SECURITY NOW!

Transcript of Episode #100

Listener Feedback Q&A #21

Description: Steve and Leo discuss questions asked by listeners of their previous episodes. They tie up loose ends, explore a wide range of topics that are too small to fill their own episode, clarify any confusion from previous installments, and present real world "application notes" for any of the security technologies and issues they have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-100.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-100-lq.mp3

INTRO: Netcasts you love, from people you trust. This is TWiT.

Leo Laporte: Bandwidth for Security Now! is provided by AOL Radio at AOL.com/podcasting.

This is Security Now! with Steve Gibson, Episode 100 for July 12, 2007: Your questions, Steve's answers.

It's time for Security Now!, our 100th episode. We need streamers. We need - we're actually recording this on the 4th of July. You'd think there'd be some fireworks. But no. Nothing.

Steve Gibson: No.

Leo: Unlike some of the other shows, we're just going to go along about our business. But congratulations, Steve. My deepest thanks for allowing us to carry this on the TWiT network. It is absolutely, after TWiT, the flagship podcast, it's the one everybody talks about.

Steve: Well, I've been really, really happy that we did it, Leo. I wouldn't be doing it were it not for you because, you know, you make me get up...

Leo: I think you were the second one we did; right?

Steve: You make me get up every morning and think about, okay, what are we going to do next week?

Leo: I'm sorry. I put you to work.

Steve: I think it's valuable. I found out when I was doing the keynote at Harvard, I was very surprised by the number of people who are Security Now! listeners who were in the audience. I mean, sure, it's that kind of audience. But still...

Leo: No, that's encouraging. That's really neat.

Steve: This matters to people. And they really enjoy it.

Leo: I'm thrilled.

Steve: I'm here.

Leo: Good. Well, it's a question-and-answer mod 4. If you divide 100 by 4 you get 25. So we've got our 21st question-and-answer session. You want to read a question or two from the mailbag before we get to our official dozen?

Steve: Well, you're going to be doing most of the reading this episode, as you do on our Q&As. So rather than doing mailbag - and in fact, some of these are mailbag-esque questions, again, because I was just finding so many good stuff.

Leo: I guess what you're saying is, instead of it being a question per se, it's a statement or an observation that is of great interest.

Steve: To some degree, yeah. But I will read something which, of course, I always love to read. I got a really interesting and fun piece of feedback from David Mitchell. And the subject was "You are the man," with about 15 exclamation points after it. And at first I thought that, you know, he was telling me that I'm the man. But it turns out that's from his mother.

Leo: Awww.

Steve: Anyway, so he says, "I listen to Security Now! fairly often. I'm a network admin part-time and full-time student. I figured I'd give credit where credit is due. I recommended SpinRite to my brother when his hard drive crashed. It didn't do much for him," and he says "(I suspected it was the hardware in his case, given the noises it was making, not much could be done.)" But he says, "My Mom is a district attorney in Alaska. When she had some power failures (a guy electrocuted himself and took out the power grid for the small town she's in)..." I guess he must have been low resistance.

Leo: Yeah, it went right through him, short-circuited the thing.

Steve: So some poor guy electrocutes himself and causes a power failure, problems for his mom's computer. And he says, "And her Windows wouldn't boot. She had serious cases before a grand jury, and the IT department was no help. Due to where she was, the best they could do was have her ship her PC to Anchorage for repair and perhaps total loss of her case files."

Leo: Oh, yeah, almost certainly.

Steve: Yeah. Well, I mean, they're not going to care. They're going to say, oh, give 'em a new hard drive. "Since it didn't do much for my brother, I figured she could use the copy he purchased" - he's talking about of SpinRite - "since I figured it was her best chance, since I couldn't do anything from the East Coast. It's a paid-for copy, and it stayed in the family, so I figured you wouldn't be too miffed at that." Of course not. So he says, "You can read below to see how it turned out. P.S.: I'm going to get on her about doing backups after this." And so this is where he attaches the mail forwarded from his mom, where her subject line is "You are the man," with the 15 exclamation points, referring to her son. So this is from Susan Mitchell, who says, "After days of system failure, to the point where I broke down in tears on the phone with our office manager (in Anchorage, of course)..." and of course we don't know what little podunk she's in. But wherever it is, if you electrocute yourself, the power goes out. She says, "I was finally able to sit down today, on my day off, and get the SpinRite on a floppy disk (had to scrounge to find one) and ... it worked. I then ran the original Windows CD to restore it, and I am back up and running. You are the man," she's saying to her son. "I know I am your mom; but seriously, I can give you very high marks and recommendations on your work, and best of all your knowledge."

Leo: Oh, isn't that nice. You know? And is it okay, I mean, is a license like a family license, is that okay?

Steve: It's the GRC tolerant license.

Leo: I like that.

Steve: I don't have a problem with that. If you need to fix your mom's computer, by all means.

Leo: But if it fixes it, you know - I'll tell you what. Since it didn't fix his brother's, okay.

Steve: Yeah, completely.

Leo: If it fixes two computers, then maybe you'd want to buy a second license. I'm trying to help you out here, Steve.

Steve: Okay.

Leo: Trying to make you some money. All right. Let's get to our questions. There's 12 great questions, as always with our listener Q&As. Starting with Evan Moore in Leesburg, Virginia, with an authentication puzzler. You know, I have to say parenthetically that the authentication episodes you've done, and now the TPM module that you did last week, has really been great, and obviously spurred a lot of interest. It's something that - you quantified something that we all kind of are somewhat familiar with and really explained why this is better than that, you know, two is better than one and so forth. And it's really been great.

