Listener Feedback Q&A #51

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-165.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-165-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 165 for October 9, 2008: Listener Feedback #51. 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!, the show that covers your security online, your privacy online and off, everything you need to know about locking systems down. Security expert Steve Gibson is our guru. He joins us every week from his fortress of security in Irvine, California, home of SpinRite and GRC.com. Good morning, Steve.

Steve Gibson: Hey, Leo. Great to be back with you.

Leo: Good you can be here.

Steve: Glad that the world is still in one piece. The market today is pretty much - it dipped down to about -243 at one point this morning, but it's back up a little positive now. So after yesterday's harrowing 508 point slide...

Leo: Wow, yeah, no kidding. It's funny, I was at the gym, working out. And as the

market was going up and down. And I kind of - I felt like as I was running, I was actually like the little chipmunk, getting the market back on track as it was...

Steve: Bouncing on your ball.

Leo: Yeah, exactly. It's kind of fun to watch the - I leave the sound off, so I don't exactly know what's going on. But it's fun to watch the market go up and down as I...

Steve: Oh, that's fun.

Leo: It's funny how, in the gym at least, and I imagine this is true even in bars and other places, CNBC has replaced the traditional golf channel, the news channels. Everybody wants to see what the market's doing. It's kind of a big story.

Steve: Yeah.

Leo: It's a fascinating story. So today we're going to make - I think we are going to make up our question and answer from last week?

Steve: Well, yes, we're going to succeed with the last week's interrupted Q&A, finish that. We've got a whole bunch of good stuff going on. I have a lot of preamble goodies, and about the first half of the Q&A were not from last week, they are responses to the Sockstress episode that we did. And we've got a really great last one, like, I don't know, what I called the really bad news or something of the week award. Someone wrote in and notified something. And I also - so I want to talk a little bit, I want to follow up at the top of the show about my conversation. I had a nice telephone conversation with the guys who came up with the Sockstress deal.

Leo: Oh, good. Oh, good.

Steve: I'll talk about that. And we've got a lot of weekly news. None that involves Microsoft, for a change. Actually it doesn't. Just all kinds of other stuff. Some follow-ups from last week, some new stuff, and then our Q&A. So...

Leo: So let's get to the news of the day.

Steve: Yeah. Well, first of all, following up on last week, the Q&A, about the first half of Q&A is sort of detailed listener feedback from last week. So I don't want to preempt that. However, what I did want to say is that, after we finished the audio and had it assembled, and in fact once your team had edited it and provided it to Elaine for her to begin transcribing, I sort of felt badly just that I had been really tough on Robert and Jack, the two security researchers. I mean, I don't think I was unfair; but I was, you

know, I was really tough. And so I wrote to Robert a pretty lengthy note explaining that, you know, basically what I had said in the podcast, that I was disturbed that they had said as much as they had in their little interview that they did and, you know, that I'd been pretty tough on them.

Leo: You gave them more warning than they gave Microsoft, let's put it that way.

Steve: And anyway, so I was delighted to receive email back almost immediately saying - oh, and I sent them a link to where they could listen to the audio ahead of time. I said I'd rather you heard it from me than you heard it from other people what I said because I wanted to give them an early heads-up, just sort of as a courtesy. And I got email back immediately saying Steve, let's talk. Jack and I are listening to the podcast now. I'd like to give you a call when we're through. And I suggested Skype, and they said, nah, phone's better, and I said okay. So we set that up. And they...

Leo: After all the security news about Skype, they may not trust it.

Steve: Well, especially we got some of that today, too, in China. I'm sure you heard about that.

Leo: Yes, I did, yeah.

Steve: So the podcast was 90 minutes, and I waited for that. And then sort of wondered, like - oh, and in our email back and forth I said, if you're still interested in talking with me after you've heard the podcast, then by all means you can give me a call. Well, they did call. And they were chagrined about what I had said last week because they felt like it further amplified the chagrin they already felt over the fact that this had all become such a big deal. We talked for a long time. I talked to both of them. And essentially - and we agreed to go off the record, so I'm not going to discuss any of the technical stuff that we shared.

But what I can say is that they never really intended this to happen. They wanted to do their presentation, as they're still going to, at the upcoming conference in the middle of the month, the middle of this month of October. And it's funny because, after listening to what I did last week, they went back before calling me and listened again to their interview, sort of feeling like, wow, we didn't think we said that much. And of course there's lots now that has been out on the 'Net in this intervening week that has happened. Again, lots of misunderstanding and various people, as happened with Kaminsky's DNS issue, guessing, like, what kinds of things might be going on.

So they were - I just wanted to sort of set the record straight. They feel like it got away from them. They didn't intend it to. They're certainly not seeking press. They're not trying to do any self-aggrandizement at all. They're really unhappy that this got - like it came to the attention of the RSnake site, which then got picked up by Slashdot. And then of course that's where I found out about it, and it just went crazy from there.

Leo: They probably assumed, as many do, eh, it's just a podcast, nobody's going to

hear it. Or maybe they said, yeah, maybe they said more than they had planned to, just as the conversation rolled on or...

Steve: That's what I think did happen. I think as the interviewer sort of seemed to be not understanding it - there was some dialogue I've seen that indicates that he really does know, the interviewer knows a lot about this; so he was just attempting to draw out more information. So, and for whatever. The good news is that - and we've had some email correspondence since our telephone conversation. Everybody who should now be talking to them is talking to them. So...

Leo: Good. So did they not say enough for somebody to implement it? I mean, I got the impression that you certainly understood enough to say, oh, I could implement that.

