SECURITY NOW!

Transcript of Episode #190

Listener Feedback #63

Description: Steve and Leo discuss the week's major security events and discuss questions and comments from listeners of 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 we have previously discussed.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-190.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-190-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 190 for April 2, 2009: Your questions, Steve's answers #63. This show is brought to you by listeners like you and your contributions. We couldn't do it without you. Thanks so much.

It's time for Security Now! on the most important security day of the year, April Fools.

Steve Gibson: And isn't it perfect that I sound like crap today?

Leo: I am Leo Laporte, and that guy, way distant far away...

Steve: Hello, Leo.

Leo: ...is Steve Gibson.

Steve: Earth to Leo, Earth to Leo, come in, Leo [crackle]. Over.

Leo: Steve's on the phone today. He's using his PDP-8 for Skype, and that's what they sound like.

Steve: Yeah, I thought I'd have more time to work out the code. But we seem to have a problem here.

Leo: That's what happens when you use a 12-bit computer to do an 8-bit job. But it's a big day. We're recording on April 1st. Of course this will air on April 2nd. But we've got to talk about Conficker. We've got to talk about that GhostNet, the giant spy network created by some foreign government. There's lots to talk about. Plus we've got 12 great questions and answers.

Steve: Yup, and another bad problem has been found in Windows kernel.

Leo: Good lord.

Steve: Yeah. It just never ends.

Leo: When will it all end? So Steve, I don't know where to begin here. This is the day that Conficker was supposed to phone home.

Steve: Well, actually, yes. What happens is on April 1st the security analysis in the industry, looking at the Conficker code, also known as Downadup - and this is actually - so far we've had three variations of Conficker, the so-called .A, .B, and .C. The .C most recent variation has been - there's been this cat-and-mouse game with the Conficker control masters or bot masters and the security industry. The key for keeping a worm or botnet - and Conficker is both - essentially alive is allowing it somehow to update itself and to avoid the authorities. So one of the reasons that this particular worm is doing so well, if you want to put it that way, is that it's able to phone home to get updates to itself.

And normally what happens is, for example, a couple websites will be determined to be, like, control points. And those will be shut down, thus cutting the worm off from any updates, or in some cases from control. Well, the earlier versions of Conficker were doing something new; and that is, they were using a pseudorandom algorithm to choose domain names which would extend all the way into the future. And so the idea would be that, instead of it having a few domain names hardwired into the code, that people could reverse-engineer and claim, essentially register preemptively. This thing just keeps generating them.

Leo: It's so smart, so clever.

Steve: Oh, it is, unfortunately.

Leo: Yeah.

Steve: I wish these guys would get a real job, in which case...

Leo: Hey, I bet this job pays pretty well. I mean, we don't know - do we know what they're up to really?

Steve: Well, we know that, I mean, it's a botnet. It spams. It attacks. It also infects. So what's happened is on April 1st, instead of generating and checking 250, a random set of 250 constantly moving target domains, on April 1st today, or yesterday if you're listening to this on our release date for the podcast of Thursday, that number jumps to 50,000.

Now, that means that the worm will be randomly checking one of 50, or many of 50,000 different sites per day, which makes it, I mean, really difficult for the security guys to preemptively block, to register and block those sites. All they have to miss is one. That is, if the bad guys choose one out of the known 50,000 that the worm will try to contact, then they're there, and some of the worms will get there. Conficker also has a peer-to-peer technology that means that they don't all have to reach the mothership. They've also formed an interlocking network among themselves. So if only one gets updated, it's able to update its peers. I mean, it's beautifully designed to survive. And so far it has done that.

Leo: It's pretty amazing. Now, did this - today was the day it was supposed to go get variant D. And the last I saw, it hasn't.

Steve: I have not been looking at it today. I've been producing this podcast all morning.

Leo: I've been watching. In fact, I started watching because, you know, it's so funny because it's very hard to sort through this because even the tech press seems to be completely incompetent when it comes to analyzing and understanding this stuff. So you're getting all these variants. One magazine said it's going to be at midnight, a rolling midnight across the world, which doesn't actually - I don't - maybe that's what it's doing. Another one said midnight GMT. So I started looking around 10:00 p.m., which is about three hours in. And the servers had gone live. The Confickers had picked the 500 servers they were going to use. But no data had been handed off yet. And so I suspect that what's happening is they're - it doesn't have to be now.

Steve: No, no.

Leo: It could be any time; right?

Steve: Right. Right. Essentially, the code that's already in place, which is the C variant, that code that's in place changes its behavior on April Fools Day of 2009 to dramatically expand its - basically its potential target domains. Now, there is still a lot more B variant than there is C. There weren't that many updates to C. It's still significant, but there's still a lot of B. The B variant doesn't do this. It's only the C variant of Conficker which changes its behavior on April Fools Day. So, and it's funny, too, because, I mean, the popular press has been - anytime something malicious has a trigger date, that sort of catalyzes the press. It's like, oh, you know, we can talk about this on March 31st.

And I've had people, my regular normal friends, sending me email this morning, oh, should I turn my computer off? Do I have to worry about anything because of this April Fools Day thing? It's like, no. I mean, this doesn't directly affect people. This changes the behavior of something which, now, moving forward, it can be much more difficult potentially to block this. But nothing happens, as far as we know, specifically today. It's just that the behavior changes today.

Leo: Well, I guess we'll just keep up to date on what's going on and fill you in.

Steve: Yup, yup.

Leo: Once again, though, I'm just impressed by the technical skills these guys have shown. It's just...

Steve: Well, you know, it's a cat-and-mouse game. And so smart people against smart people. I mean, these are smart people. And so as the security industry has come up with and gotten better at blocking the communications, the command-and-control channel of previous botnets, the botnet authors have scratched their heads and said, okay, how do we get around this problem? And then that one gets fixed. Okay, how do we get around that problem? And so...

Leo: That's why this payload thing is smart because, as new problems arise, you can update your virus.

Steve: Yeah, yeah.

Leo: Not to give them any props whatsoever.

Steve: In other happy news...

Leo: Yes?

Steve: ...we do have a new bad kernel integer overflow that's been found in only XP. So Vista people don't need worry, nor Windows 7. But it's across all of XP. It's in the gdiplus.dll, and it's the enhanced metafile, the EMF. In fact, the formal name is Microsoft GdiPlus EMF GpFont.SetData Integer Overflow.

Leo: Oh, boy.

Steve: What it all means is that, if a maliciously crafted image can somehow be shown on your computer, that's a takeover event. So put it on a web page, you visit the web page, your machine is compromised. Send somebody an email. If they're using, for example, Outlook with the - I was going to say with the preview pane. But even if you didn't, if you view the email, the act of viewing it displays the image. Your machine is compromised. Or even embedding the malicious image in an Office document, and you open the Office document, bang. And even in a PDF. I mean, anything that displays, that uses the Windows renderer to parse this enhanced metafile image can cause this integer overflow to occur.

I looked at the website for the guys that have completely dissected it. And although there isn't - it's not currently known that this is in the wild. Microsoft knows about it, but there's no patch for it. So now we wait. So now it's a question of who's going to get there first. And here we are at the beginning of April. So the question will be, will Microsoft be able to get something out the door by Patch Tuesday, which is the second Tuesday of the month. And what is this? That will be, well, they have two weeks because it didn't happen this week because this is Wednesday on the 1st. So they've got two weeks. My guess is this is bad enough that we will see a patch for this in two weeks. So there's a two-week window, given that they make it in two weeks, during which time something bad could come out.

Leo: Remember Microsoft saying with the last WMF exploit - actually, remember, actually it was we said - who thought it might be an intentional plant in the WMF.

Steve: Well, that was a...

Leo: They should really dump this code.

Steve: ...very controversial position that I took a little over a year ago. It was that - and remember that Mark Russinovich looked at the code, and he said, uh, this does look like it was in there from a long time ago. It was clever. You could imagine that, like, back before in the days when security was a concern, some guy said, hey, we have an interpreter for metafiles. That's what this whole metafile thing is. It's basically - it's a little interpretive language where it's like, move here, draw line here, put cursor here, draw circle of this radius. It's an interpretive language. That's what a Windows metafile is. And you can imagine the developer saying, you know, what if we wanted sort of like an escape hatch where we could actually run native code in the image file, not just interpretive code?

And way before the Internet, I mean, Windows metafiles were there in Windows 1.0. It's an original core component. And so security, no one had a concept of a malicious image back then, or the idea of communicating it. So it was like, that's sort of a harmless extra feature that may never be used. Well, looking closely at the metafile interpreter, the hackers said, hey, we know how to use that. So anyway, that's the history of that. I do think it was intentional. I don't think it was an intentional backdoor. I think it was an intentional feature that Microsoft forgot to remove, that was never used or required. They just forgot to remove it over the passage of time. Sort of like they forgot to remove raw sockets when they went to XP.