So here's the situation. See if you can untangle this one: We have a password ID form for students/faculty to confirm their identity outside our campus if they need their log-in credentials auto-emailed to the email address we have on file. In other words, they're off-campus, they want the VPN certificate, and so they have to fill out this form. Since we don't tell the student what email address we sent their credentials to, some students don't know what email address they should go to to locate their credentials. Some students figure out on their own what email address we sent their credentials to. This is going to be more of a problem has people get lots of mailboxes, actually. But they may not remember the credentials to get into mailboxes like Gmail or Hotmail. And then there are those who forget what their security questions are in the third-party mail site, so they can't change their credentials to enter their old mailboxes to get to our auto email.

So here's the solution they've come up with: So very soon our online university plans to implement a change in the forgotten password ID request form for students and faculty who have these issues requesting their log-on credentials. The discussion that IT and academics have had on the subject would still allow a student or faculty member to confirm their identity outside the campus. But the required fields we ask for would be their first and last name, date of birth, full social security number, and email address on file.

The controversial new field on this updated form, which is why I'm writing, would be a field asking the student to enter the email address they'd like us to auto-email their credentials to if they enter all the fields correctly. The question is, would you consider allowing the students the ability to enter any email address they choose to send their requested credentials to as a security risk? Is that a security risk? Would you be able to explain your thoughts on that? Our concern is a few of us worry that, by allowing a student to send their log-in credentials to any arbitrary email address, we're opening a security risk because, well, we're assuming the student's who they claim they are, but we're not confirming if they have access to the email address on file. Those of us who are concerned have never seen a form before that allows a user to send their log-in credentials to a different email on file without the vendor confirming the user has access to the email address the vendor has on file. Apparently we get a lot of requests for ID/passwords from our student services department, so IT and academics are trying to devise a way to reduce these email and phone calls.

The problem is really that students often are using multiple email boxes and maybe not have access to the ones that you're going to send the credential to, so.

Steve: Exactly.

Leo: I understand the issue.

Steve: Well, and the really interesting point, the reason I thought this was a great question is, I mean, it's a perfect authentication/verification/security question. His point is, the fact that we're allowing somebody whose identity we don't know for sure to specify where their credential token should be mailed to, you know, that's the gotcha because they haven't verified that that person controls that email address. And in fact, they haven't verified that this is the person. They're assuming you provide your first name, last name, social security number, and then provide the email address you want the credential sent to, that that's, you know, authenticating you to the system.

But they're concerned that they're allowing the credential to be then sent to an email address provided at the same time on that form. And I agree, that really does create a gray zone. And you might - and he refers to saying, well, how do we know they control that email address? Well, that really isn't germane either because, if some bad guy wanted to obtain these credentials and log in as somebody else, they would of course create some random Yahoo! account. They would provide that information on the form. The credentials, even if you had an email feedback loop where, for example, the server then sent a verification to that Yahoo! account, you know, click this link to confirm you are who you are, well, the bad guy who created the account would also have provided that account's email address, to which the verification would get sent. So he'd go there, click on it, and confirm the process.

So the point is, this is opening, you know, doing this opens up a clear breach in what would be otherwise the security that they had in place before, which assumed, for example, that the physical person was in the office at the university filling out this information where they wrote down the email address to which they want any lost credentials to be emailed. You know, there, where you have physical identification and proximity, it's certainly reasonable to ask for that information. But as soon as you allow it to be done with no other verification, you're in trouble because, for example, we've talked about all the problems with site injection, you know, various types of remote code injection. There could be some sort of a glitch somewhere that would allow people's social security numbers to get out.

And of course we've seen that. We hear about that all the time. So if the university lost control of that, and the student's name and social security number were all that was necessary in order for someone remotely to provide a new email address, basically they've commandeered that student's account. So I think it really does represent a problem.

You know, the only thing I could say is that, you know, they're softening their existing security to make up for the fact that faculty and students aren't taking enough responsibility for their own security, like forgetting which email account they use. It's like, come on. I mean, I understand people are going to forget things. But the user has to have some responsibility. Otherwise you're going to compromise security. And I really think this makes a substantial, a substantive change in the protocol that they had deliberately established.

Leo: And I also don't like the idea of using a social security number for an identifying strip anyway.

Steve: No, that is creepy.

Leo: That's a problem in a lot of different ways. I mean - anyway. So you're saying don't do it.

Steve: I'd say that I think the concern that he expresses is that something feels wrong to him. It's like, okay...

Leo: He's right, yeah.

Steve: ...he's right that there's something wrong. It is a fundamental change in the security architecture that they had established. And it's a change not for the better.

Leo: Yeah. Paul in Edmonton, Alberta writes: In Episode 96 you answered a question on resource depletion using VMware. I thought I'd tell you, you can buy a multiprocessor, multi-core computer, from Polywell, for instance, with eight dual-core Opteron processors, up to 128 gigabytes of DDR2 RAM. Works well for VMware.

Steve: I just love that. You know, because remember that we were talking about the problem with virtual machines is that, because they think they're in the equivalent of a real machine, they assume they have access to all of the RAM. So when you create a VM, you need to give it, you know...

Leo: A lot of RAM, yeah.

Steve: ...half a gig or a gig or a gig and a half, whatever. And in doing so, you have taken it away from your host system. I mean, it's just gone. And so of course the Opteron is the AMD 64-bit processor. And here he's talking about having 128 gigs of RAM.