Steve: Yeah. And no. And in fact I've seen now other things on the 'Net that I will talk about in the Q&A section of today's podcast, the first half of our Q&A, where it's clear that people are understanding what it is that they have found. And...

Leo: I don't want to tip your hand, but are we seeing exploits already?

Steve: No. Although, in a bizarre turn, because I've been out Googling "Sockstress," kind of trying to track the story and understand what's going on, there are postings in Programmer for Hire sites, trying to hire programmers to write...

Leo: Please write this for us.

Steve: Yeah. And I'm thinking, okay, wait a minute, you know, sort of like if you're - it's not the case that somebody who could write this would be responsible enough to use it responsibly. But the idea of somebody saying, hey, I want you to write a really bad Denial of Service attacking tool for me. It's like, okay. Anyway, that's out there. So anyway, I just - I came away with a good feeling about these guys, that they certainly know what they're talking about. They definitely found something, as I knew they had from what they were saying. And as you remember, I was wishing that they had said less. But it just sort of - it got away from them. And they're unhappy that it generated the attention it did. On the other hand, it's not clear that they were having success with the major vendors whose attention they wanted to bring it to until this happened.

Leo: So they had tried to talk to people before the conversation.

Steve: Yes. They had attempted to.

Leo: Ah, okay.

Steve: But the dialogues had not developed. And apparently they have now. So...

Leo: Yeah, I bet they have.

Steve: ...that something good came from it. Okay. So we have some follow-up news and some new tidbits. In follow-ups, actually it was somebody in the GRC newsgroups posted a nice little blurb from Google's official Chrome blog. And it reads, from Google's Chrome blog, "Google Update is automatically uninstalled on the next update check, typically every few hours..."

Leo: Ah, that's the trigger.

Steve: "...after the last Google product using it has been uninstalled. The Google Update team is working on functionality to allow Google Update to be uninstalled immediately upon the removal of the last app."

Leo: Okay.

Steve: So they've felt the pressure from this delayed self-removal, and they're going to change the behavior so that when you remove the last Google thing from your system, Google Update is taken out immediately, as well. So, and we talked about it, but there was also lots of other buzz on the 'Net about it. So it's not - I think it was the person who was posting said, way to go, Steve, thanks for you and Leo bringing this to - making enough of an issue of it that Google had to take notice. Certainly we probably participated in that. But the 'Net is a big place now. And I'm just glad that Google is responding because I think there's no reason they couldn't have this thing see that it's the last Google software that's around and just remove itself, so.

Leo: Right, right.

Steve: And, for what it's worth, as I found, it is not - you don't have to have an update check occur. You just have to have Google Update try to have an update check. So whether it's able to contact the mothership or not, even if you have an off-the-'Net machine that was once on the 'Net when it got Google and Chrome and so forth, once it tries to do an update, even if it fails, it says, okay, well, oh, look, there's nobody else here, I guess I'll leave, too.

Leo: It's getting lonely. I'm out of here. I like it.

Steve: Also, updating the continuing-to-evolve RealDVD saga, some interesting news. RealDVD has now been removed from the market.

Leo: So this is RealNetwork's attempt to create a program that wouldn't, well, from

your point of view as a user wouldn't crack DVD encryption. It has to, of course, to copy the DVD to the hard drive, but then puts its own encryption or protection on there. And the movie industry sued Rob Glaser and RealNetworks. And Rob has countersued.

Steve: Yes. Actually it's - I wouldn't say that it has to crack the encryption to copy to the hard drive because, as we know, you can rip DVD files. DVD files are just files.

Leo: They just won't play.

Steve: They're just encrypted, yeah, the files themselves are encrypted. But they do copy just fine.

Leo: But isn't that why the judge said, and the lawsuit, the premise of the lawsuit was, oh, you have to reverse-engineer our copy protection, you have to violate the DMCA to do this?

Steve: No, because they - remember, as a player, anything that is a DVD player has to be...

Leo: Has to be able to do that, yeah.

Steve: Exactly. It's got to be able to decrypt the DVD, the CSS copy protection. So there's a nice piece on Ars Technica that I'll share with our listeners because it explains exactly in, like, the right language where this got tripped up. Less than a week after RealNetworks launched its new DVD ripping and archiving product, RealDVD, a court has ordered the company to temporarily suspend its distribution. A visit to the RealDVD website, which is RealDVD.com, reveals a message from Real stating that the product is unavailable. "Due to recent legal action taken by the Hollywood movie studios against us, RealDVD is temporarily unavailable," reads the site. "Rest assured we will continue to work diligently to provide you with software that allows you to make a legal copy of your DVDs for your own use." From the moment Real first announced RealDVD - this is back reading from Ars Technica - the company was aware that there would be legal questions about the product, but seemed to think that everything would be fine since the company said it had, quote, "licensed the DVD technology for a legal right to play back DVD content."

Leo: Oh, interesting. So they had paid for that.

Steve: Yes. And just like there are DVD players, software-based players, as we know, that'll play DVDs. But you have to have the DVD disk in the machine from which it plays. And that's the difference. The same day RealDVD was released, however, the company found itself in hot water with the movie studios. Real preemptively filed a lawsuit, which is what we did indicate last week, on the day RealDVD was released in response to threats from shareholders, to which the Motion Picture Association of America, the old

MPAA, responded with its own lawsuit. The MPAA claims that Real has violated DMCA anti-circumvention rules with RealDVD, even though that product copies DVDs to a hard drive while leaving CSS encryption intact. It turns out that the association, the MPAA, is picking a nit with the fact that RealDVD doesn't require an actual disk to be in the drive when decrypting a movie for playback, therefore allowing users to rent, rip, and return movies.