Leo: Whoops.

Steve: Yeah, there's another little problem.

Leo: Now, this one, I mean, this is an integer buffer overflow.

Steve: Yes. This is not a feature. This is something...

Leo: This is bad programming.

Steve: Right. And, I mean, true, it's the programmer's fault. I will, in defense of programmers everywhere, just say that, whoa, this is - it's just so hard to find every possible way that a program can be abused. It's just hard. Programmers look at it in terms of getting it going, not in terms of, oh, how could what's going somehow be made to do the wrong thing? So it's just difficult.

Leo: Yeah. Wow.

Steve: Okay. Also, since we've last spoken, Firefox bumped itself up to 3.0.8. And the Firefox updates seem to be coming a little more often than they used to. This fixes two security vulnerabilities that involve, one of them at least, a malicious XML file. You would have to view a malicious XML file. So it's not something that's super critical. But Firefox is good about updating itself. So I would imagine our listeners have seen their Firefoxes update automatically. I know that I did a few days ago. So that's being dealt with.

Leo: That, by the way, one of the exploits they say they fixed on the Apple version, the Mac version of Firefox, is that instant exploit that was used at Pwn2Own.

Steve: Right.

Leo: I didn't realize it was not just a Safari exploit, but it was a problem in Firefox, too. So they fixed that, yeah.

Steve: Right. And I don't think we've talked ever about a Cisco router update. But there is one. I want to mention it because I know that we've got listeners in IT and who are, like, involved with networking. And I wanted to make sure they knew that Cisco had released the first in a long while update. There's eight updates that address 11 security flaws in the IOS, Cisco's router firmware. So anybody who's maintaining and keeping Cisco routers up will want to make sure that they're aware of that.

And then finally, on the security front, news came out this week about an interesting network, basically a spy surveillance network which has been named GhostNet by its discoverers. It is the topic of next week's Security Now! podcast. An old buddy of mine, John Markoff, who used to be at InfoWorld, he's now writing for The New York Times. And he wrote an article, I'll read a little bit at the beginning of it because it gives our listeners a quick snapshot of this. The title was "Vast Spy System Loots Computers in 103 Countries."

"A vast electronic spying operation has infiltrated computers and has stolen documents from hundreds of government and private offices around the world, including those of the Dalai Lama, Canadian researchers have concluded.

"In a report to be issued this weekend, the researchers said that the system was being controlled from computers based almost exclusively in China, but that they could not say conclusively that the Chinese government was involved.

"The researchers, who are based at the Munk Center for International Studies at the University of Toronto, had been asked by the office of the Dalai Lama, the exiled Tibetan leader whom China regularly denounces, to examine its computers" - that is, the Dalai Lama's computers - "for signs of malicious software, or malware.

"Their sleuthing opened a window into a broader operation that, in less than two years, has infiltrated at least 1,295 computers in 103 countries, including many belonging to embassies, foreign ministries and other government offices, as well as the Dalai Lama's Tibetan exile centers in India, Brussels, London and New York.

"The researchers, who have a record of detecting computer espionage, said they believed that in addition to the spying on the Dalai Lama, the system, which they called GhostNet, was focused on the governments of South Asian and Southeast Asian countries.

"Intelligence analysts say many governments, including those of China, Russia and the United States, and other parties use sophisticated computer programs to covertly gather information."

Leo: An interesting point. We kind of knew that these kinds of things were going on. Didn't we?

Steve: We did. But what's really cool, and the reason I want to give this next week's episode, is that their report is beautifully written, conservative, no hyperbole. It's very nicely written. It's extremely comprehensive. And so I'm going to absorb it all and distill it for our listeners and really sort of - I think we'll have a great episode next week talking about an instance of this. I mean, sure, we all sort of presume it's going on.

Leo: And we're probably doing this. I would hope we're - in fact, I'd be disappointed if we're not doing the same thing.

Steve: It still strikes me as sci-fi, this notion of cyber warfare. But I guess you need to take it seriously.

Leo: Well, it'd be imprudent...

Steve: Yes.

Leo: And I have to figure that for at least five years the NSA and others have been working on this kind of thing.

Steve: Yeah. Oh, one hopes and presumes.

Leo: Yeah. I would hope our hackers would be every bit as good as theirs.

Steve: Right, exactly. Okay, so that's next week. Errata, or sort of errata: Every so often, as our listeners know, I stumble upon something that I think is neat. I ran across an interesting add-on for Firefox that I wanted to share, which may suit some people and may not others. The way I run Firefox as my browser is, as I'm researching things or running around during the day, I'll use control-click to open another tab. And I use it sort of as a placeholder. Like I'm going to get to that, but I don't want it to distract me right now. Consequently, I end up with a huge number of tabs open. I mean, so much so that Firefox gives up trying to show them all, and I get little scroll arrows on the left and right so that I can move through them. Which is sort of a problem because I'd like to have a better view.

I just stumbled on something, I think it might have been it was suggested to me by Firefox. Maybe it was a new version or update or something. Anyway, it's called Tree Style Tab. And so if you put into your Find Add-ons dialogue in Firefox "Tree Style Tab," it'll take you right to this. What this does is, there's many different ways it can display. What I've done is I've opened up, like, a tab list on the left-hand side of Firefox, so that I've got the whole height of my screen now to show me my tabs. And when you're on a page, and you shift-click to open another tab from a link, it indents it in outline style so you can see the parentage and trace back the relationship.

So anyway, I'm still - I've only had it for a day. So I still am - my instincts are to go up to the top where tabs used to be. And I'm not yet retrained. But already I can see so many more tabs. And the hierarchical representation is something that I really appreciate. So it's funny because I told a friend of mine, and he said, oh, no, I just - I only have - I don't like multiple tabs. I just have one thing at a time. It's like, okay, this is not for you.

Leo: No, no.

Steve: So I recognize this may not be for many of our listeners. But if there are people like me who end up with, like, with just tab insanity, this thing really looks like it's going to be a great solution for that.

Leo: I think power users and techies use a lot of tabs. Kevin Rose two weeks ago on TWiT was talking about that. And he said - he just said, okay, how many tabs do each of you have open? And it was an average, I think, of 17 or 18.

Steve: Yeah.

Leo: It's very common.

Steve: And oh, one very cool thing about this, I mean, this thing's got features coming out of its ears. It's got more features than I've talked about. But, for example, you can protect tabs, which I really like because, for example, I'm still keeping a track of DEC PDP things on eBay. So my very first tab for a long time has been a tab open to my eBay page. My second one is something that I use for kind of keeping track of the stock market. It's a nice little stock market ticker viewer. Well, I've gotten used to them being there. But every so often I'll delete them by mistake. It's not a big problem to reopen it. But this allows you to protect the tabs, essentially locking them in where they are and to their page. So then you just hit refresh every so often. So anyway, it's got a ton of features. If people, like, organize their lives through their browser based on tabs, I wanted to let people know about Tree Style Tab.

Leo: I'm installing it now.

Steve: It's cool.

Leo: Sounds like something I'd use a lot.

Steve: In going through the mailbag I ran across a number of people who were a little despondent at being behind in Security Now! episodes, catching up, hearing about the PDP-8 kit that was made available several months ago.

Leo: Yeah.

Steve: And they're still available.

Leo: Oh, good.

Steve: So I wanted to - I just wanted to let people know not to be despondent. All of the first round of them were made and sent. I've got my three. And I'm not parting with any, so don't bother asking.

Leo: No begging.

Steve: But there are still some available, both the full kit for the board itself, and for the front panel, and a bag of add-on parts for the front panel. So I just wanted to let any listeners know who were saying, gee, I'm really sorry I missed being up to speed on that because I would have loved to do that. It's SpareTimeGizmos.com is the site. SpareTimeGizmos.com. Or you can probably search for SBC, as in single-board computer, SBC 6120. That's the name of the chip, which is a single-chip PDP-8 this kit is based around.

Leo: Some people are reporting in our chatroom - I think everybody went off to download that Tree Style Tabs. And some people are reporting the same experience I had, which is - and I occasionally get this with Firefox add-ons, an error saying the CRC is inaccurate or something. And just my tip to them, this has happened to me before with these beta - it usually is with beta add-ins. Right-click and download it, and you'll get a .xpi file which you can then open directly in Firefox, and you won't get that same error. I don't know why that happens. I think it's either a bug in Firefox, or it has to do with signing or something like that.

Steve: And maybe it's that the site's busy? Could it be that...

Leo: Could be we killed the site.

Steve: Because it worked for me just, I mean, I just downloaded it directly into Firefox in the normal way.

Leo: It might be a Mac thing. I don't know. But anyway, I've installed it now, and I'm loving it.

Steve: Oh, it's a good thing.

Leo: Yeah, yeah.