So the thing I thought was interesting about this is that it's - 32-bit address space we know is limited to 4 gig. You're only able to address 4 gigs with 32 bits. So by going to 64 bits, you break out of the 4-gig barrier. And so, yes, you could certainly be running a 64-bit multi, or even single, for that matter, Opteron system with much more than 4 gigs, and then have many more VMware systems emulating a 32-bit environment. So basically you're using the 64-bitness, not because you need 64-bit operations, but just because you've gone VM happy and you'd like to have, you know, 20 or 30 independent machines, each with their own 4 gigs running at the same time. So that is a very cool platform. You know, a little overkill, maybe, but certainly very cool.

Leo: I'm pricing one out right now. They're really designed for servers, obviously. They're not designed for desktop computers.

Steve: Yeah, but, I mean, it's a perfect host for multiple VMs.

Leo: I guess, if you're willing to spend the money. Chris Rydin of Troy, Michigan, has been flooded with bot spam. He says: I'm the tech guy for my company. I noticed yesterday our mail servers were being bombarded with spam letters, much more than ever before. When looking at the spam in my quarantine, I noticed they all had nearly identical headers. "You received a postcard from," and then it adds a random prefix like "a neighbor" or "a friend." It tells you to click on a link to catcher.hk, which installs a worm on your system. This is spam you can get a virus from, much like the real spam.

I noticed our email servers, which support many, many companies, reported more than - oh, get this - half of the total mail it rejects, including regular spam, was from this site. Half. It sends mail from spoofed headers to make it look like it comes from reputable companies. Upon talking with friends and coworkers, they say they got this on their home email, too. However, I find very little about these attacks being reported. The virus it said it detected was html/postcard.N@trojan. And the few sites that actually adhere to this name report that they, too, are seeing very high traffic of this worm. What's going on here? I'm not familiar with this one, actually.

Steve: Well, this is a classic spam-sending botnet. Basically this is exact - you know, the way that his server is able to see such a high incidence of this and, at the same time, other people scattered around the Internet are seeing this, is that some large botnet was contracted with to send this out. Or it is very likely, since he talks about this thing having a link to some site in .hk, so in Hong Kong, that is probably the botnet trying to spread itself. And remember we talked about, when we talked about the whole bot roast FBI effort, that now we know that large gangs of bot herders are competing with each other, trying to take over bots from each other. So networks are being used to promote their own bot, using their own network. So to do that they spray massive amounts of spam, which they can do because they've compromised hundreds of thousands of end-user systems. They spray this spam that's got links.

Now, the problem is that Hong Kong site will be taken offline very quickly, so there's a short window of opportunity during which, you know, before the site is taken down through some cooperative effort of law enforcement, the trick is how much spam can we get out to how many people who will say, oh, I got a postcard from a friend, and click the link and infect themselves with a bot. Maybe their system is already infected, in which case this bot will try to take over from any bots you had before, to now you've got bots having war on your own machine. But, you know, that's what's happened here in this case.

Leo: And by the way, I'm just scanning through news stories, and this has been going on for years. I mean, I've found stories about a postcard trojan from 2004, 2005. This has been going on forever. And, you know, it's funny that I think nowadays you don't see as many alerts going out because everybody's got pretty effective spam filters and antivirus filters. And, you know, I for one, I never see a single one of these because I'm sure they're getting filtered out by my antispam service.

Steve: Well, exactly. And that's the point was he was looking at the server.

Leo: He was looking at the antispam server.

Steve: Yes, and noticing that more than half of what the server was filtering. So it's exactly as you say, Leo, it's the case that not that many of these actually make it through.

Leo: We hope.

Steve: But again, it's a numbers game. So the more they flood out, the more are going to get through.

Leo: I'll have to ask my antispam guys at MailRoute if they're seeing a lot of them. Mark in Wapakoneta, Ohio asks: Are hard drives getting less reliable? I've been the computer fix-it guy for my family for a number of years now. I've noticed a lot of hard drive failures recently. Most of the time when one of my family members has asked me to fix their computer, it's been the standard spyware/virus stuff they always seem to be infected with, or at least that used to be the case. Up till six months ago I'd never run across a dead hard drive. But the last four computers I fixed have all needed new hard drives. Of these, three were laptops, so maybe it's the increase in laptops that's leading to more drive failures. In any case, I'd like to hear the opinion of a hard drive guru on the subject. I've been listening since Episode 1 and using ShieldsUP! for years. Well, you are that.

Steve: And we do know that, from the article that - or the report and study that Google did, and we did a podcast on this some months ago, that in fact hard drives are certainly less reliable - I think, if I remember right, by at least a factor of nine - than manufacturers are stating. And it does seem that newer drives are failing more often. I'm always worried about hard drives in laptops, in the same way that I'm worried about hard drives in iPods, that they're inherently susceptible to physical abuse. And I'll see people take a laptop running - and I hang out at Starbucks, which tends to be laptop heaven because they've got T-Mobile service there. And they just, you know, pick up their laptop, and they don't drop it hard on another table, but they just...

Leo: There's a clunk.

Steve: Yes. They don't get it that inside is phenomenal technology which has enabled a few small spinning metal platters to store literally tens of billions of bits of data. I mean, I can't believe this stuff works, and here they are, just clunk. And I just shiver whenever I see it.

Leo: It's not just that, it's also a lot hotter in those hard drives.

Steve: Oh, in fact, heat is a real problem for laptops. And in fact, you know, me being aware of this as I am, I set my laptop - first of all, I try never even to move it when the hard drive is spinning. But if so, I put it down one corner at a time, first one corner, then one side, and then I slowly tilt it back down. I mean, I'm maybe a little carried away, but I've never had a laptop hard drive die on me. So I do think people really need to be very conscious of the fact that there is a delicate, spinning piece of phenomenal technology inside there.