Leo: Right. Okay.

Steve: By not requiring the disk to be in the drive, Real supposedly makes circumvention of the purpose of the encryption possible. So circumvention of the purpose of the encryption, even though it doesn't appear to circumvent the CSS encryption itself. You know, it actually uses - it decrypts CSS, which it has a license to.

Leo: Well, really that's poor design on the part of CSS. It should check for a disk, and then they wouldn't have this problem.

Steve: Well, it's funny, too, because the executive vice president and general counsel of MPAA says, "RealNetworks' RealDVD should be called 'StealDVD.'"

Leo: Oh, geez. These guys. They have to really learn, this rhetoric does not work. It does not play.

Steve: So it says a temporary restraining order has been issued against Real while the judge takes time to review all the available documents. A decision will apparently be made as to whether the suspension will remain in place and for how long, a Real representative told NewTV. Given the tenacity of the movie studios when it comes to copyright infringement, however, the MPAA isn't likely to let the restraining order be lifted without a fight. At least Real has one thing going for it. The company managed to have the case moved from the Central District Court of California in Los Angeles (Hollywood's backyard), to the North District Court of California, which may give it a fighting chance against the movie studios.

Leo: Well, to play devil's advocate, because you know I hate DRM, they do have a point. I mean, you could use this to effectively circumvent...

Steve: Oh, yeah, rip and return.

Leo: Yeah. So I can see why they'd want to plug that hole.

Steve: Yes. Now, what's funny is that DVD Decrypter, which is freely downloadable all over the 'Net...

Leo: Then why aren't they suing them? You know why, because they're in some other country.

Steve: Well, yeah. And they're, you know, it's software, and it's out, and it's around, and there's no one to sue. In fact, they did stop its support. DVD Decrypter was being actively maintained for many years. And they came down on the guy, and so he said, okay, I'm abandoning this, I'm disassociating myself with it. I'm not doing anything more. But, by the way, it works perfectly. So, and I've taken advantage of it for legal purposes. I own content. I've got the original disk. But sometimes it's nice to be able to repurpose it, in no way that causes lost sales of the DVD.

Leo: Well, but that's the spirit of the law, but not the letter of the law. The letter of the law is you can't copy them.

Steve: I know.

Leo: It's a nasty little thing we get into when we start to...

Steve: Okay. In new tidbits...

Leo: Yeah.

Steve: We've got a little - we have some happy news about something we've talked about several times. We've talked about how the new DHS, Department of Homeland Security, guidelines for the level of search and seizure of electronic media when you cross the border, remember we've talked about it being...

Leo: Oh, we've talked about it on TWiT endlessly. I mean, it's just nasty.

Steve: Yes. Well, there is now in Congress, in the U.S. Congress, something called the Travelers' Privacy Protection Act, which is current legislation. I don't think it's yet enacted, but it's in place, and it's working its way through our various processes, which would require the DHS, the Department of Homeland Security, to establish reasonable suspicion of wrongdoing before searching U.S. residents' devices. And it would require the Department of Homeland Security to have probable cause and a court order or a warrant to hold the device for more than 24 hours. So at least the notion of having reasonable suspicion of wrongdoing, it raises the bar way above, you know, like a random strip search of anyone who is carrying a laptop. And they say, oh, you know, we want to - remember the issue came up because there were people who were using, for example, TrueCrypt to encrypt their drives, and the agents were saying - forcing the passwords to be revealed in order to look through people's laptops and hard drives, which will no longer be legal without a reasonable suspicion of wrongdoing. That is, they have to have some reason to believe, you know, you can't just be random Joe Citizen and have them say, okay, we want to see what the titles of all your files are.

Leo: Yeah. Okay.

Steve: And finally, actually two things. Skype has admitted to the reports that came out the week before. You remember, and we didn't talk about this last week, but some Canadian researchers found evidence on Skype servers that the Chinese version of Skype was going beyond keyword searching and actually logging the conversations of Skype users whose text Skype instant messages contained keywords, and was logging their names and transactions and IPs and everything. And so Skype has responded. And they said, yes, they've acknowledged that they didn't realize logging was going on. There's a Chinese affiliate that they work for. TOM is the Chinese company.

And of course this is my great concern with anything like Skype. It's why I've got the acronym TNO, Trust No One. Because the fact that these dialogues, these Skype connections are running through a central server means that you're implicitly trusting Skype not to be able to spy on you and eavesdrop, instead of having a direct point to point and not using a third party. So I just - that's another example of what happens when you do need to trust an intermediary. I mean, you just can't tell what's going on. The fact that they're doing a keyword search inherently means they're able to log. And Skype apparently knew that they had provided them with keyword searching capability. But in order to do that you have to decrypt. And if you can decrypt, you can keep records. So, not good.

Leo: So let's be clear. Skype calls themselves are encrypted and cannot be listened in on. Is that right?

Steve: No. No, this demonstrates it's not true.

Leo: But that's chat, not calls.

Steve: I don't know.

Leo: I'll tell you the reason I ask. Garry Kasparov, the former world chess champion and an outspoken opponent of Putin, said I use Skype all the time because it's encrypted, and I know the KGB or their modern-day equivalent would be listening in, so I make sure I use Skype. And I think there are a lot of people who feel that way. Especially since we know that the NSA is listening in on some domestic calls, as well, on the regular phone service. So you think that a Skype - you would not assume that a Skype call is safe.

