
# Listener Feedback #158

**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-385.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-385-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Our first episode of the new year, and lo and behold, has TrueCrypt been cracked? Steve will answer your questions, including that one, next on Security Now!.

**Leo Laporte:** It's time for Security Now! with Steve Gibson, Episode 385, recorded January 2nd, 2013: Your questions, Steve's answers, #158.

It's time for Security Now!, the show that protects you and your loved ones online. And speaking of loved ones, Happy New Year, Steve Gibson.

**Steve Gibson:** One two one three. Great date. I like that. Although 12-12-12 was good, too.

**Leo:** And then somebody pointed out that March 14th, 3.14.15, 2015, will be a very auspicious date,

**Steve:** Beginning of pi. Have some pi on that date.

**Leo:** Some pi. That'll be - really that'll be pi day.

**Steve:** Yeah.

**Leo:** Hi, Steve. How was your holiday?

**Steve:** Well, I have to say, Leo, I'm so pleased by the response to what we chose to do for the podcast over the holidays.

**Leo:** The young Steve Gibson.

**Steve:** It was a real hit.

**Leo:** Oh, good.

**Steve:** Yeah, it really was. I got a bunch of feedback in the regular Security Now! feedback mailbag that I saw last night as I was pulling questions together for today. And also Twitter was crazy with people saying they haven't laughed that hard at a moron for quite a while.

**Leo:** A moron?

**Steve:** Oh, no, no.

**Leo:** Oh, my. I hope they didn't say that.

**Steve:** And many people were, as I predicted, very surprised that so little had changed in 22 years. Indeed, a lot has not. In fact, it did spur some questions. I just answered a couple of them in today's Q&A, that many people had, saying, well, wait a minute. If that's the way things were, is that the way they still are? What's not that way?

**Leo:** Isn't that interesting, yeah.

**Steve:** I've got a couple, I'll answer a couple of those things today. But, yes, I think I have an idea for next year. People want more of that. It's like, well...

**Leo:** You got any more VHS tapes in your closet there?

**Steve:** I do. I do have old TechTV shows, Screen Savers.

**Leo:** But those are short. But we could bring those back.

**Steve:** Yes, that's a good point. It was only a short - although I have them all.

**Leo:** Well, and I should point out that, despite the fact that G4 is now history, has become the Esquire Channel, or about to, that they still reserve the rights to everything. So occasionally they'll do takedowns on YouTube when people put stuff out there.

**Steve:** Ooh, you're kidding, even now.

**Leo:** Yeah, yeah.

**Steve:** But anyway, I've got some ideas for next year.

**Leo:** Okay. We've got time.

**Steve:** It was fun for everyone. And the fact that it was a video didn't seem to slow many people down. So what I may do is keep with that concept and do something visual for the holiday break on my own and then provide it to you guys so that you can publish it.

**Leo:** Cool. You could do, like, dance lessons or something.

**Steve:** I don't think that's what I'll do.

**Leo:** Okay.

**Steve:** Although apparently I was spinning around pretending to be a screw.

**Leo:** Yeah. There was a little dance move in there.

**Steve:** A little dancing there.

**Leo:** So you said it. We're going to do a Q&A episode this week, our 158th Q&A episode.

**Steve:** Yes. And the hackers took the week, took the holidays off also, apparently.

**Leo:** Oh, that's good. That's nice.

**Steve:** There was a good conference, the CHAOS Conference, over the holidays, and it produced some interesting things. There was one report that I haven't yet followed up on

about the cost of pulling files back out of Amazon's Glacier storage. I think it's an alarmist posting, but I need to do some math because it was some - one particular reading of their FAQ led someone to believe that it could be prohibitively expensive to pull things out of cold storage. But that doesn't make sense to me. So I'll follow up on that.

But the big news is we've got a new zero-day vulnerability that was discovered, not surprisingly, in Internet Explorer, leveraging an Adobe flaw. It only affects versions 6, 7, and 8. So the very latest people, the users who are using 9 and 10, IE versions 9 and 10 are not in danger from this. But what's happening is it's being used in so-called "watering hole" attacks. That's the jargon that the industry has developed. You know we have phishing, right, so this variation, this is being on the other end of the phishing pole. The watering hole attack is, if you want to attack, like, say, the Department of Defense, you go find a vendor or a site that your typical DoD employee would visit, and you compromise that.

**Leo:** Ah, perfect, yeah.

**Steve:** So you poison the watering hole where they innocently go. So that has actually been happening. So this was discovered for the first time in the wild. Shoot, I can't remember, who was it? I did quote Brian, and I don't think - he said, "Attackers are breaking into Microsoft Windows computers using a newly discovered vulnerability in Internet Explorer.... While the flaw appears to have been used mainly in targeted attacks so far" - and this is really Brian's forte. He says, "This vulnerability could become more widely exploited if incorporated into commercial crimeware kits sold in the underground." And Brian's done a really good job of developing a persona or set of personas, I presume, that allows him to monitor the underground. He gets a lot of his information that way. And so he's able to see when these new exploits get moved into the single pushbutton crimeware kits.

And so what's going on is Microsoft has acknowledged it in their typical Micro-Speak format. They said, "Microsoft is investigating public reports of a vulnerability in Internet Explorer 6, Internet Explorer 7, and Internet Explorer 8. Internet Explorer 9 and Internet Explorer 10" - apparently they will not say "IE," no matter how much they have to type "Internet Explorer" - "are not affected by this vulnerability. Microsoft is aware of targeted attacks that attempt to exploit" - sorry, Microsoft, do succeed in exploiting - "this vulnerability through Internet Explorer 8." So 8 has been exploited; 6 and 7, however, are known to be also vulnerable. Says Microsoft, "Applying the Microsoft Fixit solution," which is called "'MSHTML Shim Workaround,' prevents the exploitation of this issue. See the Suggested Actions section of this advisory for more information."

And then Microsoft goes on to say, "The vulnerability is a remote code execution vulnerability that exists in the way that Internet Explorer accesses an object in memory that has been deleted or has not been properly allocated. The vulnerability may corrupt memory in a way" - now, we understand the vulnerability DOES corrupt memory in a way. Microsoft says "that could allow," we change that to "does allow an attacker to execute arbitrary code in the context of the current user within Internet Explorer. An attacker could host a specially crafted website that is designed to exploit this vulnerability through Internet Explorer and then convince a user to view the website." And as we know, the way it's actually being used in the wild is websites with good reputation, not created by hackers, are being modified to insert this attack.

So there is a Fixit. I just tweeted it on the @SGgrc Twitter stream [New 0-Day IE flaw

being actively exploited. Users of IE: Go here to apply Microsoft's quick interim FIXIT: bit.ly/WeEJx5].

Or you can just go support.microsoft.com/kb/ and then the magic number here is 2794220. So that's Microsoft.com, support.microsoft.com/kb/2794220. I would recommend that people do this if you are a heavy IE user. I would imagine most of our listeners...

**Leo:** The workaround is not a fix, it's a workaround.

**Steve:** Correct.

**Leo:** So what's the difference?

**Steve:** Well, okay. This just shuts down the avenue of exploitation, but doesn't actually fix the problem. So here we are with the first of the year being Tuesday, which means the second Tuesday of the year, and of the month, is the soonest it could be. So that will be next Tuesday. And Microsoft just got onto this this weekend, this past weekend. So it's not clear whether they're going to have time to put this through everything that they need to, all their regression testing and everything, to get this thing into the second Tuesday of this month. If it escalates and becomes bad enough, they could of course do an out-of-cycle patch. Maybe they'll wait till February, next month's second Tuesday update, to do this. We don't know.

But if people, for whatever reason, if you're in a corporation that's standardized on IE, and you have to be using IE, you might want to make sure that your IT folks know about this and maybe - because you can have this Fixit easily propagated through an entire corporation. The Fixit is able to do sort of a cross-network mass installation. So this may be important. It's the kind of thing - these are what the attackers jump on because they recognize there's a window of opportunity which will be closing pretty quickly, and they'd like to jump on it as quickly as they can.