Leo: I haven't had a hard drive die in a while. But the most recent one was an iPod, just a few months ago. Or even a month ago. And I've been lucky with the other hard drives, though, so far. You know, parenthetically, my new iPhone, which is not using a hard drive, as far as I know, I'm pretty sure it's solid-state memory, gets very hot, when it's charging particularly. And of course it's a hot day, it's about 90 degrees right now, and it's pretty warm. Is there the same problem for solid-state memories with heat and getting banged around?

Steve: Well, electronics really doesn't like heat at all. But there's no problem with it physically. And in fact, this is why the only iPods I have ever owned are the solid-state ones. I just...

Leo: They're very reliable, yeah.

Steve: Yeah, I was just too freaked out by the idea of having a little hard drive in something that just sort of seems like you ought to be able to not treat it with kid gloves every single second.

Leo: Well, you'd better believe for 600 bucks I'm treating my iPhone with kid gloves, regardless of its hard drive. Nitin Saini of Tampa, Florida - I know I'm mangling your name, but I apologize.

Steve: Well, but we tried valiantly.

Leo: We tried again and again - prefers TreeSize to SpaceMonger. Remember we were talking about ways to look at your hard drive space. That was on 96. He says: I have one comment. You mentioned SpaceMonger, which I've looked at in the past, wasn't so much impressed with. I’ve been using TreeSize, which also has a free version and can be downloaded as a standalone exe, less than a megabyte. And he gives us the link, it’s jam-software.com.

Steve: And I should mention that I am also a TreeSize user, and I have a paid version of it that I like very much. They’re very different. SpaceMonger is so cool because it gives you a visual, graphical, nested rectangle-like map over your entire screen of the layout of your drive in terms of the size different areas occupy. I really like that to give you a quick sense of what’s going on. What’s cool about TreeSize is it very quickly parses your entire drive and will give you a list sorted from largest to smallest of the largest hundred files on your drive. And it has lots of other features, too. It’s got cool little bar graphs, and it’ll do graphics of the distribution of file size by type and a bunch of other stuff that I really don’t care that much about. But it is just nice to be able to instantly get a sorted list by size of the biggest files on your system because invariably you’ll find stuff there, it’s like, oh, I don’t need that anymore, and look how much space it’s taking up.

Leo: Oh, yeah.

Steve: So I like them both. But they’re very different, and they serve a different purpose.

Leo: On the Mac I use a similar program called Disk Inventory X. It’s free. It gives you a color graphical map as well as a list of files. Disk Inventory X.

Geoff in Ottawa worries about unchangeable biometrics: I have a question regarding the use of biometric data as a form of authentication. I recently had a discussion with someone who’s much smarter than I in questions surrounding technology. He was presenting the argument that retinal scans or some other personally unique biometric factor is the way of the future. None of us will carry credit cards or similar tokens in ten years.

My response? Well, there are two kinds of data in this world, data that’s been lost or stolen and data that will be lost or stolen. Ooh, this guy’s a pessimist. The retinal scan, for instance, is just a series of zeros and ones that are structured to represent the unique pattern of blood vessels in my eye. If a bad guy were to gain access to that digital representation of my eye, would he not have my password forever? That’s the point here. Isn’t it better to use tokens that can be reissued if compromised? For instance, the credit card numbers we now carry? He felt this could never happen, what with the strong encryption algorithms and such we have today. Never is a long time. Am I missing something here, or do I lack some specific knowledge that will make it all work beautifully?

Steve: Well, it’s a great question, and it comes up often when we’re talking about biometrics. And the point of course is that, if we lose control of our credit card or if we lose control of our password, we’re able to change our password. And of course best password practice says change it every so often even if you don’t think you’ve lost control of it, just because it’s a good thing to do. Maybe you have and don’t know it. But you can’t change your fingerprint. You can’t change your iris or your retina or your DNA. You know, ultimately we can presume that’ll be what we’re locked onto.

So people are concerned that, wait a minute, I don’t want my retina scanned and essentially then moved out of my control because what if some bad guy did something wrong with it? You know, it’s then don’t they have me, my essence essentially, forever. And I guess actually if they had your DNA they could just clone you, and they could get your fingerprint that way. Here’s an eyeball. So it’s an interesting point. The one mitigating factor here is that these technologies will generally not save the retina scan of your retina in the same way that they don't scan and save your actual fingerprint. What's done is what's called "feature extraction," and that's the key. That's what makes this not such a bad thing is that - oh, and the case would be the same for DNA likely.

But in the case, for example, of a fingerprint - and in future podcasts we're going to talk in more detail about how this stuff works, but the idea is that the image of your fingerprint isn't saved. I mean, it's not really the way you see it on CSI where they're flashing all these fingerprints by because that's a phenomenal amount of data. What's done is the fingerprint image is analyzed for features. There are places where you have dead ends in your fingerprint, places where two ridges merge, places where a ridge stops. Anyway, the idea is that an analysis is done, and the actual image is reduced to a series of much smaller and simpler characteristics. And then what is done is this is hashed into a 128, 168-bit hash. And so that's what is stored. That's why it's safe is that the hash - as we know, a hash is a lossy, one-way function. And so from the hash number there is no way to go back. So you couldn't even go back to the features that were used in order to synthesize a representative fingerprint that would hash back down to the same value. There's no way to go back.