Steve: Oh, absolutely not. In no way would I make that assumption. The only thing that you could assume is safe is a point-to-point VPN, I mean a point-to-point encryption where you know everything about and trust the nature of the way it operates, and you can demonstrate that there's no traffic to a third party. We don't have complete documentation of the Skype protocol. They consider that proprietary. They don't want people making clones of what they consider to be their intellectual property. So, no, I mean, so long as there is a third party involved, the fact that my Skype client is connecting to a mothership, is phoning home as it is - and so is yours, Leo. You and I

have a direct point-to-point connection, but we don't know that the encryption key, the symmetric key used isn't provided to Skype.

Now, the fact is that we do know that our traffic is going between you and me. So we know that there's no copy of our traffic going to Skype. So we can, looking at the packet flow, we know that Skype is not participating in that. So that's certainly a good thing. It may be the case, though, that because instant messaging is such low bandwidth, and like Skype will hold messages waiting for later delivery, it may be that all text goes through Skype central; whereas audio and video does not. It goes point to point where it can.

Leo: Oh, I want to add, a couple of people in the chatroom are saying what's that bill number. The bill is Senate Bill 3612, for those who want to - the bill on curbing searches at the airports. Travelers' Privacy Protection Act. So if you want to send a note to your member of Congress - your Congress Critter, as Cory Doctorow calls them - that's S. 3612.

Steve: And tell them yay.

Leo: Yeah, yay.

Steve: Okay. And one more, you know, this is just too bizarre to be true but is. Did you hear about the top secret camera from MI6 that was sold on eBay? Oh, Leo.

Leo: MI6 is a British intelligence arm.

Steve: It's like our secret service, or NSA.

Leo: Yeah. So tell us.

Steve: There's MI5 and MI6. And one is internal, and one is external; right?

Leo: Right, right.

Steve: So that way...

Leo: MI6 is external and 5 is - I can't remember. I'm trying to remember from my James Bond novels. I don't know. I'll look it up.

Steve: There was a story that appeared on TechCrunch that the Washington Post picked up. And it's around the 'Net, so it's easy to find. A 28-year - and I'm reading from this Washington Post column. "A 28-year-old deliveryman from the UK who bought a Nikon Coolpix camera for about $31 on eBay got more than he bargained for when the camera

arrived with top secret information from the UK's MI6 organization. Allegedly sold by one of the clandestine organization's agents, the camera contained named al-Qaeda cells, names, images of suspected terrorists and weapons, fingerprint information, and log-in details for the Secret Service's computer network, containing a 'Top Secret' marking. Once he downloaded the contents of this camera onto his computer, he immediately went to the police to explain the situation. The police originally treated it as a joke, but within a week anti-terror officers started investigating and demanded that he not talk to the media about the contents contained in the camera.

"Journalist and author Neil Doyle told The Sun that the contents are 'MI6 documents relating to an operation against al-Qaeda insurgents in Iraq. It's jaw-dropping they got into the public domain. Not only do they divulge secrets about operations, operating systems, and previously unheard-of MI6 departments, but they could put lives at risk.' MI6 is currently trying to track down the agent who made the mistake." Oops. "If caught, the agent could face serious legal ramifications and face suspension." Oh, gee.

Leo: I would hope so. It's the external - it's the secret intelligence services. It's the external, obviously. Holy cow.

Steve: Yup. So we talked last week about the VPN that hadn't had its configuration wiped, and when connected it happily connected, you know, phoned home and connected into some poor company's internal network. And here is, like, serious images left on a camera, sold for $31 on eBay.

Leo: Good deal, by the way, on that camera.

Steve: Ah. Yeah.

Leo: Wow. Well, you know, I think this is a sign. The Internet makes all of this stuff so much more out there and doable. And I think everybody, even the most secretive agencies have to kind of pay attention in ways they didn't have to before.

Steve: So I've got two little bits in the randomness grab bag. Scott in Pawtucket, Rhode Island wrote and said - and I ran across this when I was pulling together the Q&A for today: "I realize that your mention of the Richard Matheson novel, 'I Am Legend,' and its subsequent film adaptations, 'I Am Legend' (2007) and 'The Omega Man' ('91)," both which you and I referred to last week, Leo, "was a commercial for Audible. But I thought you and Leo might find this little tidbit of info interesting. The original film adaptation of 'I Am Legend' was done in 1964.

Leo: Right.

Steve: It starred Vincent Price, called Dr. Robert Morgan in this version, and was titled 'The Last Man on Earth.' I listen to your podcast as much for the security and technical aspects as I do for the occasional science fiction recommendations."

Leo: Not so occasional, but all right.

Steve: Yeah. He says, "'Fallen Dragon' is now one of my very favorite books, and I have suggested it to many a friend. All the best, and thanks for a great podcast."

Leo: Yeah, I knew that because when the Will Smith "I Am Legend" came out there were many articles about the various adaptations of this book. And that's what actually led me to the book in the first place was, well, I want to read the original. And it is very different, you know. And everybody has their own take on it, which is interesting.

Steve: And my last little bit of randomness is that, as you may know, Leo, photos of Kindle 2 leaked out.

Leo: Now, they look - I don't know if they leaked or if they were Photoshopped. And this is something those of us who cover the Macintosh...

Steve: Oh, you mean they could be completely fake.

Leo: Oh, yeah. Those of us who cover the Macintosh are used to this happening all the time. So we might have a little bit more radar than the average Joe or CNN. I haven't seen anybody from Amazon acknowledge that these are real. But, boy, it's the kind of thing you'd expect Amazon to do at some point; right?

Steve: Well, yeah. And in fact it was you who said that there are rumors of a new Kindle on the way.