And the only other thing I had was just a note at the end of the year, an interesting article quoting a report, and the article was in The New York Times on 12/13, so New Year's Eve. The headline was "Outmaneuvered at Their Own Game, Antivirus Makers Struggle to Adapt. Antivirus Industry Is Always Reactive, Not Proactive." And that's something that we've been feeling more and more. I mean, even if you have state-of-the-art, updated, latest AV, all of these things are getting past. I mean, think about it. We never hear that, oh, unless you've got this AV or that AV, you'll be okay.

**Leo:** That's a good point.

**Steve:** That's not the case anymore. So Ted Schlein, who's a security-focused investment partner at Kleiner Perkins Caulfield & Byers, of course a major tech venture capitalist up in Silicon Valley, he was quoted in this article saying that "Existing methodologies we've been protecting ourselves with have lost their efficacy. This study is just another indicator of that. But the whole concept of detecting what is bad is broken." And we've been sort of talking about these ideas. We've talked about how the first firewalls were all open, and then you closed them to protect things. And then people

realized, okay, that model is backwards. We need to close them by default and then open them to allow things that we want.

And, sadly, that model is very awkward to apply to software because we're all using software coming from everywhere. I mean, even a very carefully managed store like Apple's iOS store, I mean, there are things getting past them all the time. So even if you closely curate what you allow, you're still going to have a problem. So my feeling is we need to fundamentally change the way our systems work so that, at the system level, they are invulnerable to this kind of malarkey. And you can sort of imagine a form of a virtual machine system where everything runs in its own space, more like an iOS model, which even that's not perfect, and it does cause people problems because you're going to trade off features for safety.

**Leo:** I know antivirus companies talk about ESET, which is a sponsor on the radio show.

**Steve:** Behavioral based?

**Leo:** Yeah. They call it "heuristics," where they're looking - they understand, everybody understands signatures are not updated fast enough, obviously, for zero-day exploits. But what if you look for virus-like behavior? Is that effective at all?

**Steve:** Well, it gives you, I mean, there's nothing that's a perfect solution. I did see somebody, I think it was either a tweet or in the mailbag, just telling me that my LeakTest program was being considered a malware by some latest version of some antivirus system. And that's exactly this, that this is the problem with heuristics is that they're inherent - because they have to recognize something they've never seen before, they might misrecognize something they have not seen before as bad when it's not. And then the tighter you make that, that is, the more protective you make the soft determination, the more false positives are going to get caught in that net. So it's just a mess.

**Leo:** I'd rather have a false positive, I guess.

**Steve:** It's a classic - well, except then you get harassed. I mean, it's like you and NoScript, how you love NoScript popping up notices constantly.

**Leo:** Right, so I turn it off.

**Steve:** And you're having, exactly, saying, okay, this is just not worth the hassle. So, I mean, what this tells us is that all of these solutions are wrong, that is, they're the best we can do given the foundation that we've created. But so it's the foundation that's the problem. It's fundamentally broken. The fact that this can happen to people is ridiculous. So I look forward to the day when you and I will connect up, and you'll say, "Well, Steve, do we have a podcast?"

**Leo:** Never gonna happen.

**Steve:** "Nothing happened, Leo. We went a whole week, and nothing. No."

**Leo:** Now, QE says, and I don't know if this is a fair thing to say, I'd rather use a more secure OS like Ubuntu. Is, by its nature, is Linux more secure? Or is it just less targeted? I mean, we thought...

**Steve:** It's less targeted.

**Leo:** Yeah, we thought Apple was more secure until it became more targeted. Then we realized how it was the same as anyone else.

**Steve:** Yeah. What we see over and over is that all of these systems have porous surfaces. They're fundamentally porous. And so it's just a matter of how badly you are willing to work, how hard you're willing to work, how badly you want to get in, that determines whether or not you can because their systems are just too complex. They have really outgrown our ability to understand every possible interaction. Look, for example, at a protocol like TCP, which is so solid and has been around for a long time. Everybody's pounded on it. Yet a couple clever guys realized that, oh, look, there's an inter-packet dependence in the way the cipher block chaining works, and we can leverage that. I mean, just little tiny mistakes that are there, not like a buffer overrun, but actually a flaw in the protocol. These exist, unfortunately, throughout our systems because they are so complex. We're asking them to do so much.

**Leo:** It's easy to say something's more secure. And I just - and it may be, until it's hammered on in the way that Windows is hammered on, I don't know if you can - we can be sure.

**Steve:** Well, and you might argue that, if it is not being attacked, it's more secure.

**Leo:** Oh, yeah. That's for sure, yeah.

**Steve:** So I would say that Linux is probably safer to use today, not intrinsically safer, but effectively safer. That is, it doesn't have to actually be more secure. It just has to be something no one's attacking, in which case you get the effect of it being more secure because it's not being attacked.

Anyway, I did get a kick out of - this is not security related. This is just we're now into miscellany. I tweeted an interesting little observation that people began observing on New Year's Day, and that is that their Do Not Disturb option, which was added to iOS 6 so that the phone would not ring and messages would not alert you and wake you up, you're able to turn on Do Not Disturb at 11:00 p.m. and turn it off at 6:00 a.m., for example. And it turns out it no longer turns off. And we were first thinking, well, maybe this is just New Year's Day. And a lot of people had fun with it, saying, oh, it knew that I

needed to sleep in to get over the hangover of New Year's Eve. But it turns out that it's still not turning off for everybody, even today.

**Leo:** And of course they ran, of course Apple ran an ad touting - with, what, Serena and Venus Williams - touting their Do Not Disturb feature that didn't work.

**Steve:** Yep. Ill-timed ad.

**Leo:** It seems dates are a problem.

**Steve:** I imagine that they will be releasing a "disturbing" update to iOS 6 at some point.

**Leo:** But what that underscores is what we were talking about earlier, which is software is difficult to make perfectly.

**Steve:** Yeah. And you know also, Leo, you probably remember this, that for the last several years they have had New Year's problems. There was one in 2012, and there was one in 2011. That just really seems to trip them up every time.

**Leo:** Apple apparently says it will fix itself on the 7th.

**Steve:** Oh, nice. They've got really rounded corners on their stuff, but that first week of the new year, that's just a little problem.

**Leo:** Now, here's a challenge for adept programmers.

**Steve:** Explain that, yes.

**Leo:** What would be the bug that would have the Do Not Disturb fail January 17, 2013. What is it - and then it starts working January 8. What is it that's - I can't think of it. But that's a good puzzler. If you're a good programmer, you should be able to figure that out.

**Steve:** Yeah, they just forgot to carry the one. Okay, so...

**Leo:** That's always a problem.

**Steve:** So a little update on my Quiet Canine project. The device has a name. And I got trademark protection for it just because. I don't know that anything commercial will ever come from this.

**Leo:** Oh, aren't we fancy.

**Steve:** But it's just the name is so good. We were playing with Hush Puppy for a while, but that's just overused and so forth. Anyway, it's called the Treble Shooter.

**Leo:** The Treble Shooter. I get it. Because it shoots a high-pitched treble sound.

**Steve:** Yeah, Treble Shooter. And we now have two designs. Well, one design is finished, which is the one that I will make a bunch of in order to satisfy my interest in determining whether or not this does anything more than this full industry of attempts. Because, I mean, the barking problem is - I've been flabbergasted by how many of our listeners are like, oh, my god, do I have a problem. And then they explain it to me. And it's like, I mean, there's divorces, and there's crying babies, and we have to move out of the house, and the helicopters are circling overhead. And it's just unbelievable what is going on just from dogs barking. And of course we know famously that arguably McAfee is on the run somewhere because he had a backyard full of barking dogs, and things escalated to homicide, presumably, from that.