Steve: A couple people asked, with regard to the multifunction YubiKey, whether or how they could get the newer version. I sent email off to Stina, which I received a reply to. Apparently they're going to be using one of the major Swiss or Swedish, I wasn't sure which, automotive key manufacturers to make their next round of keys. So they're probably ramping up production. It's not going to be available until after the summer. And they are going to be offering a discount to existing customers. So it won't - there's no real way to upgrade your key or to trade the single function for the dual function. But at least people who have keys will be able to get them at a better price.

Leo: Yeah, great.

Steve: And then one bizarre thing that just sort of crossed my radar. I was reading in the SANS security list, their newsletter, somebody was talking about - he was actually attending a security conference or event down in San Diego. And he received in his phone bill an excessive roaming charge from Verizon for his EVDO, which is what he uses to get on the Internet, of $199.

Leo: Whoa.

Steve: And it turns out - get a load of this. He talked to Verizon. They said they would remove it from the bill because he said, you know - the point was he was close to the Mexican border. And the Mexican cell towers are not that busy and have high signal strength. And so his card was captured by an out-of-country tower, which caused him, even though he wasn't out of country, to be hit with super-high roaming charges. And apparently the problem is even worse near the U.S./Canadian border. So I thought that was just, like, the kind of thing you never expect or think about. But it's like, whoops, gotcha.

Leo: Oh, man. Did he appeal it, and will they give him his money back?

Steve: Yes. They said - they did remove it from his bill, they said, but just this once. So his advice was be careful about what you're being captured by in terms of cell usage. And it may not be safe to use something that's got high roaming charges near a border point where you're not sure because there's no obvious demarcation. It's all automatic.

Leo: Yeah, it should - I guess they don't want to bug you, but they should warn you if they're going to do that.

Steve: They really ought to, yes.

Leo: You can turn off data roaming in most phones, which is probably not a bad idea if you're going to be near the Mexican border.

Steve: Oh, so that you would deny any non-local carrier.

Leo: Data roaming can be very expensive.

Steve: Right.

Leo: Basically what that says is only use Verizon for your 3G data.

Steve: Right.

Leo: And that's probably, you know, a good idea, prudent for most people, anyway, yeah.

Steve: Well, we're half an hour in. And I've got a SpinRite testimonial, but so many people in our Q&A mentioned SpinRite this week that I'm going to skip the testimonial. We'll do it next week.

Leo: Who needs a testimonial, my god. Everybody knows SpinRite's the one.

Steve: Well, this is really well written and really neat. But we'll do it next week.

Leo: All right. Yeah, we've got some great questions, 12 questions, coming up in just a second for people who - the questions are really our chance to talk back to Steve and get clarification. And Steve's so great at answering these questions and explaining what's going on, which is really great. So we've got questions. Let's get to them, lots of them, for Mr. Steve Gibson. Starting with Michael, an expat. He's using something called TOR, The Onion Router, and he wants more: Hi, Steve and Leo. Thanks for the great show. I'm an expatriate living in Southeast Asia, and I like to use TOR to visit some areas of the web. I understand why, you know, he wants some privacy.

As I have some interests in U.S. and European businesses that often restrict their services based on the IP address's incoming location, so I've configured my tor.rc to only exit through certain U.S. or European nodes, and not to have the exit node hop around. Oh, I didn't know you could do that. That's cute.

Steve: Yeah, isn't that neat? A nice little tip in there, too.

Leo: So far, so good, as long as I'm only interested in anonymity, that is, in sites that don't need a login.

The problem comes when I try to use TOR's features with sites or services that require a login. Seems to me there is no way to know if I have picked a compromised exit node - and this is really an important phrase, "compromised exit node" - that can sniff, log, or otherwise misappropriate

my credentials.

Steve, I'd love it if you would run a TOR server. I feel after years of listening now that your node would be trustworthy. I think we know that. I'd be sure to make your server the exit node of choice in my tor.rc file. What do you think? Would you consider adding a server to the TOR network?

Steve: No.

Leo: Yeah. I wouldn't, either.

Steve: Well, and here's why, and here's the problem, is that I care passionately about Michael's characterization of me being trustworthy.

Leo: Yes.

Steve: I mean, I would absolutely myself never in a gazillion years consider taking advantage of that trust. But I don't have control of the traffic once it leaves the wire in my rack at Level 3. And that's why I thought this was an interesting question, and one that we needed to discuss, is that it's certainly the case that you might have untrustworthy TOR exit points, meaning that there's, I mean, inherent in the TOR network is this notion of aggregation of traffic. There are only so many TOR exit points. And everyone using the TOR network has all of their traffic jumping from one server to the next.

And if our listeners, if we have any listeners who don't remember about TOR or don't know what this is, we did a beautiful podcast on it, how we explained why it's called The Onion Router and how each node only knows enough to take one layer of encryption off of the data and then forwards it on to the next so that it really creates a very strongly encrypted, anonymous system, and even when it comes to, like, routing the data between TOR servers.

Ultimately, though, you have to leave the network at a so-called "exit node." At that point the last, the innermost wrapper of encryption is removed by the exit node TOR server, and your traffic goes onto the Internet without any encryption capsulation. So you can imagine that anybody, any entity, whether it's governments or malicious people, not to say malicious governments, but anybody who's interested in the kind of traffic that people might want to anonymize, you might imagine that the traffic entering and exiting TOR nodes is more interesting than just the random sea of traffic on the Internet because there is a presumption that there's something, there's some reason that people want to have anonymity and the privacy that TOR potentially creates.

The point is that anyone operating such a node, like myself, can control their own node, that is, can say, okay, I know that I've made it secure, it's got its own firewall protecting it. I don't care what data comes in or out, I mean, I've got 20-plus years of history of dealing with hard drive data. And every so often, more in the old days, people would send a drive that they desperately needed repaired, and I'd sort of do it as a favor. Well, as a matter of honor I never looked at anyone's data. I don't care about that. I just - I like the idea of being able to fix it and being able to send it back. So similarly, I don't care about what is going on, if I were to run a TOR node.

But the point is, ultimately I'm connected to my provider by some wire. And it goes to a router somewhere. And off it goes. At that point I've lost all control. If I had a TOR node, and it was an exit point, then all the client traffic has been decrypted as it left. And if, for example, some government were to subpoena an eavesdropping operation on that wire, I would never know. So with the best of intentions I would - basically the trust that had been placed in me would, through no fault of my own, be subverted. So that's just not something I would want to do.

Leo: I'm going to also refer you to a blog post that I read recently from - it's calumog.wordpress.com. His post is: "Why you need balls of steel to operate a TOR exit node."

Steve: Huh, interesting.

Leo: And he says, "I totally believe in TOR," as we do. "I think it's a magnificent force for the circumvention of Internet censorship. But there's a problem. I was visited by police in November 2008 because my IP address had turned up in the server logs of a site offering, or perhaps trading in, child pornography. The date of the offense was one month after I started the server."

So it looks like the site in question had been under surveillance for more than a year. The police made what's known as a dawn raid. They threatened to burst down the door. They had never heard of TOR. They had no idea what he was up to. But just the fact that somebody had used TOR to access that server, he was the exit node, so he was the node of record, implicated him. They took his computer, went through it forensically. He was never charged, fortunately. But that's the risk you run. You don't know what people are using it for.

Steve: Right. Well, and here's Michael, who's living in Southeast Asia. He'd like to have access to services that are only available to U.S. and European IP addresses. So it's one of the things that TOR does is allow you to anonymize yourself and lock your exit node to specific IPs, which is very convenient. So it's a tremendous service. But as you say, Leo, it comes with a great risk to those who are running those exit nodes.

Leo: Yeah. And we thank them for doing that.

Steve: Yes.

Leo: And we encourage you not to do anything illegal on those servers. Listener Fred says: What's an HTTPS scanning server? Steve, I especially enjoyed your series on HTTPS, secure HTTP. But I have a question about something called an HTTPS Scanning Server. When I log onto my corporate domain every morning, I have to click through a pop-up window agreeing to be monitored. The text of the pop-up says, "Users are subject to monitoring at any time, including accessing HTTPS websites using iAccess. HTTPS scanning servers decrypt all HTTPS traffic." I am curious about their claim to be able to decrypt HTTPS traffic. I don't object to this on my corporate domain, but it begs the question, can it really be done? And what's to prevent any other server from doing it? We've talked a little bit about this in the past.

Steve: We have. And I don't want to belabor the point. But I wanted, first of all, Fred says he doesn't object to this on his corporate domain.

Leo: Doesn't matter.

Steve: I want to make sure that he understands this is not just accessing his corporate domain, but accessing anything outside his corporate domain that is being scanned. So the idea is that, first of all, I did like this question because it demonstrates a concrete instance of this happening. We've talked about it. I talked about having gone to RSA and talking to a company that is offering these services. I wanted to make sure, because I think it's important, that our listeners really get it that HTTPS can be decrypted on the fly if your system has been configured in advance to allow that.