Leo: Amazon has acknowledged that they are working on a Kindle 2. They have said that it will be out next year. But I don't know how forthcoming they've been on details. And I think so far it's kind of all rumors.

Steve: Well, and they don't want to cannibalize their current sales.

Leo: Well, that's the risk, of course. So you don't want to say it's coming out in December because nobody will buy Kindle 1.

Steve: Right.

Leo: So the pictures, the so-called leaked pictures I don't think have ever been admitted. I don't think Amazon has ever said that those are actual pictures. It's from a very big rumor site called Boy Genius Report.

Steve: Right, right.

Leo: And, you know, again, the Mac people are going, yeah, yeah, we've seen it all before. So we'll just see. You know, it looks like something that you would want.

Steve: Well, I'll tell you, Leo, I have fallen in love with my wacky little wedgie.

Leo: Yeah. I don't mind, I've gotten over the form factor issues, which are many.

Steve: Yup, yup. And the fact that it's just weird. I mean, it's a bizarre shape. But I look at this one - and this one, for the sake of our listeners, it's sort of a - it's, like, traditional looking, very smooth, very slick, rounded corners. It reportedly has, like, sort of a polished, stainless steel back, like an Apple iPod does; a keyboard in a single place instead of being broken into a left-hand and right-hand sort of thumb keyboard.

Leo: Oh, yeah.

Steve: Very different navigation buttons. What's missing from the photo - which actually, now that you mention Photoshop, raises my eyebrows - is there isn't the LCD selection stripe shown on the right-hand side.

Leo: There's a scroller, but no stripe.

Steve: Yeah. So maybe they've built the scroller into the eInk surface, or who knows what they may have done. But, Leo, I would keep what I have.

Leo: You would, even looking at this new one.

Steve: I have zero...

Leo: Did you look at the new Sony one? The new Sony one's quite nice.

Steve: With the touch screen.

Leo: Yeah, touch screen, built-in light. Also expensive. It's more than the Kindle is and doesn't have wireless.

Steve: And again, unfortunately, it's Sony. I mean, I owned two of them before the Kindle came along.

Leo: Me, too. Me, too.

Steve: Leo, the fact that this thing, I've now got - I've added a couple blogs to my four newspaper subscriptions. And it's just, I mean, I love - even if I couldn't read books, the fact that this thing is newspapers and really interesting blogs - and, boy, the selection of newspapers has exploded, and the selection of blogs has exploded. There is just so much content now. And it just, you know, and the blogs are, like, a buck or two a month. So if the...

Leo: I subscribe to one. I subscribe to Salon because I never read any of them, and I still don't even read that enough. I feel like I'm wasting my money. Are you pretty religious? You read it every morning, and you get through it, and...

Steve: I like having them there.

Leo: I like having them there, exactly, yeah.

Steve: No, no. But, I mean, but I do spend several hours in the morning reading the newspapers. And I'll jump around. I'll wonder what The Wall Street Journal has, or The New York Times, or the Financial Times. And I just sort of - I'll scan the front pages. Often, depending upon my mood, I'll spend many hours or just a few. So, but, I mean, the sense I have coming away is I really know everything that's going on that I care to know. So it's super useful. And I have a...

Leo: Super, super useful.

Steve: Super useful. I have a very short little SpinRite anecdote from a David Lisney in Hertfordshire in the UK.

Leo: Oh, I say.

Steve: And he says - the topic was, oh, it just came in through our regular feedback submission, which is why I know where he's from. He said, "On the 'SpinRite Saves Lives' subject," he says, "I used it on a PC running a paging system for the London Fire Brigade. This is the largest brigade in the world as far as the population in their catchment area." I guess that means, like, the size of the area that they're responsible for because he talks about as far as the population. He said, "The machine which was running our paging system, upon which we depend, had chugged away 24/7 for 10 years. It had finally suffered a failure of CPU fan and crashed. After replacing the CPU fan, the PC still did not boot up. I ran SpinRite on the failed drive. It raced through, and in under an hour the PC was booted and running again. My colleagues could not believe it. Ironically, this PC was due for replacement the following week. The story is not as exciting as some you have received, but it did save a lot of head scratching. Needless to say, the recovered machine was fine for the remaining week and has now been put to use in a noncritical area, still running the same repaired drive. Thanks for SpinRite. Love

the show. Regards, David."

Leo: Well, isn't that nice. Happy story.

Steve: Another nice story, yup.

Leo: A happy story.

Steve: Somebody happy with SpinRite.

Leo: All right, Steve. Let's - whoa. Are you ready for the questions? There's two of me and none of you. Let's bring you back. Are you ready for the questions here?

Steve: Brings new meaning to the term "talking heads."

Leo: Talking. I can talk to myself on the show, and in fact I do quite a bit. Let's see here. I had the email, and I had it open, then I guess I closed it. Question 1 from John Schember in Florida. He's wondering about Vista and TCP exploits. He writes: Vista includes a new network stack - we talked about this when Vista first came out - developed just for the OS. According to the latest episode of Security Now!, Vista should be vulnerable to this newly discovered TCP stack attack. Since Vista is supposed to be the most secure Windows yet, why did Microsoft design a new stack that was inherently insecure? Well, nobody knew about Sockstress; right? I mean, that's something nobody really knew about.