So all these various biometrics, given that they're implemented correctly - and of course, with anything dealing with security you've got was this done right, was this done wrong, were passwords stored or were their hashes stored, for example, on the server, as we've talked about recently. In this case you're using the same thing. You're taking the biometric data, and you are doing a feature extraction. Then you're hashing those features into a token which is a unique representation of the person without being directly connectable to the physical representation of the person. So that's why it's a good thing.

Leo: It's turtles all the way down.

Steve: There you go.

Leo: I guess, I mean, I guess the features would be immutable, though; right? So if they got your eye, they got your features.

Steve: Yes. Okay, right. If there was a place somewhere between the scanner of your retina and the hash coming out...

Leo: The database.

Steve: Well, no. The hash is the database. The database would just be a hash. But if there was a...

Leo: Oh, I see. And the hash is not reversible. So what you're really saying is the system's secure.

Steve: Yes. But it's certainly the case, if someone got the actual image of your retina, on the other hand, then the question would be how would they...

Leo: It would have to be pretty severely compromised.

Steve: ...synthesize a retina that would fake a scanner. I guess what they're saying is, okay, if you had the database of hashes of retinas, then you'd have...

Leo: You could reverse it.

Steve: ...everybody's hash. And the point is that the technology would - every time the technology's applied, the same retina would generate the same set of extracted features that would generate the same hash. Although that's not necessarily the case either because different features could be brought out. So you'd have to have...

Leo: Right. Some hashes, you know, one machine might work differently than the other.

Steve: Exactly. And so you would have a different mapping of physical features into hashes, depending upon technology. So that means that, even though you've got one single finger, the different scanners would extract different features in a different way or choose different features to hash. And so there again you have a one-to-many mapping. So that it's not just like once something is compromised, it's compromised forever.

Leo: Makes sense. Robert James - huh?

Steve: So again, it's a good thing.

Leo: It's a good thing. Robert James in San Francisco wonders about ID Vault: I've been wondering about the ID Vault I keep hearing advertised by Kevin Mitnick and Leo. I did one ad once on KGO in San Francisco. But now I'm - this is the problem with doing any ads on the radio at all. Is this basically like Roboform or Keepass, or is it something even more secure? Would you be able to address it on your Security Now!? The link is guardid.com/idvault.

Steve: And here, Leo, I figured you were an expert, so this was a question for you.

Leo: Well, you know what, actually you'll be interested in this. And really their whole premise is exactly what you were talking about. In fact, if you go to their website, they'll say, "Security experts have long known that two-factor authentication, which requires both something you know and something you have to access your online accounts, is much more secure." They obviously listen to the show. What it is, is it's a smart card in a USB dongle. So it's coded to you. It is not a thumbprint recognizer or anything like that. But it is a second form, it's an additional form of authentication, equivalent to an ATM card.

Steve: Okay.

Leo: So that makes sense; right? And then they add software on top of it so that it does things like remember passwords and so forth. And that password database is stored on the computer - I believe it's stored on the computer. I don't believe it's stored on the ID Vault. But it's encrypted by the key in the ID Vault. So you'll see these on USB keys sold by other companies, as well, where there's some storage on the USB key, and that's encrypted by hardware in the storage. Same idea. In this case it’s just a key on a card, and it’s a second form of authentication. So you add it to your password, which you know; right? Something you have and something you know, and you’re that much safer.

Steve: So it’s two-factor authentication where you’ve got a portable key ring hardware token.

Leo: Exactly.

Steve: Neat.

Leo: Yeah. And it encrypts all of your identity information, passwords and so forth, on the hard drive of the - again, this I’m not sure about, but I believe it does that on the hard drive of the computer.

Steve: Okay.

Leo: Well, wait a minute. No, I’m sorry, it does have storage on it. So you take it on the USB key. So it keeps a database - it’s like Roboform AI, but it keeps the database on the USB key with that secure chip, and it’s secured there. So you plug it in, you choose your account, you enter your PIN, and then you’re signed on, and you’ve got your passwords with you. Does that make sense?

Steve: You know, Leo, as we’ve been talking about this, and as I talk about these questions, I just - I cannot wait for the day where we have OpenID-style sign-on of some form. It is just going to be so nice. I think, in fact...

Leo: This could do that, of course, because this would be your OpenID key, basically.

Steve: Yes. Exactly. You could have it set up so that your OpenID authentication server looked for that information on the USB dongle and got your physical presence verification there. But, I mean, just to have - I don’t even care if it’s literally OpenID, where I’ve got my own token. Just to use the same token to log in anywhere, oh, that would be so nice.

Leo: Wouldn’t that be nice? I guess really the main thing about ID Vault is that they have a million-dollar guarantee. If someone steals your financial credentials and uses them to access your financial accounts, they will reimburse you up to a million dollars. So if nothing else...

Steve: I’d like to see someone try to collect on that.

Leo: Claim that? Yeah, that’s true. I don’t know if they have any examples.

Steve: And the problem is that there are so many ways that that process could go wrong that you’d have a hard time proving to them that it was theirs...

Leo: It was their fault, yeah. Yeah, that's a good point, yeah. Moving on to Laura in Arbor Vitae, Wisconsin. I'm sure I'm pronouncing that wrong, but that's how you would pronounce it if it were Latin.

Steve: If it was career vitae; right?