So Fred's question says, I don't object to this on my corporate domain, but what's to prevent any other server from doing it? Well, it's the configuration, the use of this so-called iAccess approach, where all of the clients in the corporation are accepting a certificate from the gateway which allows a secure - essentially allows the gateway to function as an impersonation of the remote server. Essentially you're saying that your gateway is trusted like a so-called certificate authority. And it's that, like a root authority, so the gateway is able to essentially sign the certificates of sites that you think you're visiting. In fact, you're visiting the gateway. Your traffic is decrypted there, analyzed, and then reencrypted for its transit across the Internet. So I just wanted to - I liked this because it says this really does happen. Here's an example of it. I'm impressed that every time, every morning when he logs on, he's reminded...

Leo: That's a good thing. Because they don't have to do that.

Steve: No, they don't. And I did like that about it very much, that it's like, okay, we're going to - you're going to say "Yes, I agree," every single morning at the start of the day. And we're going to remind you that this is what's going on.

Leo: I always encourage - first of all, employees have to remember they're using company equipment. And so the courts have consistently ruled that companies have the right to do any kind of monitoring they want. But I always encourage - without notice. But I always encourage companies to make a written policy, to post it, and to tell employees, regularly inform employees of what they're doing. That probably has a good kind of preventative effect anyway.

Steve: Well, and I've recommended that it be put on a strip like a piece - sort of like a Scotch tape across the top of the monitor so it's just in front of you all the time. Because it's been well established that companies have the right to do this. However, the psychological effect of learning about it when you didn't know, I mean, that's what causes people, employees to, like, drive their cars through the front office of the company and really become outraged. I'm reminded of the original discovery of adware from that company Ad-Aware. And it was - that's when I coined the term "spyware" and wrote that first little antispyware gizmo Opt-Out back in the old days. And people were phenomenally upset. I mean, even though the fine print, they had agreed, that they didn't really know this was going on really upset them. And so you don't want to surprise people this way, so I just say it makes so much more sense to be upfront.

Leo: I agree. I agree. And it has a good deterrent effect, as well. Jack Jensen, Tampa, Florida, says "I have unwanted company." He says: I'm trying to get help. I'm hearing the sounds of - this is the kind of question I get on the radio show, by the way, Steve, all the time. I'm hearing the sounds of mouse clicks and keyboard typing, not mine, coming from my speakers. Even with the browser closed. I tried Spybot, Malabyte, SUPERAntiSpyware, CCleaner. Running free of Comodo Firewall. I do have snapshots of denials of Firefox requests, while I'm using

it, and also of my active connections, if that would help, and I could attach it. Thanks, and HELP!

Steve: Yeah. As you said...

Leo: What's going on? Because I'm curious.

Steve: Well, I don't know specifically. But if this is going on, and Firefox is denying things or is acting oddly, it really sounds like he's not in control of his own machine.

Leo: He's hearing typing and mouse clicking coming from his speakers.

Steve: Yeah. There's something...

Leo: That's bizarre.

Steve: ...very wrong in his system. He's obviously tried all kinds of antispyware stuff. The one thing that I would recommend is, due to the nature of contemporary malware, we've got this problem that we've talked about often called "rootkits." Once something gets into your computer, it can be extremely difficult to see it because you're using the operating system and trusting the operating system to do the seeing for you. That is, anything you run on your computer is a client of the operating system, which uses the operating system's services. I mean, even doing a directory listing, you ask the operating system for a list of the files in a directory.

Well, if the operating system itself is compromised, that is, something has crawled in underneath it, then that something can filter out the response. We saw this famously with the Sony DRM, Sony's Digital Rights Management, that installed a rootkit which hooked the response to directory listings and removed itself from those listings. So no matter how much you tried to look, you couldn't see what was really there. The only way to deal with this is to take the - well, there are many ways to deal with it. In my opinion, the most straightforward and in some ways easiest way is to take the drive out of the machine and make it a data drive of another machine, and then scan that drive. You don't want to run anything on this data drive because you don't want anything to have a chance to get off of this infected drive.

But the problem is all of these scanners are scanning after the infection has taken hold. So they may not be able to see it. I did this just a couple weeks ago with a friend's laptop that was infected and found everything on it by catching it as a data drive and running the scanner on a good machine against that data drive. And it was able to find it. And there was no chance for the bad stuff to get in there and prevent it from being seen.

Leo: Yeah. I'm wondering, maybe, could it be just kind of RF leaking into the...

Steve: I was thinking the same thing when he - but it sounds like it's autonomous, like it's happening when he's - when it's not he who is clicking things. And he also says that Firefox is...

Leo: That's the thing that's worrisome.

Steve: Yes.

Leo: Those certificates or whatever.

Steve: Denials of - so denials of Firefox requests. So Firefox is trying to do something.

Leo: Well, this was a symptom of Conficker, a very strong symptom of Conficker, because one of the things I think all variants do is block you from going to antivirus sites.

Steve: Yes, I was just going to say, if you hadn't said it, that I forgot to mention when we were talking about Conficker, people have been wondering how to easily determine whether they've got it. Well, try to go to Microsoft.com...

Leo: There you go.

Steve: ...SANS.org or Symantec.com because Conficker blocks at least those three and more. And your computer just will not go there. So if you find that your machine will not go to SANS.org, then game over.

Leo: It also won't allow you to do Windows Updates.

Steve: Right.

Leo: Yeah. Game over, man. You know, I think, given that we know that there's probably somewhere between 9 and 15 million computers that have Conficker on them, that it's a very good likelihood that people listening to this show don't have it.

Steve: I think that's true.

Leo: But that they know people who do. So it would be the good thing to do, the Good Samaritan thing to do to talk to your less sophisticated computer friends about this and spread the word.

Steve: Yes.

Leo: Paul Harding in Calgary, Alberta, Canada wonders about external drive recovery. Well, you've got the right guy. Steve, I have a question you might want to address for all SpinRite potential customers. Although I personally don't own SpinRite yet, I have heard now hundreds of testimonies that praise the wonders of SpinRite. I have no doubt that SpinRite is a fantastic product. My question is in regards to the usefulness of SpinRite on external USB, eSATA drives and NAS drives. Currently on my computer I have 5.6 TB of storage. You know, a few years ago that would seem out of, you know, out of control.

Steve: Oh, mainframes didn't have that much storage.

Leo: Yeah. But now it's like, eh.

Steve: Okay, double drive.

Leo: Eh, big deal.

Steve: I got that in my bathroom.

Leo: I just ordered 12 gigabyte drives for our NASes, to upgrade our NASes. It was a terabyte, I mean, 12 terabyte drives. It was a terabyte, three [indiscernible]. We were going to upgrade the NASes. And a terabyte is kind of - it's $100. It's the default size now.

Steve: You know, I'll just interject here, I had Mark Thompson visiting twice over the last couple weeks. He was driving through. He went up to the Game Developers Conference, the GDC up in...

Leo: Oh, I wish I'd known, I would have loved him to visit.

Steve: Yeah, well, he was with a bunch of friends and back-to-back meetings.

Leo: Oh, okay.

Steve: He's got all kinds of things going on.

Leo: He's busy, yeah.

Steve: So he's super busy. But he happened to mention that he's using Western Digital drives. And I remember the days when nothing could make me use a Western Digital drive.

Leo: I know that's changed, though, hasn't it.

Steve: And that's my point. I wanted to give them real props. And I know - I heard you speaking, maybe it was to Andy, about some, like the 1 GB Black - the Black Caviar...

Leo: Caviar Black, yeah, that was to Ryan Shrout we were talking about those.

Steve: Right. And so I just wanted to say that, I mean, I know that there are - it is a moving target. And it's the case that a company that, I mean, hard drive storage is on the edge. It's always on the edge. Because if they could fit any more data in there, they would. And so they do. Always to keep themselves on the edge, being as competitive as they can. But so as a consequence it's sort of - you can go through a bad spot, a rough spot where your process just isn't nailing it down, or that it's got problems after a few months. And after being burned by a succession of WD drives, I swore them off. But this was 15 years ago. On the other hand, those memories are slow to die.

Leo: They are.

Steve: I think it's clear that this memory should be dead, and that I need to give Western Digital another look.

Leo: You know, Egghead had a deal of $119 for 1.5 TB Seagates.

Steve: Wow.

Leo: But I was a little...

Steve: Oh, Seagates.

Leo: I was a little slow to buy the Seagates. And I said you know what, I'm going to buy the Western Digital Caviar Greens instead.

Steve: Yup, I think they're now the - I think they're the sweet spot.

Leo: Yeah, yeah. And the Greens are good because they run a little bit slower, cooler, lower energy, and perfect for a NAS.

Steve: Yes.

Leo: We're going to put four in Drobos - four in a Drobo because we're going to give that away. And we're going to put four in the NAS, and we're going to use the other four for recording the shows. I've been using Hitachi. I know you like the Hitachis.