**Leo:** See? You could have saved a life here.

**Steve:** So this is a problem for people. The first design came out better than I had ever imagined it could. It's very, very few components, even less than I showed you, or even fewer, sorry, components than I showed you when I was holding up that little first iteration version. Really, really cool. It takes a 9-volt transistor radio battery, from which it only draws 100 milliamps, and it produces an 80-volt beautiful sine wave across this capacitive piezo tweeter, so using the tweeter as part of the circuit. So I'm really happy with that.

And it runs at a fixed frequency with a button that you can press to drop the frequency down, just so you can appreciate how loud this is for the dog. The problem is, if it didn't make any sound at all that we could hear, I mean, it's still powerful, and it's still potentially harmful. So I want people to appreciate the amount of sound this is making, even though they can't hear it. So the second button just drops it down from 16KHz that we can't hear, or barely hear - kids can hear it - down to 8. And, boy, it is so...

**Leo:** That's a very good idea that you did that. I think that's great.

**Steve:** Yeah. It's like the physicists who were working on the development of the H bomb for the country, the Manhattan Project. They all died of cancer because something that they sort of only dimly understood was really affecting them. And I don't want anyone's hearing, anyone, human or animal, their hearing to be harmed by this. So it takes respecting how loud this is to treat it properly. We know about guns and knives. They're just obvious. We have real-world understanding of them. But this is something different. And this is too powerful to be used casually.

**Leo:** Good for you. Good for you.

**Steve:** And I've got about 150 of our listeners so far have said, oh, my god, I have a problem that I need help with. I can't make that many. But what I'm really wanting to do is to find the people who have already used other, like there's a market of sonic deterrents, none of which work. And so my question that I really want answered is does this one work? Is this any different? Or do the dogs get used to it? Who knows what. I do know that my buddy Mark, who's got that overkill one that I made, he used it yesterday because these two dogs were in the back barking. His brother is visiting over the holidays with his dad. And they just tapped the button. And his brother actually thought the dogs had been vaporized because they were instantly gone. It's like, where did they go?

**Leo:** Oh, my god, it's a ray gun.

**Steve:** They were just gone. And he says it's like the Keystone Cops, the dogs trying to both squeeze at the same time through the doggie door that can only accommodate one of them. So anyway, that project is moving on. The second device I'm working on now, since the first one is perfected, is I still want to do, before I drop this - because once I get it documented and put it up on GRC, I'm going to move on to other things as I always do and probably never come back. So while I'm in this mode, I want to nail it. So what I want to nail is a microcontroller-based version that could do all kinds of other cool things, like only send out a short burst; maybe at some future time respond, hear barking and respond to the bark in order to do an automatic training mode, those sorts of things. So I'm going to - I want to nail down an efficient microcontroller-based power amplifier before I move on to other stuff.

**Leo:** Unless people - I was watching football over the holiday break. And lest people think this is cruel, I saw an XFINITY ad. You know, they have the ad for - I think he's, like, stopping the DirecTV guy from coming. He's, like, powerful, some football player, And there's a dog comes out and barks at him, and he throws a stick at it, hits the dog in the head. The dog runs off. I thought, if they can do that in a national TV commercial, then there's a little yelling, a little sound, it's like a yell, it's like "Get out of here," is fine. It's not going to hurt the dog.

**Steve:** Well, and I have to say, Leo, after reading the torture that dogs are really putting people through, I mean, I don't want dogs to be hurt.

**Leo:** I blame the owner. It's the owner's fault, not the dog's fault, obviously.

**Steve:** Well, many owners are at their wits' end, too. They have apparently - the owners are not trained, and presumably the owners could be trained, but that doesn't seem to be happening. So, I mean, it really - I am a dog lover. I grew up with dogs. But, I mean, many people are - and the problem may be the owner, but it's the neighbor owner. And many people have written to explain that they have begged and pleaded and offered to do anything if the owner would not put the dogs in the backyard and drive off for the day, but the owner just doesn't care. It's like, eh, tough, too bad.

**Leo:** Yeah, the owner's not there. They don't hear the barking.

**Steve:** Yeah. So in a little bit of health-related news, I mentioned this before we began recording, a number of our listeners who followed the whole low-carb deal sent me links. And of course I follow a lot of these things anyway, so I had seen them, too, but thank you for sending the links to make sure that I had seen it. This ketone that I have mentioned, which I have been measuring in my blood now for nine months because I have been continually in ketosis since that happened to me on April 17th, for the last nine months, the one that you measure with the blood test is called beta-hydroxybutyrate. The other one was acetoacetate. Beta-hydroxybutyrate is the product of fat metabolism, which appears in fasting or starvation or, in moderate amounts, in ketosis.

Well, a report came out, produced by a huge number of PhDs, I mean, it's like, my goodness, you've got to be kidding me, about their discovery of the mechanism by which beta-hydroxybutyrate extends your life, that is, extends lifespan. And we know that caloric restriction, CR, is known to extend lifespan. But the problem is you have to almost starve yourself. Well, it turns out that one of the mechanisms may well be that, in caloric restriction, you are also in ketosis, and you're producing beta-hydroxybutyrate.

What they found was that the presence of beta-hydroxybutyrate blocks a class of enzymes known as histone deacetylases that would otherwise promote oxidative stress. And we also know that people in ketosis have much higher levels of antioxidants in their bloodstream. We haven't understood exactly why until now. But the idea is that burning fat produces lower levels of free radicals than burning carbohydrate. And now we understand also that one of the side effects of fat burning is the production of this beta-hydroxybutyrate, and it directly blocks the promotion of oxidative stress. So another good little piece of happiness for all of our listeners who have been experimenting with reducing the level of carbohydrate in their diet.

I do have the worst New Year's pun, courtesy of a frequent tweeter of ours, Simon Zerafa. He tweeted, "I'm going to make my New Year's Resolution - 1388x768."

**Leo:** I've heard that one.

**Steve:** Oh, god.

**Leo:** It's good. I like it. It's a little low. I'd actually like a higher resolution, if you don't mind.

**Steve:** I did notice, when I turned on my machine that monitors and reports on the number of remaining days of Windows XP support, that we still have, as of today, 460 days remaining of Windows XP SP3 updates.

**Leo:** That's the enhanced support, though. I mean, that's like if you're a - I think that it's expired if you're just a normal person.

**Steve:** No.

**Leo:** No?

**Steve:** They changed it about a month ago. You're right, but they just changed it, and they extended this for everyone.

**Leo:** Everybody. Good. Because a lot of people still use it. My question is how long is Windows 7 going to last? At least 20 years, I hope.

**Steve:** I hope so, too. Because my second part of this is I am liking Windows 7 more and more.

**Leo:** Oh, it's great.

**Steve:** I've been using it a lot. I've set up a couple of Media Center systems in the last couple months. I finally retired my triple TiVo configuration because it was an old standard definition. I went to HD, and I ended up building a Windows Media Center box, and a couple more of them for some friends. And I'll be quite happy to jump over Vista, from XP to 7 here, maybe about a year from now.

And, lastly, somebody sent me a tweet that I could not find again. I think it was a tweet. Turns out they have been cataloguing all of my SGgrc tweets for the Security Now! episodes where I've been putting links in my Twitter feed.

**Leo:** Very nice.

**Steve:** And I went to the page. It's beautiful. It's by episode, here's all the links and all the tweets for that episode. But I can't find the page again, and I can't find the tweet. So could I ask whomever that is to tweet me again? And I thank you for it, and I will share it with our listeners, hopefully next week when I see it.

**Leo:** And if you're in the chatroom, let us know, and we'll let you know. But it doesn't look like it's somebody in the chatroom.

**Steve:** And I did find a nice story from a listener of ours, sent on December 30th, Sunday, named Jenn. He's in San Diego. He said, "Fixing emergencies before they become emergencies, and lots of thanks." He says, "My SpinRite story: My company was trying out whole-disk encryption software, and my team served as guinea pigs. The first step was to run chkdsk, which found tons of errors on my two-plus-year-old laptop. In spite of that, it said it was clear to install the software and encrypt the drive. I did, and then it refused to reboot. Completely bricked. Help desk wasn't any."