Leo: Yeah, curriculum vitae, yeah. Arbor Vitae in Wisconsin has a VPN question. She writes: Here's your background for my question. When I connect a VPN to my employer it disables my ability to print to my local network printer, which is a 192.168, because I've been assigned an IP in my company's scheme, 172.16. I understand why this is happening, but here's the question: If I use Hamachi or OpenVPN at home when I travel, will I have access to my other devices - TiVo, printer, other computers - on my home network, or just the one I'm VPNing into? Oh, she's got a little idea here.

Steve: Well, it's a great question, and it brings up two things. First of all, Hamachi of course, as we remember from our classic "Hamachi Rocks" episode, Hamachi runs on the interesting five-dot network, which is not formally available, but it's an unassigned IP space on the Internet. So Hamachi gives every single machine that's running Hamachi its own 5.x.y.z IP address. But for exactly that reason she is restricted to only accessing that machine through the Hamachi network because, when she's accessing that machine, the machine is, for all intents and purposes, on the five-dot network. And therefore she's unable to access devices in her home 192.168 network. Although, if the printer were connected physically to that machine, then she would be able to access that machine that way.

Also the other nice thing about Hamachi is that it is a bridge for MAC addresses, I mean, yeah, for MAC addresses. That is to say it is a MAC layer, Ethernet layer link, not an IP layer link. What that means practically is that network broadcasts, that is to say, ARP broadcasts of MAC addresses are able to flow across Hamachi. So things like Windows filesharing and printer sharing are able to work across Hamachi.

Now, OpenVPN is a different animal. And in fact, you know, I promised a year ago, a year and a half ago I think it was, that I would have the OpenVPN how-to assembled. I've gotten stalled on that because OpenVPN feels to me like the wrong solution. I'm trying to find something that works much more easily and does the same thing. But OpenVPN can work with what's called a "tap driver," or a "tunnel driver," tap or tun. In the tap case you're able to get the same sort of functionality that Hamachi provides, where a machine connecting to OpenVPN appears on your network like another machine on your network, and then you can see all the network resources that are available. I'm sorry, it's more like what her VPN would see, not what Hamachi does. In the tunneling driver you don't have that kind of capability.

Anyway, so unfortunately there isn't a super simple, clean and clear answer. I'm trying to find a solution that will provide users with exactly the capabilities that she wants. You can get it with OpenVPN if you set it up for a tap driver. But it requires bridging the OpenVPN adapter to the system adapter. You need third-party software to do that under Windows 2000. You can do it under XP. Presumably you can also do it under Vista, although I haven't verified, but I'm sure you can. But it's really complex, much more than I would like it to be. And OpenVPN is, too. So I'm looking right now for a better solution for that.

Leo: So I imagine that IT guys often do that, though, in the office because they want you to be able to use the network printers and so forth once you've VPN'd in. Yes?

Steve: Yes, exactly. You VPN into your corporate network. Now your remote machine is exactly like it’s on the corporate network. And you’ve got access to the internal IM and chat and printers and all the resources, so just like the machine were on that network.

Leo: But that doesn’t happen when you log in via VPN to your system.

Steve: Well, it can, if you set up OpenVPN correctly. And it’s not the default way. You’ve got to fight everything in order to make it work. So it’s like, okay, rather than trying to explain this, maybe I’ve just got to find a better solution. So I’m looking.

Leo: Seems like there ought to be, yeah. Jim in San Diego writes: A lot of IE vulnerabilities have been related to ActiveX. Do Firefox add-ons present the same potential danger as ActiveX, or are they inherently more secure?

Steve: Well, that’s a great question. Okay. Backing up a little bit, remember that ActiveX is Microsoft’s OS integration technology. It’s really just a DLL implementing their original protocol for essentially creating components. So it’s a componentized DLL. The thing that’s most frightening about ActiveX is that, until very recently, Windows and IE would download these in their default configuration on the Internet and instantiate them, that is to say, download and run them in the default configuration without the user’s knowledge. Finally with IE7 you’re given a bar at the top of your screen when a site tries to run an ActiveX control on your machine. And I see that all the time. So it demonstrates how much of that was going on.

Leo: I think it used to ask you for - didn’t it ask you for a certificate or something, or show you a certificate and ask for your approval in the past? I think it was supposed to.

Steve: Whatever it was...

Leo: It wasn’t effective.

Steve: ...it was happening in such, yeah, things like CoolWebSearch, which was a notoriously horrible ActiveX control, was finding a way onto users’ browsers all the time.

Leo: I think the design was that the user still had to approve it. But it wasn’t as clear about what you were doing.

Steve: Well, and there were things...

Leo: And there were holes.

Steve: There were things like, do you want to download and run ActiveX controls? And so, okay, if it’s not signed, don’t. But, you know, bad guys would sign theirs.

Leo: Right, exactly.

Steve: And then it would be signed. So it was like, thanks a lot, that provides me no help. Anyway, getting back to Jim's point, having completely beaten up on ActiveX now, a Firefox plug-in has potentially the same problems because it is software that you are running, and it's in a very dangerous place. The browser is a dangerous place to have add-on components. Even the browser itself has enough trouble staying secure. Now you're talking about adding third-party components to a high-traffic, Internet-exposed, dangerous place that's running on your computer. It scares me a lot. So I think, yes, although we see more problems with Active X because of Windows and IE still being the majority platform, there's as much a danger theoretically over on the Firefox side.

Leo: I'm going to give you the alternate point of view, though. Because remember, Steve believes that JavaScript is dangerous. So, I mean, all of these controls on Firefox are running in XUL. So essentially they have the same permissions JavaScript does. They're not ActiveX. They're not DLLs. They're running in a scripting language.

Steve: Oh, okay. In that case that is potentially better, although again...