Steve: I do, very much.

Leo: They bought the Deskstar line.

Steve: Yes, from IBM.

Leo: Yeah. Let's see. Moving along. Oh, we're continuing to answer a question. I forgot. We haven't even gotten to the question yet. I know a little extreme for a home user, he says about the 5.6 TB. However, I edit video and regularly max out my storage space. 1.5 TB of this space is completely external drives, 1.5 TB in a NAS. Can SpinRite maintenance be used effectively on the USB or the NAS? Thank you for your informative podcasts. I've been a loyal listener since the beginning. And as an aside, I have a great story about how your podcast helped to catch a criminal. What?

Steve: Mmm.

Leo: I'm waiting for the trial to begin and will send you the story once I am not in a place that could compromise the case. By the way, you should do an episode on computer forensics. I would love that if you did that, Steve.

Steve: We'll do it.

Leo: Yeah.

Steve: Relative to SpinRite - and maybe other data recovery utilities, I really can't speak authoritatively about anything other than SpinRite - there really is a difference between external serial interface like - traditional serial interface like USB and Firewire, or network and eSATA. What I mean is that eSATA is essentially the same, the external SATA, serial ATA, that is just like the ATA interface, but it's been serialized. So there's the same total access to the hard drive's guts for eSATA, SATA, and PATA, the parallel ATA. Those are the best way of letting SpinRite have access to the drive, is anything that's ATA, because there's a whole vocabulary of commands which are not about transferring data in and out, but are about the inner workings of the drive, which give SpinRite far more intimate connection to the drive. It's able to do things that it cannot do if you are over a network, where basically your API, so to speak, your interface is read this block of sectors and write that block of sectors. That's pretty much all the remote coupled or the USB and the Firewire interfaces allow is read and write data. They do not allow nearly the same level of recovery and intimacy.

So we recommend to anyone who really wants to maintain their drives, or especially data recovery, it is worth taking the drive out of that configuration, sticking it onto a motherboard, and running SpinRite against it there. The results, I mean, we hear also people doing data recovery remotely through USB and Firewire. It can work. But if it were me, I would always go to the trouble, if this was important data, to give SpinRite the best connection it can have for the drive.

Leo: Yeah. I think I told somebody that on the radio show this weekend that, yeah, you want an internal drive. It's the best.

Steve: My neighbor said - she was out watering the lawn yesterday. She said, hey, I hear you have six Palm Treos, or Palm Pilots.

Leo: Not the Pilots. I think I said the T, what was it, the TXes; right?

Steve: Yeah. Yes.

Leo: Yeah, I think I said the TXes.

Steve: I said, what? How could you possibly know what my Palm Pilot count is?

Leo: Here's what happened. A guy called up, was furious. He had a Palm Pilot. And basically Palm, it's over for Palm, that Palm OS, and they're going to go to their webOS. In fact, the jury's still out whether that'll be enough to keep them alive. And he said, but what about us? I love my Pilot. And I said, well, I have a friend, Steve Gibson, who felt the same way, loved the TX so much he bought a half dozen of them and put them in the freezer.

Steve: Yup, they're still there. They're in the refrigerator, just waiting in case I ever need them.

Leo: [Laughing] Now you're looking pretty smart, Mr. Gibson. I love that. I use you as an example fairly frequently on the radio show.

Steve: Of over-the-top extremism.

Leo: Yeah, the geek, the real, the true - what a true geek is all about. Mike Nicklin in Eureka, California writes - he's got three questions. One, do you accept cookies? Two - this is like a congressional hearing - do you worry about them? Three, should cookies be accepted just to keep the hassle down? Three questions, three quick answers. So let me give you number one. Do you accept cookies, Steve Gibson? The world wants to know.

Steve: I accept only first-party cookies, never third-party cookies.

Leo: Number two - we'll let you explain further in a minute. Number two, do you worry about cookies?

Steve: No.

Leo: Number three, should cookies be accepted just to keep the hassle down?

Steve: That's not necessary, Senator.

Leo: Thank you. Speak into the microphone, son, and tell us your position on cookies.

Steve: I've been exposed to all the extremes of cookie handling. There are a lot of people who just don't worry about it, don't care. They just say, well, I have no control over what happens on the Internet. All of my privacy and my rights are gone anyway. So I'm not worrying about it. I've got better things to do than worry about it. They're at one extreme. The other extreme is people who, I mean, really focus on cookie management and go through their cookie list and think, where did this come from, and they delete it. Or they, like, set their cookie files to read-only, so their browsers bang on the door, unable to store a cookie there. And then they selectively let them in. I mean, there's all kinds of policies.

I'm much more of a middle-of-the-road cookie person, which is why I don't worry overly. But I take a simple countermeasure, which is just say no to third-party cookies. All browsers allow you to turn them off. Sadly, not all browsers even do that correctly. But that'll be a topic for a future show since I have got a cookie forensics technology now working on the site for some time which does a good job of allowing people to see exactly what their browser is doing.

But for most people, for Mike, I would just say go find the setting to disable third-party cookies. It's in IE, it's in Firefox, it's in Safari. Everybody's got it. And just turn that off. You may want to flush your cookies after you turn it off, and restart your browser because that way you've gotten rid of the debris that you've accumulated. And then don't worry about it because you've really dealt with the major source of tracking. Now, Flash cookies are a different matter, and they're becoming a little more pernicious over time, and we'll be dealing with that, too. But just say no to third-party cookies, and you've really done 99 percent of the work for 1 percent of - you've solved 99 percent of the problem for 1 percent of the work.

Leo: And I have one more question to you. Did you ride here in a private jet? All right. Moving on. Lee, Lee W., in West Milford, New Jersey, refuses to elevate his rights. Steve, I'm a faithful listener of Security Now! since Episode 1 and Leo's other TWiT shows. It has replaced my radio for my two-hour daily commute to work. I have learned a tremendous amount from your shows. I'm also a proud owner of SpinRite. I tell everybody about it. It's saved me and my family members several times. I've been following your discussion about the Microsoft Malicious Software Removal Tool and decided to download this month's release, KB890830. I launched the tool as described by the Microsoft site and then was surprised by the prompt, "You must be logged on as a member of the Administrators group to run the tool."

Sure enough, it does say that on the download page. You and Leo have me trained well for security, and I simply won't use my PC as admin, period. I use a limited user account unless I need to install applications or other items. It just strikes me as funny that Microsoft didn't design a tool to run for limited accounts, where all software should be designed to run for security purposes. Most average users should use a limited account to protect themselves. I did not try it with the automatic updates in the limited account. I just thought I'd share this with you. This is actually a really good point.

I love the show and want to thank you for your hard work and dedication you give to the world by trying to make it a safe and secure Internet. I love your precision - yes, I agree - and great care you have for getting things right. Keep up the great work. P.S.: I, too, use Firefox NoScript plug-in and Block All and run only the main sites as temporary allow. I sometimes enjoy figuring out which scripts allow a page to work.

So I've modified my recommendation, by the way, to run as a limited user because of User Access Control and the way Microsoft works. But let's talk about rights elevation.

Steve: Well, I would say to Lee that I'm really happy he's taken the philosophy and the position he has.

Leo: I was going to say, yes, yes.

Steve: But there are some times when running as an administrator does make sense. It's what you have to do. And he cites when he needs to install software, he needs to run as an administrator. Well, the fact that MSRT will not run as a limited user is really demonstration of the beneficial limitation normally of being a limited user. I mean, you would want it not to be able to run because the whole point of being a limited user is that what you're able to do is limited. But unfortunately this is a trusted, authentic from Microsoft tool that needs to get to the deepest roots of the operating system in order to do its job. By definition, it cannot do that from a limited user account. And you don't want it to be able to do that.

So but philosophically I just - I wanted to suggest that Lee back off a little bit and to recognize that it's not the case that there is no role for administration other than just installing software. There are things, and MSRT is a perfect example, where the nature of it is that its operation requires full access to the system. You don't want your typical user to have that most of the time, but you do want the administrator, the trusted administrator of the system when running trusted software. Then it absolutely makes sense.

Leo: Yeah. And, you know, you can right-click on any application and run it as administrator.

Steve: Yes. You need to provide those credentials. And so it's very easy to do. You don't have to go through logoff/logon process.

Leo: And that would be the preferred way to do it because then you don't forget to log out and so forth.

Steve: Right, right.

Leo: So but I still run as limited user, even though nowadays running as - because of User Access Control, even nowadays you don't really run as administrator, even when you're running as administrator. Right?

Steve: Well, yes. And you know, this is another perfect example of the evolution of security.

Leo: Right.

Steve: Today it is much more practical to run in a non-administrative context than it was when we first started recommending it. I mean, UNIX has always been this way. Windows was never this way.

Leo: Right.