He said, "No big deal. I of course backed up my data before step one, so I got a loaner laptop and was up and running again from my backup in no time. But I would have to

either get my computer wiped and rebuilt or would need to order and build a new one. When I reported the problem to my boss, also a fan of this podcast, he said, 'You've been meaning to buy a copy of SpinRite, haven't you?' 'Yes, sir.' I got myself a copy and started running it.

"After SpinRite completed its task, I was able to boot the computer, and it ran perfectly. I was impressed and have continued to use it to fix my family's computers before they break. When I've had questions, I've been impressed with how fast and thorough the response was. (Thanks again, Greg)," he says. "I've even recently taught my 10-year-old how to run SpinRite. More thank yous."

He said, "I have several professional certifications that require me to do continuing education. Since I've been listening to this podcast longer than I've had the certifications, I've often used Security Now! to meet this requirement. Recently one of these continuing education submissions came up for a random audit. I needed to provide a summary of the event and how it related to my work. Easy. I just looked up the transcript, skimmed over it to refresh my memory, and quickly wrote up the response. So thank you, Steve and Leo, for the great podcast, and thank you, Steve and Elaine, for the transcripts."

**Leo:** Well, thank you.

**Steve:** And thank Jenn for the nice note.

**Leo:** Yes, for the nice note. We're going to take a break. When we come back, we have questions from you the listeners. Some good ones, too.

**Steve:** Well, actually I want to say that we need to finish this one because #10 answers a question many people have been worried about. There was an article, it got Slashdotted and a few other places, about a company that had broken the encryption on whole-drive encryption.

**Leo:** Ooh.

**Steve:** So all kinds of people were saying, oh, Steve, my goodness, is this true? And so we will answer that in Question #10.

**Leo:** Don't fast-forward. Stay right here because we still have other questions. All right, Steve. I'm going to do some reading, if you will do some listening. Question #1 of our listener-driven potpourri comes to us from Tim D. in Detroit, Michigan. He had some holiday episode questions. As you mentioned, we have a few. So last week, Steve's holiday episode was a video of him from 1995?

**Steve:** '90.

**Leo:** '90, explaining how hard drives worked. And it prompted Tim to write: Do

modern spinning disks still allow you to access the low-level formatting? If not, how do they account for long-term drift? Do the seek motors move in increments much smaller than the width of a track? Does the servo information you mentioned allow the drive to correct for drift? What is it, Steve? Tell me, tell me.

**Steve:** So, many people asked this question in various forms in the mailbag. They got a much better sense from last week's holiday episode for the inner head and platter level interaction going on in a disk, which really is what I spent 45 minutes explaining in detail. And that left them wondering, okay, well, that was 22 years ago. We've talked about how little has changed. But the question of course was, well, okay, what has changed?

Probably the major change was when drives became intelligent. And that was with this conversion to the so-called IDE drive. We went from MFM and RLL and ERLL technologies. There we had a controller that had to be mated with a specific technology of drive. And the controller and drive together knew what kind of encoding would go on the drive and so forth.

Well, what happened was, essentially, that controller function moved into the drive. So it was no longer something that you plugged into the motherboard, and then it worked in tandem with the drive where actual, like, flux data level signals were moving across the cable between the controller in the computer and the drive. Instead, the controller itself moved into the drive with the move to so-called - and in fact IDE stood for Integrated Drive Electronics. We're integrating the controller with the drive. Now we simply ask the drive, which contains its own controller, for certain sector numbers. Give us this block, and we'll give you this block, and the drive takes care of it. So that was a huge change.

Well, what that allowed the industry to do, then, was essentially anything they wanted to, downstream of that so-called IDE interface. They were free to change the sector count in the drive because one of the things you could ask the drive was how many sectors you got, and the drive would just give you back a number. And then the computer will go, okay, thanks. We'll use zero to one minus that number as our sector numbers, and ask you for them that way. So that was a huge move because the drives stopped being just dumb electromechanical units, and they became smart themselves.

Then of course the world went crazy with density because this moving the controller into the drive and creating an IDE interface allowed the manufacturers to start cranking up the density, with it having no effect on its interaction with the computer, which was key. One of the things they did to solve the problem that I describe in the holiday video of long-term drive alignment drift, in that case, in the case 22 years ago where you had a stepper motor and essentially a rack-and-pinion gearing system or sometimes a metallic band wrapped around a capstan that was controlled by the stepper motor, there, over time, the mechanics could change, which would cause the tracks to drift. So the technology was completely changed to a so-called "servo-based" technology.

The first drives to do that dedicated a single surface out of however many surfaces the drive had to a so-called "servo platter." The servo platter, which was really just a servo surface, was only one half, one side of the platter. It contained sort of the master positioning reference for all the other heads. So what that meant was that, as alignment drift occurred, the same drift would occur in the access to the servo tracks that would occur to the data tracks. And so that helped to compensate for alignment.

But the problem was initially the servo was over on one end of a pack or a stack of

surfaces. Then they realized, oops, if we put it in the middle, that would be better because then it would be sort of - it would be physically, in terms of a sort of a rotational drift, it would be closer to the heads. And that was called - there was something called "tower misalignment," which I did mention in the holiday video, the idea being that you call the stack of heads a "head tower" because they're all moving together in and out as one. But the problem is even though you have fixed the alignment problem for the one head which is your positioning reference, if the tower skews a little bit, the head furthest away, at the other end of the tower from the servo head, that could still have an alignment problem.

So as they kept inexorably increasing the track density, chasing, cramming more and more data on drives, they ended up doing what's called "embedded servo," and that was another huge sea change in drive evolution was we embedded the servoing data into, essentially intermixing it with the data that is being stored. There was a technology called a "wedge servo" where two different phases of sine wave were recorded in a chunk of the disk. And if the head was exactly in the middle, the idea was that the amplitude of one phase of sine wave went down from left to right, and the amplitude of the opposite phase went down from right to left. So if you were exactly in the center, those two 180-degrees out-of-phase sine waves would cancel each other, and you'd get a null signal. But if you moved off center, you would start seeing a signal of one phase or the reverse phase, and which phase it was told the head which direction it was misaligned and by how much.

So what happened was, when we went to embedded servos, the ability to actually lay down new sector headers, as I describe in last week's video, we lost that. And we really arguably didn't need it anymore because the drives got factory formatted with a very sophisticated format. There were machines in the factories called "servo writers" which were the only things able to lay down this very delicate, super-precise servoing information on the tracks. And once that was done, they shipped the drives out, and they just ran by themselves from there on. So that began.

So essentially, as soon as we went to this embedded servo technology, the command to low-level format the drive, which had been in the interface, it was changed to a "zero the sectors." So there was a command called "low level format drive," or "low level format track." And all it did, though, was clear the data in them. It never actually rewrote the sector headers.

And I should mention that, by the time this happened, we had solved the interleave problem. The sectors had to be interleaved, as I described last week, because the systems were not fast enough to accept the data at the speed the drive was able to provide it. So sectors had to be spaced out from each other. Our machines got faster. The buses got faster. And moving the controller into the drive solved some of that communications bottleneck. So then drives were all interleaved at 1:1, meaning just a linear stream of sectors. The successive tracks were skewed rotationally so that, when you changed heads, because there's a delay in changing heads, that next track would be rotationally rotated downstream so that you'd have time to change heads and get reservoed and aligned before you encountered the first sector of the next track. But so the answer is, no, we're no longer able to actually low-level format drives. We've lost that ability decades ago, and obviously we're surviving without it.

