they have second thoughts and doubts. Sometimes it's a day or two later that they ring the bank and say, "I think I might have been a very silly person." So they're quite slow in getting around to responding. And by that time, it may well be too late. The fraud may have already happened.

Julia Allen: I like your suggestion about having people report in because sometimes an individual event by itself doesn't convey much meaning, but when you start to see a pattern, or if you see the same type of an approach across a variety of attempts, then, as you say, it really the integrated view that can give you more of an idea that you may be under attack.

Gary Hinson: Yeah. I think on that line, one of the best ideas I've seen is to get people to start reporting near misses as well as actual incidents. So you have an incident reporting process, typically it goes to the IT help desk or whatever it's called. You ask people to call up with near misses as well. So things that just about scrape through without causing a problem but something that perhaps we ought to learn from and try and avoid turning into an incident next time around. If you can get people to report that kind of stuff in a very open and helpful kind of manner and especially if you follow up those calls and thank them for the information and say how valuable it is and say how useful that kind of stuff is, you encourage this general corporate culture of providing information about attacks of all sorts.

Julia Allen: Have you found some particular approaches or ways of conducting the awareness and training that allows this to be addressed on a regular basis? Are there some good tips there that our listeners might take advantage of?

Gary Hinson: The first thing I would say is to look at the traditional way of doing security awareness and training, which basically doesn't work. This is the way that so many companies over so many years have tried and then kind of given up on, and I call it the sheep dip. The idea is you put everybody through a training course of some sort, probably once a year. You force them to go through the course, you say, "You must attend, otherwise you won't get your bonus," or whatever. So they come to the course basically because they have to, not because they want to. They get somebody at the front telling them a load of stuff, and they go away again. The information from that course sticks in their mind for a short while but it gradually decays and within a month or two, it's probably gone. They probably don't even remember going on the course. So that kind of once a year approach is pretty hopeless.

For that reason, we promote the idea of much more regular contact with employees. The technique that we favor in particular is using a monthly topic. So we'll pick a topic like, for example, social engineering or viruses or contingency planning, or something like that. Give a load of information out in one month to the various people in your company and focus on that topic. So there's something interesting there, something a bit different to what we normally read about, something that they can get their teeth into. If they're interested in it, they can get more information, they can go to additional resources. And then you move on to the next topic the next month. And if you're

clever about it, the next topic refers back to the previous month, and refers forward to the next month. So you start to get continuity between all the topics all around the area of information security.

Julia Allen: Do you find that this is best done face to face in some kind of an event, or do you do it via a website or virtual meetings? What techniques have you found to be effective?

Gary Hinson: All of the above. One of the things about getting through to people is people have different communications preferences. So some people like to read things. They're bookworms; they like seeing stuff in print, whether that's on paper or on the web. Some people don't respond unless they've got somebody telling them something. So they like to go to presentations and seminars. Some people like to get involved and like to do stuff. So we like the idea of case studies where we get people to imagine themselves in a scenario where, for example, a social engineering call has just come in. They imagine them taking that call and what do we do next? We ask them some questions and get some discussion going. So that's a way of getting them more involved in the topic. And some people like, I would say, the fun side of security. So they're not that keen on information security and maybe not IT. It's probably not really something that they consider is important, but you can still get them engaged through things like puzzles and crosswords and competitions. We promote a whole variety of techniques. It's a bit of a shotgun approach, hoping that you'll snag everybody. To be honest, you will not snag everybody every month. But maybe if you get some interesting topics in the year, you'll find something to interest everybody in the year.

Julia Allen: Well that sounds great, Gary. Just by way of closing, are there some resources that you recommend where our listeners can learn more?

Gary Hinson: Yes, certainly. I've got two or three to talk about. The first one is EDPACS. That's EDP Audit Control and Security newsletter, which has quite a bit more information specifically on this topic. Secondly, I've mentioned Kevin Mitnick a few times. Mitnick's written two or three books, *The Art of Deception* (*Controlling the Human Element of Security*) and *The Art of Intrusion* (*The Real Stories Behind the Exploits of Hackers, Intruders & Deceivers*) are two books that I would recommend. They're quite easy to read. They've got a lot of useful information in there. And the last one, on the books, is one by Ira Wintler, called *Spies Among Us* (*How to Stop the Spies, Terrorists, Hackers, and Criminals You Don't Even Know You Encounter Every Day*). Ira's done social engineering attacks as part of his penetration testing activities for companies. In his past he was a spy, so he did this for real. And the case studies in part two of *Spies Among Us* – really well written, very worth reading for any executive. I'd definitely recommend that. Last of all is our websites, which you've mentioned.

Julia Allen: Yes, we'll include all those links in our show notes, and I'm so very appreciative of your time and your expertise, and the work that you're doing on behalf of the information security community. So I thank you very, very much for your time today.

Gary Hinson: Thank you very much, Julia. It's been a pleasure.