Leo: Same risk. Same risk as JavaScript.

Steve: And scripts have buffer overflows all the time, too.

Leo: But it's my understanding, and I may be wrong, but it's my understanding from talking to Firefox developers that these extensions are running as XUL. It may be possible for you to attach an application to it. That's an interesting question. But it's my understanding that it's really running in a markup language, basically.

Steve: Okay. Well, what I would recommend - and, you know, certainly anytime you've got things going on with your browser, protecting your system from that is a good thing. One of the plans I have in the future is to get the author of Sandboxie on to spend an hour with us talking about Sandboxie. I've contacted him; he's interested in doing so. It's just a very well-behaved, very nice sandboxing tool that is not in your way. It's much lighter weight than having to deal with a virtual machine and running your browser in a VM. There's a lot of overhead associated with doing that, as is the case with any VM where you're running a whole 'nother instance of an OS in the VM. I want to talk about sandboxing because I think it is potentially the not-received-enough-press-yet solution for all of this.

Leo: Right, right. I agree. Moving along to Ipigi of Montreal, Canada - Ipigi, that's probably it, Ipigi: I've been working at an ISP for the past five years, three of these as a customer service representative and now as a technician. I also have my own small business on the side where I troubleshoot computers, be it just basic support, all the way up to virus and spyware removal, data recovery, network layout, and security. My comment is about Episode 97, the bot roast. You mentioned that more responsibility should be assigned to the ISP - I said that, I'm not going to put that on Steve - in helping to track down users with an infected computer or even actively intercept and log traffic to make sure none of their users are infected with network-aware malware.

There are two problems I see with this. If the ISP is small, it probably does not have the infrastructure required to intercept such traffic and scan it efficiently. Not only that, but I'd be surprised if they even had the financial resources available for such hardware and manpower to be made available to them. In the case of a large ISP that it would be nearly impossible to log and inspect all traffic from their users. The mere size of the traffic logs would quickly be unmanageable. As an example, our company logs every connection and disconnection from our service. It logs the name of the client, the IP address assigned to that client, the phone number of that client, date and time, as well as the IP address of the hardware where the client is currently connected to internally. This information alone in a text file generates around 20 megabytes of log files a day. Just imagine if all traffic had to be logged for every single user. Imagine if a large DSL or cable provider had to do the same.

In my opinion the FBI and ISPs should instead make an agreement. The FBI should provide the IP addresses and times where malicious traffic was detected. The ISP could then inform its clients in any way they deem necessary that their computers have been compromised, ask them to have their computers checked for malware. If a client did not comply and further infractions reported, then the service of that client should be suspended.

Finally, it's also my opinion that computer owners should be responsible for their own computers' health. Good luck. To make a comparison, you would not blame a mechanic if the owner of the car did not take good care of it and bring it in once in a while for a checkup. In my experience, most people believe they have virus and spyware protection when in reality they are running an old version of Norton Antivirus that expired three years ago and is severely out of date. I agree with you on that. I've been confronted many times by clients who said they had proper virus protection, lulled into a false sense of confidence by the free antivirus trial that came bundled with their system.

Steve: Well. When I'm talking about ISPs taking responsibility, I'm not sort of referring to them doing it in what I guess I would call an "old school" fashion.

Leo: I don't think they should be logging all traffic. I don't think anybody's proposing that.

Steve: Exactly. What I meant, and I guess I wasn't clear enough about it because I'm sure that he would have heard me if I'd said it...

Leo: Well, we made an off - I think it was probably me. I made an offhand comment that really all of this could be eliminated if ISPs would pay attention.

Steve: Well, okay. What I think we need, we need another class of hardware at the border. We need something like a router, but basically an anti-traffic flood router. I mean a router which itself is watching all of the traffic that is crossing its boundary and keeping some track of it. And there is technology like that. It's certainly available. I didn't mean to say that, for example, logs should be generated and then parsed and action should be taken because, frankly, by the time that's happened, the damage has been done, and it's over. What you want is you want technology very much like sort of bandwidth-limiting technology where users are able to send a certain amount of data per interval, and beyond that they can't. I mean, essentially, you know, most users have not only higher bandwidth downstream, but most of their use is downloading stuff. Much less is going in the other direction. Certainly filesharing is an instance where you have more symmetrical flow, or email. But in the case of email, and certainly in web surfing, your upstream bandwidth can be much smaller because you've got much less data transiting in the other direction.

So I guess what I'm talking about is - I mean, and we're years away from this. But it's entirely possible to have equipment that the ISP runs which is next-generation router equipment that is itself responsible for detecting when one of their customers, or more at the same time, of course, is sending out a massive amount of anything, whether it's ICMP pings to DoS somebody, or spam, you know, out to remote port 25 in order to spam somebody, and just say, whoops, this is a violation of our terms of service, and disconnect it. As somebody has said, I think, that we were responding to in the posts, you know, route that person to a web page so that anywhere that customer goes, they go to an ISP web page that says, we believe your computer is infected with something malicious. Please contact our support people.

Leo: There are all sorts of issues with this. I mean, as soon as Comcast started blocking port 25 we heard a lot about it. People get upset about it. So it's tricky. I understand it's not easy for ISPs. But I think there are things they can do.

Steve: Yeah. And the idea that the FBI would work with the ISP, there again, there's too much delay. There's the need for too much interaction and too much time passage. But more importantly, then the FBI would only be dealing with the machines that were attacking end-users that the FBI was involved in pursuing, rather than a solution that basically nips it in the bud, that prevents this stuff from having a damaging outbound effect from the ISP on, which really is the right place to do this. We just need some hardware that can automate the process.