Steve: Windows didn't have this notion. So what happened was that people who tried to use good security practice and be a limited user, they kept running into their software that wouldn't work right that way. It took a while for the software to catch up because - and it took pressure from the users, users saying gosh, you know, I need my email client to run in a limited context. And it was just sort of because that wasn't being done often originally, that programs that did not need administrative rights assumed them. And then they would break when you ran them in a non-administrative context. Well, we've had years of that now, and software has caught up. So that original barrier to running as a limited user is pretty much gone now. And that's a really good thing. But it's another, again, it's an example of this just taking time. Unfortunately this is all evolution.

Leo: Welcome to the wacky world of security.

Steve: Welcome to reality.

Leo: Yeah.

Steve: Yeah.

Leo: Ben Pfountz at Virginia Tech - what a great school - needs to reach out and touch clients. He says: Hi, Steve. Long-time listener here. Keep up the great work. I have always wondered if there is a client-server type of application for administrators that works similarly to what gets installed on a zombie machine.

I'm thinking I'd set up a nice dedicated server somewhere, install a lightweight "presence and remote control" service/daemon on each client machine. The service or daemon would maintain a TCP connection back to the

server, just like a zombie would, updating statistics and allowing me to remote-control the machine.

This would be beneficial to me because I'm finding myself supporting machines in more and more locations, and it's often difficult or impossible to connect to the client machines when necessary because of NATs or firewalls or blocked ports, that kind of thing. Having the client machines maintain a TCP connection to one of my servers would be very nice because it would bypass most of the network issues I've been having. In other words, it establishes an outbound connection and keeps it open.

I was thinking about how this kind of app could be created. Then I remembered you discussing your new VPN application. I know you're currently in the design stage, so I thought, let me throw this idea your way to make sure you put it in. I am the only IT administrator for our department, so this kind of functionality would be great for me. Thanks. That's a really interesting idea. What do you think?

Steve: It's in there.

Leo: It is.

Steve: Yeah. CryptoLink is the forthcoming VPN that he's talking about from GRC. And I was very enamored of the ease of setup that Hamachi offered.

Leo: Right.

Steve: Yet I also heard many people complaining that, when Hamachi's servers went down, the whole Hamachi network went down. So it's nice to have it when it's there. It's a problem to depend upon it if it happens not to be there.

Leo: Well, they were doing kind of a triangle, right, so you would contact the server and - that's kind of what I know GoToMyPC does.

Steve: Right. And so CryptoLink will have a number of different ways to operate because, if you want the super-simple drop-in operation where nothing needs to be configured, I mean nothing, and where both endpoints are behind NAT, then because both are behind NAT routers it's not possible for either endpoint to connect to the other because they can both get out of their own NAT, but they can't get into the other NAT. It'll get blocked.

So what's necessary is for a third party, a so-called "rendezvous server," they'd connect to this rendezvous server. The rendezvous server is able to look at the nature of the packets coming to it and then inform the other side about how they should try connecting in. And so you're able often to knit a cross-dual-NAT connection, but you'd absolutely need that external third party. So the beauty of that is, it just works. No router configuration. It just works.

So I absolutely want my system to have all the features that Hamachi did from that standpoint. But I also wanted not to depend upon that because, trustworthy as I always intend to be, and will be to the degree it's in my control, I don't want to force anyone to trust me beyond trusting that I wrote the code correctly. So another way for CryptoLink to work will be in a so-called TNO, Trust No One, mode where if, for example, someone like Ben knew he wanted inbound connections, he could configure his router at his end to allow incoming connections. Then anybody anywhere, even behind NAT, and independent of firewalls and ISP port blockings and everything, I mean, I'm designing this so that it will always get through, they would then be able to make incoming connections into his network. So it'll definitely, you'll get essentially the best thing that all these approaches have to offer.

Leo: I can't wait.

Steve: I can't either. As soon as I get this other backlog of stuff done, I'm onto it.

Leo: Poojan Wagh in Chicago, Illinois, our question number eight, wonders about password strength meters: I'm wondering how the password strength meters, those bars that go from red to green - some sites have them - depending on the length of your password, work. I've noticed that Google, when signing up for an account, or Yahoo!, or Microsoft, or anything that has an indicator of how good a password is, I'm wondering is there some mathematical formula? If so, what is it? If I type in "7EDAHR7J," Google says it's strong. However, if I type in "01234567," Google says it's fair. Clearly there's more than just length involved.

P.S.: I used SpinRite last week. It was especially timely since all our music and videos were on the computer, and I had just bought my wife a new iPod so she could entertain the kids on a spring break trip. Luckily SpinRite came to the rescue just in time so she has a bunch of music and kids' shows on her iPod for the flight. So what's the story on password strength?

Steve: Well, okay.

Leo: I think they're making it up.

Steve: They are. Which is not to say that's bad. What I like about these meters is that it's a nice means of educating the typical user. Lord knows our listeners, I mean, we spent the first three months of this podcast, Episodes 1 through 12, talking about passwords, I think. So there's nothing that our users, that our listeners don't know about password strength. But many of us though there are, we're certainly the minority. And the idea of rating a password as you type it in, of course that requires some JavaScript to be watching while you're typing. So the JavaScript is running an algorithm to perform a so-called "heuristic." A heuristic is sort of the fancy term for a rule of thumb. And so there is no standard for what makes a good password, or when this goes from red to green.

It's encouraging, for example, that Google was smart enough to know that 01234567 had some problems. You don't know what the algorithm was that had to decide that. It might be special casing, looking for sequential number stream. It might be like looking for a sequential difference between them. The idea is it's sort of anything that the programmers came up with that are sort of guidelines for the obviousness and the guessability of the password could go into this meter and inform the meter how good it should think what's been entered so far looks. You'd like to see upper and lower case. You'd like to see mixed numbers and letters. And you'd like to see length. So my guess is, if you see mixed case, mixed numbers and letters, and sufficient length, then that's a pretty good password. You can write a little bit of JavaScript that will look at those things and rank the password accordingly. And I just think it's good. I think if someone puts in their name or their...

Leo: "Sexy."

Steve: Yeah, exactly. It's going to say, eh, try again.

Leo: [Buzzer sound]

[Talking simultaneously]

Steve: ...nice little closed-loop feedback that I think, again, this is the way we move slowly forward.

Leo: Basically there's no standard algorithm for determining this.

Steve: No.

Leo: But as you say, the most important thing is random. Random is good.

Steve: Oh, random is the best.

Leo: The more random, the longer, the better.

Steve: Entropy. We like entropy in our passwords.

Leo: Entropy. Of course it's also harder to remember.

Steve: Yeah.

Leo: Dain Nilsson, a YubiKing Winner - yay, congratulations, Dain - asks about using hash functions as ciphers: Hi, Steve. First I want to thank you for the kind words you said about KeyGenius and for explaining it so well to your audience. I have one grievance with you, though. When announcing the winners of the YubiKing competition, you stated that I was some guy in Switzerland. I'm Swedish, not Swiss. No hard feelings, though. After all, you liked my entry. And if it weren't for you I never would have heard about the competition in the first place. We have a lot of Swedish listeners. Despite my bad Swedish chef voice. They put up with that.

Anyway, on to my question. On the HMAC episode you mentioned that hash functions are freely available, are not encumbered by intellectual property, and haven't had the export restrictions that ciphers have had. Now, this got me thinking. Couldn't you take a hash function and modify it slightly to use it to encrypt and decrypt data? I mean, I came up with a scheme. I'm curious to see if you think it would work or, rather, if there are any security vulnerabilities with it. So here's what I came up with. Oh, boy.

We know that, if we have a pseudorandom stream of data, this data can be XORed with the plaintext to produce ciphertext. A hash function will always generate the same output for a given input. So to encrypt data using a hash function we could do the following: Generate a random IV.

Steve: That's the initialization vector.

Leo: Initialization vector, okay. Supply a passphrase to be used as the encryption/decryption key. Append the passphrase to the initialization vector, and generate a stream of data that is at least the length of the plaintext according to this function: f(IV+pass) = hash(IV+pass) + hash(hash(IV+pass)) + ... + hash(...hash (IV+pass)). In other words, keep hashing it till you have enough. Now, each time you produce that hash, you're going to hash the hash, so you'll get a different result each time.

Steve: And it'll be longer.

Leo: Yes.

Steve: Because when you add the hash to the end, and so that's longer by the hash's result length than what you had before.

Leo: So you hash the IV+pass, then you hash the hash of the IV+pass, plus then you hash the hash the hash, and on and on and on.

Steve: Right.

Leo: Till you get enough ciphertext. Then you XOR the plaintext with this data stream to get the ciphertext. So now he's going to say again what I just said, which is start by generating a hash of the passphrase and IV together. Then to make the produced hash long enough, keep hashing it again and again, each time appending the result to the stream.