Steve: Well, yeah. It's one of the things that I keep seeing out on the 'Net. And I tried to make the point last week, and I want to reiterate it, and that is that this is not a bug in anyone's implementation. It's that in order for TCP to have the power and flexibility that it has, which is truly phenomenal when you consider it's a complex protocol, but it's able to - a single protocol is able to handle massively wide variations in network connection. It can handle very slow connections; it can handle very fast connections. It can handle connections with a short delay or with a long delay, that is where - remember that we're just sending individual packets, in the final analysis, individual packets from point A to point B, and routing across the 'Net can go a few hops or many hops. It can go to a satellite transmitter, up into the sky to a satellite and back down where you have really much larger latencies. And sometimes the way some of these work you have, like, downloading to you comes down by satellite, but then uploading goes back over a terrestrial link, typically modem. So there you've got very different operations. So each direction in TCP can basically run at its own speed and adapt itself to the network.

Well, in order to pull off that magic, the protocol has to first of all be very sophisticated, and it has to be very adaptive. And so what these guys have done is, by really understanding the guts of the way TCP is implemented, they've been able to say, you know, what would happen if we had one endpoint that was behaving in a specific way deliberately, and then what would happen if we had a lot of those. And so it's a combination of sort of stretching the correct operation of TCP out to sort of an edge case, and then doing it a lot. And so it's not at all that Vista's stack is vulnerable. In fact, it's

the sophistication of modern TCP stacks that allows them to be as adaptive as they are, that actually is sort of the - is the source of part of this vulnerability. But the vulnerability, unlike much of what we talk about, where we talk about vulnerabilities that are the result of design flaws or implementation flaws, this is just a sophisticated protocol that is inherently prone to abuse. And so this is what makes fixing it tough because...

Leo: Because this is in the nature of how it works.

Steve: Yes. To give an example, and this is again something already out on the 'Net, and it's something that had occurred to me, but I can say it now that it's been - it's now being bandied about. One of the things that happens with fast networks that also have latency, there's something called the "bandwidth delay product." And we've discussed this in the past, the idea being the bandwidth delay product is the size of the pipe times the delay from endpoint to endpoint. And what it essentially refers to is the amount of data which is in flight between the sender and the receiver. And so if you have a fast link, and it's got a long delay, well, then you end up with a large bandwidth delay product.

Well, the way TCP works is that it acknowledges the receipt of data back to the sender periodically. Not every single packet that comes in is acknowledged, but there are various systems. And in fact the acknowledgement routine is adaptive, as well. Well, when an application wants to send data, it sort of dumps it into the TCP layer. It says here's a bunch of stuff I want you to send. So TCP allocates buffers to accept it and then says okay, fine, to the application. Okay, fine, we'll send that for you. The application goes off about its business.

Now, remember that in their original podcast they talked about working TCP in the face of lost packets, high packet loss situations. Well, TCP, the sending end, has now accepted this data from the application. The application has gone on. The application is able to assume that the data that's given TCP will be sent reliably. That's TCP's job. Well, what this means is until all the data in these buffers has been acknowledged by the other end, that local TCP stack that is on the server side, which is typically serving large files, it has to hold that data. It can't - there's no mechanism for it saying, oops, wait a minute, to the application, can I get some of that back because the other end said it never received it? It doesn't work that way. TCP has accepted responsibility for delivering it and now must. Which means it has to hold that data.

Well, there are ways to cause TCP to believe that the network bandwidth delay product is very large, meaning that it will expand its buffers in order to hold all the data - basically it's got to hold a copy of all the data in that it could be in flight. And it has to hold it until it's been acknowledged. So you can imagine that this is a way of causing the local TCP stack to require that each instance of a connection end up with a large buffer. And so, again, there's nothing about TCP that's broken. It's the way it's supposed to work.

Leo: It's the nature of the beast. So how do you fix it?

Steve: Well, that's why this is such a tough problem.

Leo: Redesign the protocol?

Steve: Well, we can't do that, of course. Because, I mean, it's out there. It's, I mean, it is the protocol. It's taken quite a while to put it together right and get it all working right. What'll end up happening is that, to the degree that this is used in an attack - and we've got a couple questions that we'll be encountering next that question this, that is, even why it makes sense to use it in an attack. But you could imagine, and as I said last week, that the vendors of firewalls are no doubt immediately revving their products to be aware of these kinds of abuses. And so we can imagine that, with time, there will be probably preemptive protection against this.

Leo: Preempt - in other words, the firewall has to handle it, not the stack itself.

Steve: Well, Microsoft has, in their later servers, in IIS, they've got more - IIS, which is Microsoft's web server, is more finicky about connections. In IIS v5 there is a timeout for, like, a connection that's not doing anything. And I think Microsoft, I think they default to 900 seconds, which would be 15 minutes. And so that's a long time to wait before giving up that connection, before the server closes the connection preemptively. Certainly you could bring that time down. But they've got some other, in IIS 6, which is available in their later server operating systems, there are also, like there are limits you can place at the rate at which data is leaving the system. And if it seems too low, IIS will just say, eh, this seems kind of fishy to us. Let's just shut down this connection. And so there are things you can do to hopefully not introduce false positives, because you don't want to be hanging up on valid users. But at the same time you want to be less tolerant about strange connections which might be attempting to abuse your service.

Leo: Very interesting. Paul Cousins in Regina, Saskatchewan had a worm counterattack idea. He says: Hi, Steve and Leo. I've been a long-time listener of Security Now!, and your show about Sockstress gave me an idea. This is going to be a little black hat and probably bordering on illegal to actually implement this. But it's still fun to speculate. Is there any way this new attack could be used by security firms or even individuals as a way of taking worm-infested computers off the Internet?