**Leo:** Interesting.

**Steve:** And lots, lots of data. I mean, when I think about, you know, we showed, what, a

couple weeks ago the guy screaming at his drives, causing them to generate errors. I mean, that's how the servo system is trying as hard as it can to keep the head on track. But if there's vibration present, that's going to cause some trouble. And you can imagine, too, the engineers are doing everything they can to, like, not have it be like a weight at the end of a pendulum, which would be very prone to vibration. They have a counterbalance so that the head mechanism itself is not prone to a pendulum sort of effect. But even then, a rotational vibration, or a vibration with a rotational component, that is, something in the same axis as the head is moving, there's no way to prevent that from causing a problem. So, yeah, it's a challenge.

**Leo:** Well, now that we've explained that, let's talk about grown defects. This is a question from Ron Frazier in Cumming, Georgia: Could you elaborate more about how grown defects occur in modern drives? I've often wondered why drives that previously work would fail. I had a discussion on a forum recently where some people claimed the spindle bearing wear is a prime culprit. That would explain why, for example, that reallocated sector count could start creeping up, then go up dramatically in a short time. I'd be very interested to learn the best practices for monitoring a drive's health so you can exchange it while the data is still readable. I was also horrified to learn that drives do not normally do a read after write verification. Now I know I want to run SpinRite to refresh all the data a couple of times a year.

I found a cool Windows utility called CrystalDiskInfo which can monitor the SMART system, including the reallocated sector count, and send out an email when there's a problem. The only problem is it's got to run with admin rights. So I can't figure out how to install it on my dad's computer for his use since he doesn't run with admin rights usually. Any thoughts? Thanks for the podcast. Always interesting. Sincerely, Ron.

**Steve:** So in last week's holiday podcast I explain about defects, that despite our best efforts to create an absolutely flawless surface, at the molecular level almost, I mean, at the level of the size, the scale that we are recording magnetic flux, any variation in the surface causes a problem. Even just a plating thickness variation, where it wasn't exactly perfectly uniformly plated, a slight variation in plating would cause the amplitude to drop. And with today's technology, which is on the hairy edge of working at all, that's all it takes to cause that area not to be read back.

Now, we know that we have error correction technology, which we've always had, invented by IBM when they created the so-called Winchester technology, the original technology for flying heads over a spinning disk surface, because they wanted to get back the same data they wrote. And they recognized that sometimes the system was so tightly engineered they weren't going to get it back. So error correction is able to fix the problem. We talked about error correction. Actually it's the same kind of error correction, Reed-Solomon codes, which are used in the QR, the square 2D barcodes, the QR codes we talked about, that uses error correction so much that you can put corporate logos and things right across the QR code, and it just still reads it correctly anyway.

So defects are found in the factory. And in the old days, when drives were not smart, they had a list printed on the outside. This is one of the other things that happened when we moved the controller into the drive to create the so-called Integrated Drive Electronics, IDE drive, and everything since then has been that. The drive would come with its defect list built in. So no longer were we able to see what was going on. The drive would just handle those.

What happens with so-called "grown defects," if you Google the phrase "grown defects" you get pages and pages about it. This is a real phenomenon, is that there is a mechanical interaction between the head and the drive surface. They're not touching, but they don't have to touch. The technology is one where there is a surprising amount of downward pressure on the head to press against the disk. Except that the, I think it's Bernoulli, it's not Venturi, I think it's the Bernoulli effect is one where the spinning disk drags with it the air that is immediately next to it. Just because of friction, it pulls the air with it. And that air ends up being pulled in between the head and disk interface so that there's an air bearing.

So even though the head is being pressed with surprising pressure against the disk surface, I know when I've taken drives apart and felt how hard the head is being pushed against the disk, I think, wow, this is amazing that it actually isn't in touch with the drive. Once it spins up, the air is pulled under, and the drives float. They fly off the surface. And that's why the term "crash" came into common use. A head crash is literally when some disaster, like your dog running through the living room and knocking your computer over or something, or the laptop falling off the back of the couch, when some disaster creates such a mechanical event for the drive that the heads do, that they get past that air bearing and they bounce off of the surface, that's really not good. So thus the term "crash." The heads have crashed into the surface of the drive.

So what happens is, even though the head is not in contact with the drive, just the act of the pressure being put on the head, and the head and this air bearing passing over the surface, if you think about it, at a sufficiently low level that's going to create some mechanical flexure of the disk. There is essentially - the disk is not absolutely rigid. It can't be. So the head will actually dimple the disk microscopically as it's flying over it. And over time that constant flexing does cause some wear or variation of the surface.

Remember that we're not putting fewer bits on any disk than we can. No engineer would, I mean, the engineers actually want to. Management doesn't let them. Management says, no, those other guys, they've got the density up to this. We've got to compete. We need it to be cheaper and higher density. And the engineers say, no, we can't do that. And of course management says that's what you always say, just like Scotty in the engine room. So go away and figure out how we can double the density again. And sure enough, a couple weeks later, the engineers come back and go, well, okay. I think we know how to do that now.

So this has been the history for a long time. So the point is that we are always, always on the edge with this technology. We're putting as many bits as we can possibly justify onto the surface. So over time there will be problems. Now, one of the beauties of the IDE technology is that it's no longer up to us to manage the growth of these defects. And this is one of the main beauties of SpinRite is that - and I've talked about this before. A drive doesn't know it has a problem until it tries to read a sector, and its ability to read the sector is no longer just black and white because, thanks to error correction, and in fact thanks to the prevalence of error correction, many sectors are requiring error correction in order to read back.

And so the drive monitors how much correction a sector requires over time. And over time the latent defects that were not detected by the factory, or that were just so small they were easily corrected and weren't a concern, they grow. And at some point they get so big that the drive is no longer able to comfortably correct them. It worries that, if it got much bigger, the defect would outgrow its maximum correctability.

So at that point, all by itself, the drive says, okay, we've got to move this somewhere else. And that's the sector relocation. It marks this region as unusable for now, for ever

on, and it takes a spare sector, moves the data there, and essentially arranges so that, when the software asks for it, it'll go look in the relocated location. So that's how defects grow and how these drives, which are increasingly smart, are able to fix it. And of course SpinRite's benefit is that it's able to come along and get the drive to recognize the problems before the sectors become unreadable or uncorrectable, as it's called, so that's that benefit.

His last little question was about the SMART system. I learned a lot about the SMART system. What happened was - SMART is Self-Monitoring Analysis and Reporting Technology, SMART. And it's when we got the IDE drives that manufacturers like Compaq that were big users, consumers of these drives, they said we need some way of knowing what's going on in there. Now that you've moved the controller in there, you've got this integrated drive electronics, this IDE. It's a black box. How do we know how good it is, how solid it is, how long it's going to last? We'd like to know, before it dies, that it's getting kind of flaky in there. But thanks to having moved the electronics in there, we can't see into that anymore. We need some visibility into the drive. That's what the Self-Monitoring Analysis and Reporting Technology gives us, theoretically, is an API, a means of asking the drive about things going on inside. It's a classic case of politics. The manufacturers did not want to provide this information.

**Leo:** Oh, interesting.

**Steve:** The OEMs, yes, the OEMs desperately wanted the information. So the manufacturers gave them as little as possible, as little as they could get away with, and lots of excuses for why it was impossible to do more, which was complete, uh, malarkey.

**Leo:** Malarkey, there you go.

**Steve:** Malarkey, that's what it was.

**Leo:** Bull-pucky.

**Steve:** So the problem with SMART is that - and I learned this when I added that technology - SpinRite 6 is the first version of SpinRite to dynamically monitor the drive's SMART system while it's running. And what I learned was that the SMART system is only useful when the drive is under load, that is, when it's doing work. And that's the beauty of SpinRite's use of the SMART system. There's a SMART analysis page in the SpinRite UI which shows you in real-time, for example, the amount of error correction the drive is doing per megabyte of data read. And it shows you the high point, the low point, and the average over time. So you're able to judge, literally in an analog fashion, judge the quality of the current quality of the drive when it's doing work.