Leo: Yeah, yeah. And I'm sure that that hardware even exists. I just think that ISPs do bear some responsibility in all of this. How much and how they implement it, of course, is up in the air.

Steve: And also, to address his last point, I agree with you, Leo. I mean, yes, it would be nice...

Leo: It would be great.

Steve: ...if users were taking more responsibility. I mean, I have to imagine that anyone hearing our voices over Security Now!...

Leo: They know what they're doing.

Steve: They're not bot infested themselves. They may be seeing other people who are and dealing with their friends' malware infestations. But again, it's not the responsible people, it's the irresponsible people that get this stuff hosted on their machines.

Leo: And Steve and I both make it kind of our goals in life to help people, to teach people, to explain all this stuff so that they can do it. Certainly that's why I go on the radio six hours a week. I spend a lot of time - and why you come on the radio show. I mean, last week on the radio you talked about making good passwords. I mean, we try to teach people this. But I am not at all confident that the majority of users will ever get good at this. And that's why, unfortunately, the burden rests on people upstream - the ISPs, Microsoft, and on and on and on. Or maybe you should just take computers away from people and just, you know, they should have to take a test before they get one.

Steve: Yeah, then how will they listen to our Security Now!. Oh, I guess we wouldn't need Security Now! then, would we.

Leo: We wouldn't need it. That's exactly right. We didn't need it...

Steve: Let's go back to the Stone Age.

Leo: Chris Grant in Riverside, California wonders about phone radio network security: Is there good, reliable security implemented on GPRS/EDGE networks that AT&T uses? I'm suspicious that devices like my 2125 WM5 Smartphone - that's a Windows Mobile. That actually is a really nice one, the Cingular - or for that matter the new iPhone, most likely transmit passwords for email via clear text. What is the security of this EDGE network we are using?

Steve: Many, many people have asked questions about wireless telephony security. And it's something I've been poking around in, trying to get some good answers to. So I can answer the question, is this stuff going into clear? The answer is no. All of these connections, these digital telephone connections are absolutely, definitely encrypted. I just haven't been able to find any good information that talks about how. It's an outstanding issue for me to research. When I finally track it down, we will certainly do a podcast about it because it interests me greatly, too. I would like to know every little dirty, nitty-gritty detail about the encryption used on our cellular systems.

Leo: But we don't right now.

Steve: But we do know that it is encrypted, and his stuff is not in the clear. Of that I'm sure. We just don't know if it's absolutely unbustable for anybody listening in. I have a feeling a good job has been done on it because it's recent enough that there would be no excuse for it not to have been done correctly.

Leo: Yeah. Ryan Jones of Melbourne, Australia wonders about fingerprint capture: You guys have been talking a lot about bio security lately, and I have a question for you today that came from a customer of mine when I was at work. What is in place to stop someone from accessing your fingerprint that is stored on a laptop that has a fingerprint scanner on it, especially when a lot of machines have been compromised with trojans and so on? I imagine it could be detrimental if someone got your fingerprint, especially with it being used as a common method of security. Well, this is kind of the same question, isn't it.

Steve: Well, it's a little bit related. He's asking about, if it's stored in the laptop, what prevents it. And I love this because it exactly speaks to last week's topic of the Trusted Platform Module, the TPM, the idea being that that sensitive stuff is unavailable to any sort of external theft. It cannot be stolen because our favorite functions, our one-way hashing functions are used to say, here's a fingerprint I'm submitting for comparison to the one that you have stored. Is it the same? And by taking the one that's stored, hashing it through a one-way function into a hash, and the one you've just swiped on the sensor, and similarly hashing it down, you then compare them and say, oh, look, same guy, here's the password associated with it, which then it provides to the BIOS to unlock the BIOS and the hard drive. And at no point could any of that be stolen. It's very cool stuff.

Leo: You're safe.

Steve: Indeed.

Leo: You're protected. And we want to thank you all for writing. We appreciate it.

Well, we've completed a hundred episodes, Steve Gibson. Amazing, amazing stuff.

Steve: Yeah, I'm really glad, Leo.

Leo: And then he passes out dead.

Steve: We're on for 101 next week.

Leo: 101, and what are we going to cover next week?

Steve: I think we're going to get to CAPTCHA. I think that's a really interesting question, proving that you're human. And boy, that audio that you played last week, you know, I'm not even sure I'm human if I have to listen to that and figure out what that was. Wow.

Leo: Should be a lot of fun. We thank you all so much for listening. We thank you for your donations, which starting from day one have kept this podcast afloat. We also thank the great folks at AOL Radio who have been giving us bandwidth since day one. I should add it up someday, but it is hundreds of thousands of dollars, literally, in bandwidth absolutely donated from them. So we really than AOL for 100 episodes.

Steve: And you know, given that we're on our 100th episode, thank you. I want to thank, I mean, I really want to deliberately and explicitly thank our listeners for supporting me and my efforts at GRC with the purchase of SpinRite. I hear from people all the time who say, hey, I bought a copy of SpinRite because I want to support what you're doing. And, you know, then my hard drive died, and boy did I need it. So again, I don't want anyone to spend $89 just for no reason or for support. But it can come in handy. But I just really wanted to say thank you to our listeners on the occasion of this 100th episode. I absolutely acknowledge and I feel the support because it's had a clear effect on SpinRite sales. And so...

Leo: We love doing it. We love doing it, and that's really the bottom line.

Steve: Yeah.

Leo: Thank you, Steve. We'll talk again next week.