Now, the thing is, this is reversible; right? Because to decrypt you just do the same thing, the same IV and passphrase used during the encryption. So the IV would have to be supplied together with the ciphertext. We'd have to say, here is the ciphertext; and you need this, the initialization vector. Have at it.

Steve: Right.

Leo: And of course the passphrase you have to provide. It seems to me that this should work pretty well. Have I missed something vital? I haven't thought of a use for this, since standard encryption is pretty readily available nowadays. I guess he's thinking this would avoid intellectual property, export restrictions, the kinds of things that current ciphers have.

Steve: Right. I thought this was clever. But I wanted to share this with our listeners because it demonstrates one of the perennial gotchas in cryptography, which is this stuff is really hard.

Leo: You can think you've got a good one.

Steve: Yes. My golden rule is, keep it as simple as possible, as simple as you can, but no simpler. And, I mean, this hashing of the hash of the hash of the hash of the hash will definitely generate something. But, he says, are there any security vulnerabilities with it? I have no idea. I mean, and a cryptographer might instantly go, oh, [sputtering], yes, that's not - it doesn't work because when you iteratively hash a hash, the following thing happens, you know, all the bits in the middle go to zero on odd Tuesdays or something.

I mean, it's truly the case that this stuff is complicated. And it's often discovered that when you repurpose something that was designed in one specific way for one specific purpose, when you repurpose it to something else, it just has horrible problems. And it looks great until the geniuses, the crypto geniuses sit down who really understand this stuff, and it turns out that, like, all the middle parts might cancel each other out. I mean, I'm just making this up. I have no idea. But I would never do it. Not without a crypto genius showing why this is a really good idea.

Leo: Right.

Steve: So you never want to just sort of come up with something because lord knows what you end up with in the crypto field.

Leo: The object lesson in this is WEP encryption, where perfectly intelligent engineers...

Steve: Yes.

Leo: ...created what they thought was a sound system. But any crypto expert could have looked at it and said, ah, here's the flaw.

Steve: Yeah. That's a great case in point.

Leo: Yeah. Those are smart people. I mean, it's not like these are - we were saying they were bad and dumb. But they just weren't crypto experts.

Steve: Right. And in fact Dain's observation was that, hey, since hashes have always been free of intellectual property claims and patent restrictions, why not do a hash-based crypto? I think there have been some hash-based cryptos. I'm sure I've run across them. But at the same time we now have AES. So that problem is solved. I mean, we have an absolutely bullet-proof, super-strong solution free of all that. Oh, and the other problem, hashing the hash iteratively would be very slow. It ends up being an expensive way to generate pseudorandom data given that it really was high-quality pseudorandom data. There are certainly easier ways to do that.

Leo: Right. Moving on to question 10. Richard Frisch in Weston, Connecticut wonders about password overload: Steve, you've often talked about passwords, but I have situation that is significantly worse than most, and I wonder if you know of a solution to password overload. I have a client who does accounting, bookkeeping, and other functions for over 60 different clients. Ironically, one of them is Rube Goldberg, LLC. She has more online accounts and passwords than I could shake a stick at. As, by the way, does our bookkeeper. These bookkeepers go from client to client, and each one has their own set of passwords. I believe she has more than 200 account names and passwords she needs to

know. Some of the passwords are static. Some must be changed periodically. Sometimes she works from her office computer; sometimes she is at the client's. Now I'm thinking I should probably ask her how she does this because I want to make sure she's doing it securely; right?

Steve: Yeah, for your sake.

Leo: For us, because she has all our account passwords. Almost all of the work is on PCs, but a few clients have Macs. Right now she records all this information in a handwritten journal she keeps unsecured at her desk. This is not good security, and it's a real pain for her, to boot. Do you have suggestions for a better, easy way to handle this? Love SpinRite and the Security Now! show. Hello to Leo. Boy, now I have to ask Lisa how she's doing it.

Steve: Yeah, and I don't have a specific solution. I mean, we know that the industry is full of all kinds of solutions. But I just sort of wanted to step back. I like this as a case study. And obviously you can relate to it because you've got an accountant in a similar situation. Look at all of this nightmare for the single need, the single purpose of authentication. That's what - this all is, am I me? And again, here we are in 2009, recording this on April Fools Day. I'll just bet you, I don't know when, but a decade from now, if someone's listening to this, they're going to be going, passwords? What? 200 passwords written down on paper? You've got to be kidding me. I mean, I don't know whether it'll be that you just have your wrist scanned or your retinas flashed or everybody's got...

Leo: Or a YubiKey.

Steve: ...YubiKeys built into them somehow. I mean, this problem, we are going to solve this problem. Today we haven't, and it's escalating. I mean, the need, the pressure to solve it, as demonstrated by this question, is becoming overwhelming.

Leo: I mean, just from a purely pragmatic point of view, I probably - I think a password manager would be the right way to go. There is - I've recommended RoboForm on Windows and 1Password on the Mac. But there is an open source one called Keepass.

Steve: Yes.

Leo: It's cross-platform. It's open source. It uses AES and Twofish, so it's very secure. So what you do is you basically, you've got a database of all these passwords that is secure in itself. And then you can put that on a USB key. You can have multiple copies. There's no reason not to because the database is itself secured. I should probably help our bookkeeper with something like that.

Steve: Yeah, one, I mean, one question or issue, I guess, is whether you carry the database with you, or you place it in the cloud. And of course that's a choice that's up to the user. It's certainly the case that, okay, the problem with the cloud - "cloud" is putting it out on the Internet somewhere - is that it's no longer necessary to trust that entity. That is, we can pre-encrypt the data before it goes there. I do that with Jungle Disk. I have an Amazon account, their S3 service, their storage service. And Jungle Disk uses a key that Amazon never sees. So everything there is pre-encrypted. So it's no longer to trust them except you need to trust their availability. You need to know that when you're going to log on, that they're going to be there to provide your vault.

And you also - you need to be in a situation where you have access to the cloud. If you're logging on, for example, to TrueCrypt in a full-drive encryption mode, then you're not on the 'Net at that point. So in that case you need a solution that can be sort of an offline authentication solution. So I think it's one of the problems is there isn't a single easy answer that solves every problem. But boy, do we need one. I mean, we just need to solve this authentication problem.

Leo: Yeah. I mean, I like how you took it to the higher level issue, which has to be resolved at some point. John Paquette in Framingham, Mass., wonders about Granola: Dear Steve, I often hear you using the word "granularity" as a synonym for "resolution," as in, "Oh, they've really bumped up the granularity of this so that users will be able to automatically disable AutoRun on drives of unknown types, on removable drives, on network drives, CD-ROMs" - that was from last episode. Doesn't "granular," like "grainy," mean the opposite of "resolute," "precise," or "articulate"? Oh. Love the show. Heard them all. Own SpinRite. Well, "granular" is a geek term, I think.

Steve: I guess. And I think he raises a good point because, okay. I'll say something is more granular. Now, what I mean is more finely grained.

Leo: Not more grainy.

Steve: Right. Well, no. More, well, okay. Do I mean...

Leo: I guess it is.

Steve: ...fewer big granules? Is that more granular? Or is it more granular if you have many smaller granules? And I decided it has no meaning. More granular? That's a non sequitur. I think...

Leo: It's only in conjunction with the idea of globular.

Steve: [Laughing] So granular and globular are antonyms?

Leo: Yeah, well, or - yeah.

Steve: No. I don't think. I don't think you could have more or less globularity.

Leo: You're right. Globular, granular, you're right, yeah.

Steve: You know, I think it's...

Leo: It's a colloquialism. It is not a...

Steve: It is not a precise term.

Leo: Precise term, you're right, yeah.

Steve: Yeah, granular just means able to be broken up. But I don't think being more granular, does that mean the grans are bigger, or are they smaller?

Leo: Well, for the first time, because I usually don't pan back so far, but I'm watching on the video, I see the Oxford English Dictionary behind you.

Steve: [Laughing]

Leo: So I think - and I have a copy here, too. I think we should look this up.

Steve: I actually did.

Leo: Oh, you rock. And?

Steve: And I'm using it wrong.

Leo: Well, but it's colloquially used. You're using it properly in the computer science context.

Steve: Ah, but what does "more granular"...

Leo: I agree. It's a meaningless term.

Steve: Yes, exactly.

Leo: But as with all language, it's understood what you mean.

Steve: Yeah, well, John didn't think so. And I think I'm...

Leo: Okay, we won't use it any more.

Steve: I'm agreeing with him.

Leo: That's good. No, you're absolutely right. What does it mean? It's meaningless.

Steve: Yeah.

Leo: Wow. John. You've made a significant change.

Steve: Good job.

Leo: I love the OED. You have the full set there, it looks like.

Steve: I do. It's just, oh, my god, they weight about 15 pounds each. And they've got little tiny thin paper. And oh, if you ever wanted, I mean, you can trace back the history of the words, into how the usage has evolved over time.