The SMART system means nothing when you're not asking the drive to read things because it's only in reading that it has a potential problem. So it's one thing, it's sort of nice to have the SMART system around in the background. But unless you actually are watching it while you do a scan across the drive, it's not going to tell you that much. And of course the manufacturers know that. They're like, they're not wanting to actually demonstrate that, like create a means for judging drives, because then manufacturers would reject some of them.

**Leo:** I've often wondered why SMART was so dumb. Now I know.

**Steve:** They didn't want to do it. Yup, they wanted to keep it a black box. Don't worry about the man behind the curtain. We'll take care of getting your data back.

**Leo:** How interesting. And disappointing.

**Steve:** Yeah, politics.

**Leo:** I guess you have to do it at the low level. You couldn't write a third-party app that would give you that information.

**Steve:** No, exactly. There just isn't - we're now dealing with the so-called ATAPI, that's the AT Attachment Packet Interface is the formal name for the spec. And I think we're at version 8 now, and it's taken decades to get here, because they update it for SD cards and thumb drives, and it's got all kinds of stuff now, and a bunch of legacy stuff that is dragging along from two decades ago, you can imagine. But that's the interface. And the drive makers actually even support the minimum of that. There's, like, there's a whole ton of really cool stuff.

And it's like, for example, the drive wipe was a recent addition to ATAPI, and we've talked about that, where if you spared sectors that were defective out of service, they still contain some data. So if you just do a overwrite of the drive, you're no longer able to access the spares, that is, the spared-out sectors that were defective. And of course we've talked about this in the context of SD drives, or SSDs and thumb drives, that as we take sectors out of service, we can no longer get to them through the user interface. Well, the manufacturers understood that, or actually the ATAPI committee understood that, so they added a low-level wipe command that does a much better wiping of the drive than just writing zeroes to it. It's able to go in and kill all of the spared out sectors, as well. And it's taken manufacturers a while to get around to doing it because they just sort of don't feel like there's a big need.

**Leo:** We don't need that.

**Steve:** Yeah, well...

**Leo:** Nobody needs that. Come on, who would need that? Question 3, Brian Phillips, Hammersmith, London, U.K., says that Steve is more famous than he thinks: Steve, I'm watching your Christmas special, "Back in Time," on my Apple TV. I got to the part where you were deep in discussion about hard drive sectors when suddenly my partner enters the room and says, "What are you watching?" Before I could answer, she then belts out, "Is that Borat?" Well, I laughed so much. I could see the resemblance immediately after she said this.

**Steve:** Oh, great.

**Leo:** I just thought I'd drop you a note and let you know you're a Hollywood star, Steve. All the best to you and Leo, and keep up the good work. Yes, he does look quite a bit like Borat.

**Steve:** Sacha Baron Cohen.

**Leo:** I don't think you have that suit, though.

**Steve:** No. Love our listeners.

**Leo:** Yeah, they're...

**Steve:** And their partners.

**Leo:** Maybe you do have that suit. All the best to you and Leo. Keep up the good work. Happy New Year, Brian. Philip, Taunton, Southampton, England wonders about multiple keys, please: Long-term listener of your show. I always enjoy your thorough discussion on how things work. One thing that might be worth looking into is PGP whole-disk encryption, as it has a special option that allows you to have an additional decryption key for the disk, meaning that two different keys can decrypt the same disk. Oh, that's interesting. It would be interesting to know how this works and whether it has any implications to the security of the disk.

**Steve:** So, yes. PGP does allow you to establish a second key. What it's actually doing, though, is not quite what it sounds. The way whole-disk encryption works - and this applies to the encryption built into encrypted hard drives, that is, hardware-encrypted hard drives, that is typically supported at the BIOS level, where a BIOS password is given to the drive to unlock its ability to decrypt the surface, and TrueCrypt, which is the whole-drive encryption tool that I've looked at carefully and we've talked about on the show many times. The idea is that you would like to give the user the ability to change the password that is used. If you think about it, you can change the password on these drives, if you want to. Well, that means that the password is not what's used to encrypt the drive because, if it were, you couldn't change it. Or you'd have to...

**Leo:** Oh, good point.

**Steve:** Right.

**Leo:** So it's a way to unlock the actual encryption key.

**Steve:** Exactly.

**Leo:** So don't confuse encryption key and password.

**Steve:** Exactly. There's a level of indirection, we would call it in programming terms. The idea is that you generate a big blob of very high-quality pseudorandom bits, and that's used by TrueCrypt or PGP or the drive. That's used as the actual key. That data drives the encryption of the surface. And it's super high-quality random, so you're good there. Then your password is used to encrypt that. So what that means is you can change your password anytime by reencrypting that under a new password. And, presto, now that's the password you need for decrypting that pseudorandom blob, which then gives the world access to your drive. So that tells us instantly how you allow two passwords to do that. You simply store two versions of that blob of pseudorandom data encrypted under each different password, and then one of the passwords will decrypt one of those blobs. So no significant implication on the actual technology of security. Of course you could argue that the more passwords you have overall, sort of in an absolute fashion, the more passwords you have, the lower your security is because there's more ways to get to the drive. But still it's a nice hack. And we'll be talking about this when we talk about this recent...

**Leo:** Hack.

**Steve:** ...alarm over a company that has decrypted these things in Question #10.

**Leo:** All right. Question 5, Alex Waters in Amherst provides a useful wakeup call. In Episode 383, Question 7 asked if a company installed a root CA, Certificate Authority, in their employees' machines named Equifax, if that could live side by side with the real Equifax cert, and what the lookup chain would be like if that were done, and a lot of other cruft. Cruft because it's a moot point. The company can wipe out the root CA store, replace them with 200-odd CAs named identically to the ones our browsers trust, then dynamically regenerate the chain of trust at the router, inspecting the foreign cert, generating a clone, then signing it with the fake CA that matches - or, in the case of Google, generates the intermediary CAs, sign those with their fakes, and then generate the site cert and sign it with an intermediary, and on and on and on and on.

Since they don't have to be cryptographically secure - that went out the window the second the company decided to decrypt everything you're doing - this is a cheap operation. Costs you nothing. I mean, we've got a 256-bit symmetric key that's all zeroes with some metadata attached. They could even reuse entropy over and over if the former didn't work. As an even greater point, the company owns the hardware. They could, although unlikely, write a utility that hooks the OS GUI controls and overrides them. To that end, they could see that you're in IE, hook the address bar, see that the string there starts with HTTPS, change its underlying value to HTTP, but still render it as HTTPS. Ooh, that would be mean. That would be evil. It's their hardware. They can do what they want. From a personal privacy perspective, someone else's computer is, of course - I guess this is really the real point - a totally unreliable platform.

**Steve:** And I thought that was a good wakeup call. And Alex, of course, is absolutely right. The moment you are using, as he says, somebody else's machine, the only way

they could install a fake Certificate Authority in the first place is through some means, presumably IT department did it, and you sit down at your desk, and that's what's in front of you. So I liked this because it steps back a little bit and says, yes, Steve, it's nice that you explained how the technology works, but let's remember that none of this matters because it's not your machine, you don't want someone to do that to your machine, but they have the right to do it to theirs. And so the fact is, if your company is up to these shenanigans and is determined they want to spy on you, they will be able to do so.

**Leo:** Yeah. And we should point out it's perfectly legal. Courts have ruled again and again that companies don't have to tell you, even, when they're doing this. They can just do it.

**Steve:** It annoys people, but that's their right.

**Leo:** Ricardo Ramirez, Orlando, Florida wonders about card apps and mobile wallets: Steve, blah blah blah blah, since Episode 1, blah blah blah blah blah blah blah. You know what that all means. It's ditto. It's ditto. I've been thinking about pinning - of pulling my wallet - putting, I'm sorry, I'm thinking about putting my wallet on a diet. By this I mean removing unnecessary cards, like my bar-coded library card and a couple of customer rewards cards. I've been considering moving those cards to one of the many card apps for my iOS device, like 1Card or CardStar. Apple even has that capability on iOS 6 called Passbook.