Now, this idea comes up from time to time, I have to say. It's not anything new, it's just a new way of doing it. The idea being you leave your computer open to connections. When you identify a connection from a worm or bot-type program, you launch a Sockstress attack on them. As I say, this is probably bordering on illegal - not bordering, it is - as you would effectively be killing off computers of innocent users. But I wanted to hear what Leo and you thought of this kind of concept being applied to the real Internet, what with many issues of Internet background radiation, as you like to call it. Thank you for a great show. What is your take on that?

Steve: Well, okay. The first part of his thought, that is, that we are leaving our computer open to connections, I discussed that last week, this notion of a tarpit where you would deliberately have listening ports on tasty ports that worms want to attack.

Leo: This was the topic of the very first Security Now!.

Steve: No kidding.

Leo: The HoneyMonkeys, or whatever they were - remember, Microsoft was doing this?

Steve: Right, right. Windows HoneyMonkeys.

Leo: I'll look it up. It wasn't HoneyMonkeys, it was something like that.

Steve: It was - that's close, though.

Leo: It was TWiT.tv/sn1. And of course I don't have a very good description of it in the show notes, alas.

Steve: It was a funny episode, too, because you were funny on that one, Leo. You were wonderful, I remember. Anyway, so the idea being that a tarpit accepts a connection and then manages to stall it...

Leo: Strider HoneyMonkey.

Steve: Strider, that's what it was.

Leo: That was Episode 2, Security Now! Episode 2.

Steve: So the idea is that the computer accepts a connection and then stalls it, basically holding onto that connection rather than either saying no, I'm closed, and allowing the scanner to go scan somebody else, basically it sort of stalls it and keeps it in a so-called "tarpit." Now, the reason Paul's notion doesn't work is that we've got a connection, well, there's nothing we can really do to it without engaging it in a protocol. So in order to attack a machine, we need it to be the server, and we are the client; where in the scenario of a worm scanning us, it's looking for our open ports, that is, open service ports at our end, and then hoping to abuse a vulnerable service on our machine.

So really there's no way to switch this around on the fly. When a worm attacks us, if it had open ports, that is, the only way that Paul's idea would work would be we get a connection from a worm. We know that - we assume it's a worm. And that's one of the first assumptions you need to make. And it's not like your ISP scanning you to check for your own security or making sure you're not doing things that are against their own terms of service. And there has been some of that going on. But the idea would be you accept a connection. We know it's not spoofable. So whatever it is that's hooked to you, you've actually got their real IP. Well, then you'd have to do a port scan of them, right on

the fly. You do a port scan of that machine in order to look for any open ports, and then launch Sockstress attacks against those open ports. But again, bad idea because it's illegal, and you do not know who you're attacking, essentially.

Leo: In fact, I remember there was a virus, I can't remember which one it was, that the author said, oh, no, my plan was to get the bad guys, not the good guys. It doesn't - you're not allowed to do this. It doesn't matter. You still go to jail. And it isn't a very good idea. It isn't.

Jamie Scanlon in Venice, California was also thinking about TCP stack attacks. What a surprise. He writes: Steve, I just heard your discussion on the Sockstress issue. What would happen if the TCP stack would just not care if the client were not able to deal with a connection because the buffer was full? My logic would be, if you don't have buffer space, you shouldn't be making the connection in the first place. I know this is not in the protocol, but how much of the Internet would break if this were implemented? That's kind of an interesting idea - just ignore it.

Steve: Well, yeah. In fact, that's another setting that many servers have. I'm sure we've, you know, all of us who've been surfing the 'Net for a while have typed in a URL or clicked a link, and we've come across a dumb, blank, white page that just says "Server too busy."

Leo: Yes.

Steve: What that "Server too busy" message means is that there's an administrative limit, that is, whoever it is that's running the server, for whatever reason, said we're going to accept only this number of connections. Presumably they're saying somebody who can connect to us, we want to guarantee a certain level of service to them. And that's better to do, it's better to deny newcomers connections that we couldn't adequately service than pull the whole server down. The idea being, okay - and it's an arbitrary limit. It's like whatever, based on experience, maybe a hundred people, a hundred connections total can be hooked up to the server at this time because of, you know, bandwidth limitations or whatever. And so the idea is we'd rather just say, oh, sorry, we're too busy, try again later, than slow everybody down to the average of whatever number of connections are coming in.

So that's an interesting policy. It wouldn't break the Internet. It would just maybe annoy people. I have a high limit at GRC because we've got good bandwidth and my server's not heavily loaded down. And I've never needed to impose a limit. But I certainly could if that ended up being a better approach. And again, we've all probably encountered that as we've surfed the 'Net. So certainly some people do. They just say our resources are strained. We're not going to entertain any newcomers.

Leo: Seems like there would be a more graceful way to degrade, but maybe there isn't. Maybe you just have to say no, go away.

Steve: Well, you could say would you, you know, we can accept a connection, but how would you like being, like, the last person on the block to get any data?

Leo: Slow, it'll be slow.

Steve: But, you know, it'll be literally we will service these in first-come, first-served order in some sort of fashion. So stay around if you want to. How badly do you want what we have?

Leo: Are you in a hurry? Donald Stone in the UK really spun up his propeller beanie for this next one, Steve: As I understand the problem, the source of the attacker is verified; but the attacker then states I'm too busy to receive any packets at the moment, thus causing the server in question to start a timer to increment the time between attempts to successfully communicate with the client, the fake busy client. It is these timers running which results in the resource clogging. Surely a workaround to this problem would be to ignore connection attempts from the client IP when a threshold of failed connections has been reached, kind of like what we talked about, but also to discard the previous connection attempts from the client IP. There should be a limited time during which any rejected IP address should be blacklisted by the server in order that, should it be a genuine communication error, the client would have to wait five minutes or so, some arbitrary amount of time, before attempting to reestablish a legitimate connection. Surely this would reduce the impact of this form of attack considerably. The botnet attack scenario would still pose a risk, the big distributed denial of service attack, albeit a higher risk now, as it does anyway. That's kind of the same thing; right?