Leo: Well, that's what made me think of it is what we need is a geek OED that we could trace back to the first geek usage of granular, so we could understand what the context was and why. Yeah, the OED, the only time I really got mad at my wife - I bought the OED, and I love it. I thought the kids would use it. They never do. I kept going, oh, let's look it up in the OED. But I always wanted it for myself. Always wanted it. So I finally bought it. It's fairly expensive to buy all, whatever it is, 20 volumes or 30 volumes. I bought it after reading "The Professor and the Madman," Simon Winchester's book about the OED. That'd be a great Audible recommendation. It's really fascinating.

Steve: I've come so close to getting it on disk a couple times. If you look around, though, apparently the access software is really bad.

Leo: Yeah, you know it's going to be.

Steve: Yeah, exactly.

Leo: It's fine to have a book. It's a beautiful - it's a beautiful thing. Where was I going with this?

Steve: And you don't have to worry about booting them up, either.

Leo: Oh, the one time I really got mad at Jennifer, after buying the OED, I found it in the garage on the floor.

Steve: Ooh.

Leo: She said it was taking up a lot of space. I said [flabbergasted nonverbal reaction]. So now it's here in my office. With another set of volumes that was taking up space, equally unused, The World Book Encyclopedia.

Steve: Ah, yes.

Leo: We're old analog guys, despite our digital heritage. Are you ready, my friend, for the last question?

Steve: Let's wrap this sucker up.

Leo: Have you ever read that book, "The Meaning of Everything?"

Steve: No.

Leo: You have the OED. That was what inspired me to buy the OED. I said, I have to have this.

Steve: Well, I also love that whole concept of a bunch of guys sitting around deciding what words are new.

Leo: Well, the first thing they had to decide was where do we start? Where does English begin?

Steve: Yeah.

Leo: And I think that they didn't want to go past Old English. So I can't remember exactly where they started. But that's a puzzle. And then every word, they had to go through almost all the canon, every written word in English; and they made out little slips of paper with the word, the first use, the date of the first use, an example sentence.

Steve: Wow.

Leo: I mean, it's a mind-boggling thing. So what you have there on your shelf, and what I have on my shelf - it's over on his upper left shoulder, upper left corner for those of you who are watching at home, yeah, there it is - is in many ways one of the most significant books ever written in the English language. It's just amazing work. And that book will really turn you on to it. It's incredible. I can't figure out if "The Professor and the Madman" is the same book or not. I own them both. I should really listen to them and figure it out. I think they might have retitled it.

Anyway. Last question. Rick Hughes in Sykesville, Maryland brings in his Q&A Tip of the Day, "An Easier Way to Use Drive Snapshot." Now, I got a very nice note, by the way. Did you get that note from the author of the program that you recommended?

Steve: I did, and I'm puzzled about it because I was so sure that I once spoke to someone named Paul Terrell. And I thought that that's what TeraByte, the T-e-r-a, was from Terrell. I guess I'm wrong because the guy at TeraByte said...

Leo: I wrote that.

Steve: ...who's that, I never heard of him.

Leo: David F. of TeraByte Unlimited said: Steve mentioned on Listener Feedback 62 Paul Terrell [indiscernible] TeraByte. I have no clue who this is. The president's name and primary developer is me, David Flicek.

Steve: Yeah, anyway, it's funny, too, because I meant to go do an email search because I have all my email from, like, forever, and find out. Because it would have been an email dialogue about - I thought about...

Leo: Well, you know what I suspect happened is that he, well, of course he would know the guy's name.

Steve: Yeah.

Leo: Maybe he acquired it?

Steve: That's what I thought. I thought the same thing. I just, you know, anyway, I'll solve the mystery.

Leo: I bought it immediately after you talked about it, and I've used it now to image a few things. And the main reason I bought it is because Drive Snapshot makes it hard to make a bootable disk that you can then reinstall. And that's kind of the key with an image. What we're talking about is image software that makes a ghost image of your hard drive.

Steve: Right, and Drive Snapshot also runs on DOS, so it's only able to see what DOS is able to see. Whereas Image for Windows, TeraByte Unlimited's Image for Windows, it brings along its own 1394 Firewire and USB drivers. It's able to enumerate the bus, find the drives. And it's also both - it's OS agnostic. It runs on FAT file systems and NTFS file systems.

Leo: Right. Yeah, no, I've been very happy with it. But I also own Drive Snapshot. On 188, you and Leo were discussing the difficulty of restoring from an image backup using Drive Snapshot's DOS-based restore disk. As we just said. An easier way to use Drive Snapshot is to make a BartPE disk, which is at nu2.nu/pebuilder. This is a free, bootable Windows XP CD that natively understands NTFS, USB drives, networks. If you copy your licensed Drive Snapshot .exe file to the BartPE disk, you can just boot up from the CD and use the normal Drive Snapshot GUI for backing up and restoring to any drive - this is really what I should have done - without having to fool with the DOS drivers. I've been using Drive Snapshot this way for a couple of years with no problems. I know you've switched to TeraByte's Image for Windows, but Leo and others who still use Drive Snapshot might be interested in this approach. Thanks for all the great Security Now! netcasts and GRC software. I've been using SpinRite since version 2. What an excellent product.

Steve: You know, we've never talked about BartPE.

Leo: We should.

Steve: Yes, we should. I had, believe it or not, my first occasion ever to use it within the last week. I can't even remember now why. But I got myself painted into a corner somehow, I mean, I'm old school, so I've got old DOS tools that normally still - I've got Partition Magic and Drive Image, and I'm very comfortable with all of that. So I think that's kept me from having to mess with BartPE. But somewhere in the last week I needed it for something I was doing. And I thought, well, okay, I've heard about it. I mean, I knew all about what it was. I'd never done it myself.

So I just wanted to bring it up and mention it to our listeners. I'm very impressed with what it does. The idea is that it needs access to your original Windows installation CD. It's from there that it gets the original Windows installation files. It's not able to get them from your computer because it doesn't know whether they're intact, what shape they're in, whether you can trust them. I mean, BartPE is often used, for example, as an AV platform for running Windows-based antivirus. So you wouldn't want to start with an infected Windows boot.

But what this is able to do and what it does is it builds an ISO image of a bootable CD with files that it gets from a real Windows setup/install CD. And it just works. It's a script that pulls the files together. You're able to add your own, put your own stuff in a directory so that you can customize it that way. And you end up with this CD that boots and runs Windows from the CD, and then you can do other things that you want to with it, as if you had a running Windows. And if Windows won't boot, it's often a way to have some tools that allow you to figure out what's going on.

Leo: I tell everybody to make a boot CD for sure. This is a good one. Emergency Boot CD is another one. And just to be fair, TeraByte also offers a plug-in for BartPE. And when you download TeraByte, you can download the plug-in. Makes it very easy to make it part of your BartPE build, so that in future, when you build future BartPE disks, it'll just automatically be included.

Steve: Cool. I think I'll probably put it and Drive Snapshot...

Leo: Do both.

Steve: Yeah.

Leo: You probably have images lying around from both programs.

Steve: Oh, yeah.

Leo: Well, Steve, we've reached the end of 12 fine questions and true from our great listeners.

Steve: And so nice to have flawless telephone quality audio, Leo.

Leo: People are going to the chatroom, going what's wrong with Steve?

Steve: Argh.

Leo: We will get back to Skype next week. Just a little...

Steve: We absolutely will.

Leo: ...Skype hap, as happens frequently. But, hey, I understood every word you said. So that's the most important thing. If you want to get 16KB versions, in which Steve will sound even worse, and transcripts in which Steve will sound exactly the same, you'll get all of those at GRC.com. That's of course the home of Steve Gibson and the Gibson Research Corporation. There you'll find all his great free security software, like ShieldsUP!, Shoot The Messenger, DCOMbobulator, Wizmo, and the world's finest hard drive maintenance and recovery program, a must-have for everybody - can you put that on a BartPE disk? I bet you could.

Steve: Well, you wouldn't need to because it doesn't need to run under Windows, so.

Leo: Yeah, I wish you - yeah, SpinRite [indiscernible].

Steve: It makes its own bootable.

Leo: I wish Snapshot would do what you do with SpinRite, which is it makes its own bootable.

Steve: Right.

Leo: You can make a boot disk, but you have to - it's complicated.

Steve: It is.

Leo: SpinRite is available at GRC.com. Get yourself a copy. It's well worth it. And we will convene again next week for more security information. In fact, next week we're going to talk a little bit more about...

Steve: The GhostNet.

Leo: The spy network.

Steve: The uncovering, the discovery of this really multinational, 103 different countries, spy network. Which, incidentally, is able to turn on the camera and the microphone that are built into various machines in order to see what's going on.

Leo: Oh, boy.

Steve: I mean, we're really talking spy technology.

Leo: Yeah, yeah. All right, Steve. We'll see you then.

Steve: Talk to you then, Leo.