One thing worries me though. As often said online nowadays, if you're not paying for it, you're the product. How much can we - by the way, you're paying for it. You bought an iPhone, didn't you. How much can we trust these apps and companies? I don't want this to become a large honeypot of my personal information for greedy marketers. I also always err on the side of preserving my privacy. I just read the License Agreement for CardStar, and Item No. 2 almost made me break out in hives. Ooh, I've got to read that now.

Could you please look into these apps? If not, could you mention on the podcast in case someone has a good, security-minded, preferably TNO solution? I am on iOS, but I'm sure others would also be interested about Android solutions. As always, thanks for the great podcast. Greetings from sunny Florida, Ricardo. And Happy New Year to you and Leo.

**Steve:** So, Leo, this is one for you.

**Leo:** Well, I don't know these other two programs. And I agree, if - I don't know, 1Card and CardStar, if they're free. But, okay, first of all, as you will, I'm sure, aver, you could even have encryption on these things, but the company always has access to that data if they want to put a backdoor in there.

**Steve:** Potentially, yeah. If you don't trust them.

**Leo:** I would say that, if you're going to trust somebody with a wallet, you would trust Apple. By the way, you are paying for that. You're buying an iPhone. You're giving them a lot of money.

**Steve:** Yeah. I think you're right. Passbook probably is - given that it has the functionality that you want, I guess I would wonder, why go with a third party?

**Leo:** The other issue, of course, is that, even CardStar or 1Card, I don't know who these are, but it would be so damaging to their reputation. And plenty of people watch, use Ethernet sniffers to watch the traffic. If there were some sort of nefarious traffic going on, it'd be hard to hide that. And ultimately, if somebody found out, it'd be devastating. Maybe less so for CardStar and 1Card. Absolutely devastating for Apple.

**Steve:** Yeah, and so I would say never for a moment consider using some app you've never heard of from Russia or China.

**Leo:** For sure, yes.

**Steve:** I mean, absolutely stay with a mainstream, well-vetted company with a reputation because that's, I mean, that's the value to them. They cannot have that blown. So some random app that offers all kinds of features, people get seduced all the time, get suckered into this. But it is dangerous. So just say no to that.

**Leo:** And I would point out that there is something for Android. It's called Google Wallet. And again, Google has some incentive not to get in trouble, and makes money other ways. Unfortunately, Google Wallet doesn't work on a lot of Android phones. And neither Google Wallet nor Passbook will do exactly what you want them to do, necessarily. They may not store all the cards you want them to store. Apple, for instance, still doesn't support credit cards, although will at some point, I suppose. I'm just going back...

**Steve:** And you wonder if they just don't want to.

**Leo:** I think they may not. That may be why; right?

**Steve:** They don't want the liability.

**Leo:** I'm just looking at the CardStar license agreement. I don't see anything particularly horrible.

**Steve:** Well, don't break out in hives because we need you to read a few more questions.

**Leo:** Oh, Borat, whoa. Where did he come from? Okay. Back to work. Kent, at an undisclosed location someplace on Earth, asks us, please revisit UPnP. Hey, guys. (Insert normal gratuitous comments here.) It's been some time since you last really talked about Universal Plug & Play, and I do recall what was said back then. But as a new PlayStation 3 owner, I'm wondering what the state of UPnP is today? The PS3 - the Xbox does, as well, I might point out - wants UPnP enabled. Some games won't run or are limited if UPnP is disabled. Is it safe to turn it back on for my PlayStation? Or just avoid those games and features of the PS3?

I'm using a D-Link DI-604 router, blah blah blah. I can go on and on on that one. I don't think it really matters. He wants to know which is better, a WRT54G with DD-WRT installed or a D-Link DI-604 router using stock firmware. He's also thinking about using pfSense on an old PC, and wants to know which of those three options would be best. All my PCs run Windows 7 or XP with Avast, SeaMonkey with NoScript, Thunderbird instead of Outlook, and I'm the only person using the goodies on my LAN. I operate on a semi-paranoid level of security - I would say so - and I'm no newbie to computers, going all the way back to TRS-80 days, which I still have and still works. Thanks.

And then of course I should read this second one from Russ Greeno in Buckingham, U.K., who also says: You did a huge podcast on UPnP in the past and how it's a security risk if left on. What about Extended Security UPnP settings? These are something the new routers come with. Is that any better? Russ Greeno.

**Steve:** Okay. So we haven't talked about Universal Plug & Play for a long time. Kent, I was glad you acknowledged, Leo, that he is very security conscious, and he's running AV, he's using less targeted systems, he's using SeaMonkey instead of Outlook, and he's using Thunderbird, really taking care of his systems. The only danger is, that we're aware of with Universal Plug & Play, is we absolutely know that malware exists that now leverages Universal Plug & Play to maliciously open ports on a router that causes the beautiful and natural firewalling that a NAT router provides to become porous.

Years before this appeared, on this podcast, we said this is going to happen. I could not believe that Universal Plug & Play was designed with no security. Just nuts. And so that's why I've had my UnPlug n' Pray utility around for so long. It's just like, okay, unless you know you need it, turn it off. It's just one of those things that is a problem. I do know that, at least for Xbox, you can manually configure static port mappings through your router to your Xbox that essentially manually do exactly the same things that Universal Plug & Play does easily.

But I also don't think anyone should get too worked up over this. I mean, he's, like, avoiding games and features of the PS3 that use Universal Plug & Play. No. You're in control of your network. The big danger is, if something gets in, then it's able to, like, open the front door through your router to allow other things to connect to you. That's not a good thing. But again, you have to have something in your network first. You have to have malware there which is using and leveraging Universal Plug & Play before that happens.

**Leo:** Right. So it's a hole only for malware that exists on your LAN already.

**Steve:** Right.

**Leo:** It's not a hole that allows inbound malware. It just would allow a malicious program, without warning you in any way, to open ports on your router and say come on in.

**Steve:** Yes. Basically the same features that are available at the user interface of your router are available programmatically. And what's really been controversial...

**Leo:** The port-forwarding features or the opening, yeah.

**Steve:** Yes, exactly. And what's really been controversial is that many routers don't even show you what Universal Plug & Play ports have been forwarded. They just sort of do it secretly. I mean, really, that I could never understand, that they wouldn't provide you a means of seeing what has been done to your router. But they don't. And Russ Greeno's Part 2 of this, extended security is something I have not yet researched. But I've got it on my list of things to look into and talk about. So I will definitely see what, I mean, I'm delighted to hear that some routers are offering extended security. I just don't know if it's any good or how it works and what it does. I will find out.

**Leo:** You know it was Microsoft that invented UPnP, and I think it was for the Xbox because the - and the reason this is an issue on gaming machines is if you want to play games with other people, you have to allow incoming, uninitiated-by-you incoming connections. So this UPnP opens ports and says, yeah, yeah, these games can be played. If you just were opening the connection yourself you wouldn't need this. So it may be, I don't know, but you might investigate. Maybe I can get some of the functionality. Only I can start a game. I can't join a game somebody else has started, that kind of thing. I don't know. It's just a guess.

Anyway, our next question is on punch card. It says - no, it's not. Somebody sent me punch cards. But you know what's funny about these, because of their historic nature, do not fold, spindle, or mutilate, those of us who learned programming in the old days...

**Steve:** You mean like these, Leo?

**Leo:** You have a few?

**Steve:** Oh, yeah.

**Leo:** You have some lying around?

**Steve:** I still code on punch cards.

**Leo:** You know what I like about these? They're pink, and they say "Be My

Valentine. TCU Computer Center." They've got a heart. I don't think you could probably use these. And there's an arrow, too, I just noticed.

Anyway, Mikael Falkvidd in Sweden suggests that Australian museums may not be filtering. We talked in 383 about an Australian arts museum that was apparently filtering SSL connections, according to our correspondent. He's saying they just require a login so that the user accepts their terms of service for the use of their network. To do this they redirect the traffic to their own server until - this is the key - until the user has logged in and clicked Accept. Since the user went directly to an HTTPS URL, the certificate, the museum's server certificate, didn't match, of course, the Gmail URL. But the solution would be just to go to a non-SSL page, log in, get redirected, accept the Terms of Service, and then you can go to SSL, and it will act normally. I suspect that's the case. You'd have the same thing at Starbucks or anywhere where you had to accept Terms of Service. You get a start page.

**Steve:** Yes. This was a case of me not seeing the forest for the trees. I immediately went to the technology. And a large number of our listeners said, uh, Steve, ever been to Starbucks? Oh, yeah. I support, yeah, we know I've been to Starbucks.

**Leo:** Doh.

**Steve:** Anyway, so yes, everyone who wrote in, thank you. I'm sure you're right. This was just their network wanting to get a Terms of Service agreement. And if [Andrew] had gone to an HTTP site, he would have been redirected to their page. Instead of freaking out the Google's Chrome browser, he would have been bounced over to their "Accept these terms, please," and then he would have been fine.

**Leo:** Of course. I should have spotted that, too.

**Steve:** Ah, just, you know...

**Leo:** Well, but that's what happens is we go - you said "jumped to the technology." Well, let me tell you.

**Steve:** Exactly.

**Leo:** And that might have something to do with Question 10. Jason in Indianapolis, Indiana says, "Oh my god, oh my god, oh my god. TrueCrypt has been hacked." What are your thoughts on this Slashdot article? I've been listening to Security Now! because I know this will be covered.

"Russian firm ElcomSoft" - they've been around for a long time, and they've published a lot of cracks. As far as I know they're reputable - "on Thursday announced the release of ElcomSoft Forensic Disk Decryptor (EFDD), a new forensic tool that can reportedly access information stored in disks and volumes encrypted

with desktop and portable versions of BitLocker" - that's the Microsoft solution - "PGP, and TrueCrypt. EFDD runs on all 32-bit and 64-bit editions of Windows XP, Vista, and 7, as well as Windows 2003 and Server 2008." And he provides us with the link: elcomsoft.com/efdd.html. What? What? What?

**Steve:** Okay. So, yeah. Everyone went nuts with this and ran around, I mean, I got a ton of questions about it. So the good news is nothing has been cracked. Nothing has been hacked.

**Leo:** So what does this thing do?

**Steve:** Well, for $300 it would be useful to maybe the FBI. They'll buy it. Maybe, like, private security companies that want to do forensics. What it does is it automates what would otherwise be an arduous sort of theoretical process of accessing encrypted hard drives using only the techniques we've already talked about on the podcast. For example, it supports the Firewire hack. And we've talked about how Firewire, because it uses a - essentially the Firewire is an extension of the bus on your computer. Firewire has so-called DMA, Direct Memory Access. And so a device that is designed with Firewire can plug into a Firewire port and see into all of the memory on your computer.

Well, what we know is that, while any full disk encryption is in use, that is, while your TrueCrypt volumes are mounted, BitLocker is unlocked and so forth, the keys to do that must exist in memory. They're there in use. So that means that, if something gave you access to the system's memory, and it were clever, and so these guys, I take my hat off to them, $300, you don't have to do this yourself, they have done it for you, you can use this with a Firewire interface and suck out the keys while the system is in use. That is, because the keys have to be there in order to dynamically decrypt and reencrypt the data as it moves in and out of the drive.

Now, hibernation is another issue. When you hibernate your system, it writes RAM, a static image of RAM to the hard drive. Now, if the keys were in RAM at the time of hibernation, they're still going to be there in the hibernated file. And so that's another thing this ElcomSoft Forensic Disk Decryptor does support is, if your system did not wipe the keys before hibernation, then they'll be there. It can get them out of the hibernation file.

On the other hand, none of this is news to the TrueCrypt guys, for example. There's an option in TrueCrypt which wipes the keys prior to system going to sleep or hibernating. And so you know you're safe if, when you bring your system out of standby, or you wake your system up from hibernation, if you can do nothing until you give it the passwords again because the option in the user interface is, specifically, wipe the keys. And it overwrites RAM, removing the keys, the cached passwords, and the decrypted keys that we were talking about in Question #4, removes them from memory prior to going into hibernation.

And the page, I mean, this is a little bit of typical Slashdot hair-on-fire stuff that you see there all the time because it's sort of - also I guess I'd have to blame ElcomSoft for being a little bit obfuscatory here. They talk about all of their ability to decrypt everything and complete decryption and real-time access to encrypted information and zero footprint operation. And then, when you get down to the fine print, they say "three ways to acquire encryption key."

**Leo:** They bury that, that little lead there, the most important part, yeah.

**Steve:** Yeah.

**Leo:** We need the original encryption keys in order to do this.

**Steve:** Exactly.

**Leo:** Oh, well, of course you need the original encryption keys. If I'd known that - okay, yeah.

**Steve:** Yeah. So what this has done, this makes it easier for what could have been theoretically done before to now be done automatically. They know BitLocker, PGP, and TrueCrypt. They're able to do a quick search through memory to find, to lock onto the data structures that contain this, and they can suck them out. But only if they're there. And they're only there if the system, one way or another, has them in memory and is using them actively. And if that's not the case, they can't help you.

**Leo:** Right.

**Steve:** So nothing to worry about here. Nothing to see.

**Leo:** Whew. It's the same as it ever was.

**Steve:** Right.

**Leo:** Steve Gibson is at GRC.com. That's where you can ask questions for future Q&A episodes. We do this about every other episode. GRC.com/feedback. It's got a form and everything. Don't send him email. You don't even know how to send him email, trust me. I don't know how to send him email. It's a magic process. However, you can go to GRC.com/feedback and ask a question.

**Steve:** Please do.

**Leo:** Maybe he'll use it on the show. You can follow him on Twitter, he's @SGgrc. If you're at GRC.com, you should certainly pick up a copy of SpinRite, the world's finest hard drive maintenance and recovery utility. You might also want to take a look at some of the freebies over there, not to mention he's got special versions of the show, the 16Kb audio, which only Steve has. He does it all by himself. He grinds it down to a fine powder. The transcriptions, which Steve pays for, from Elaine, which make it easier to follow along, also searchable, which is nice. We have the larger

files, the audio and video files of the show, at TWiT.tv. You can find that there. Watch live, 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 1900 UTC, every Wednesday. Or download after the fact. Best would be to subscribe so you don't miss an episode. In fact, you really want an archive of these. I think a lot of people save every episode, all 300, what is it, 385?

**Steve:** Imagine, imagine how many total bytes in the universe are being consumed by archives of Security Now!.

**Leo:** Well, if it's the 16Kb version, it's barely a megabyte.

**Steve:** There you go.

**Leo:** Or something. I don't know. It's barely a gigabyte, probably.

**Steve:** I will say that I ran across people wondering or talking about their dog problems in the Security Now! feedback email. The problem is, I get so much feedback, I mean, hundreds of pieces of email, I never have time to read them all. I read enough to generate our Q&A episode, basically. So there is a different feedback page for the Quiet Canine project, and that's GRC.com/tqc - as in The Quiet Canine - /feedback. So instead of GRC.com/sn/feedback, for Security Now! feedback, it's GRC.com/tqc/feedback for The Quiet Canine feedback.

**Leo:** You've got to start calling it the new name, Dude.

**Steve:** Well, my feeling is that the Quiet Canine is sort of - it's a statement of the goal and the project.

**Leo:** Treble Shooter is just the brand name.

**Steve:** Treble Shooter is the name of the device.

**Leo:** Trademarked. Copyright. All rights reserved. Keep your mitts off.

**Steve:** It's just a nice name.

**Leo:** I like it. It's funny. All right, Steve. Thanks so much. Have a great week.

**Steve:** Talk to you next week.

**Leo:** We'll see you next week on Security Now!.

**Steve:** Thanks, Leo.