Steve: Yeah, well, Donald got a lot of value from spinning up his propeller beanie on this one. He's exactly right. This is the kind of thing that you could imagine security device vendors, I mean, there are many, many vendors of security devices. For many years when I was sort of a newbie at this, actually when I originally got set up with Verio, I said I knew I needed a firewall, and I wanted industrial strength. And they set me up with a WatchGuard. WatchGuard is a vendor of, you know, actually I think it's called the Firebox or something. It's literally a bright red painted box. And so that was my, early on, my first exposure to this. And so it's that kind of a vendor that I'm sure right now are busy introducing a series of upgrades to their products that will add awareness about these kinds of problems, if they're not already there.

I mean, there very well could be sort of degenerate TCP connection strategy and logic in this kind of border protection device. And, if so, that's an example, what Donald said is one way it could work. You literally monitor your connecting clients. And if you see some that misbehave, you decide, okay, you're bad. And in fact, now that I think about it, I remember that that machine that I had, that Firebox, the WatchGuard Firebox, did have a blacklisting facility. And I don't remember exactly how it worked. But actually I think I do now. It was by client IP. And because I remember I came up with a wacky hack that Andy, my Verio engineer, really thought was kind of interesting.

There was something that clients could do that could upset me. And what I did was I used my own raw socket system to send a packet at the Firebox, spoofing that remote IP, and contacting a service that I wanted to, like, trigger. And this thing, whatever it was, I spoofed the remote client to tell the Firebox that they were bad and attacking me, and to cause them to go into the blacklist. And so it was a way for me to add people to the Firebox's blacklist using my own logic rather than the Firebox's logic. And it worked great. I don't remember, I mean, this is in the dark ages. This is a long time ago. But so, yes, certainly that kind of approach could work in this case.

Leo: Very good. Another one for England. Score another one for the British Empire. And now moving along to Mat Ludlam in Weybridge, London. He makes another great point, this time about Sockstress and botnets. Love the show, been listening since Episode 1, favorite netcast, et cetera, et cetera. So if a botnet has 10,000 machines...

Steve: As many do.

Leo: ...then if each of them simply opened 100 connections to the same host at about the same time, holding them open would generate about 1,000,000 connections, which is probably enough to cause it to fail. Yes. 100 connections is probably well within what a Windows client machine can handle; but 1,000,000 is probably well beyond a server, particularly if it's tuned for HTTP requests, which tend to be short. Whilst I appreciate that messing around with user mode TCP stacks allows one machine to do this, surely if you have a botnet at your disposal then you don't need any of that. I'd be interested in your comments.

Steve: Yes, and I put this in because I wanted to acknowledge Mat's notion. And also there's been a ton of dialogue about exactly this on the 'Net.

Leo: Well, what he's saying is that who cares about Sockstress, there's always the DDoS. Right?

Steve: Well, and he's even saying in this case not even flooding, but just connections. I mean, a Windows system...

Leo: Oh, I see, it's not a SYN flood, it's just a million connections.

Steve: Yeah. If you had 10,000 machines, and many people do...

Leo: Many bad guys do on their botnets.

Steve: Many bad guys have 10,000-machine botnets. And every one of them could easily just open 100 connections, just say hi there. I mean, not even performing any fancy TCP Sockstress exploits, just opening 100 connections. And he brings up that that, you know, do the math, and that's a million connections. Which would be very rough for your typical website to handle, as you and I know, Leo. When we mention something live, and you've got TWiT Live listeners, it immediately brings those sites down.

Leo: Right. We do that all the time. And not intentionally, I should say.

Steve: No, and not a million.

Leo: A few thousand.

Steve: Few thousand, exactly. So I've seen this echoed in many forums online when people are talking about Sockstress. They're essentially saying there are already bots that are able to do DDoS. And they're distributed, and then they use a lot of bandwidth. And the point being, so what if there's, like, some undisclosed 'nother way of using TCP? We don't need another way. And this way doesn't allow you to spoof. And spoofing is nice if you can do it because then you're not giving away the IP of all of the clients in your - all the zombie machines in your botnet. And so big deal. It's like, eh, sort of like another...

Leo: It's just another tool in a hacker's tool box. But the point of this is that it requires - doesn't require a botnet, essentially.

Steve: Correct. Well, and in fact the counter argument, if I were to raise one, would be that there are high-value, like, gambling sites that have been subject to extortion and are the repeated subject of extortion, as we've talked many times, for example, by literally organized crime operating out of the Ukraine and various places, where in order to keep these sites on the 'Net they've moved themselves onto very expensive, very large pipes behind organizations that specialize in anti-DDoS attacks. Now, one hopes that it has not just bandwidth flood protection because, if not, then here is a new way of providing a non-high-bandwidth attack, which could be new and might succeed in pulling these sites down. So the fact that this is new doesn't necessarily mean it has no value to attackers. It's just, sure, there are existing ways to attack most sites. This is a new way that might give a new way around. And again unfortunately, you can imagine where there's big money here in this kind of extortion, that that's where the attention is being focused on. Oh, wait, maybe we can develop this new idea into some non-high-bandwidth attack and succeed with our extortion where we're no longer able to otherwise.

Leo: Well, anyway, there you go. It's just one more way they can hit us. And that's the point, really. I mean, it's not like one is to the exclusion of the other. They're all good if you're a hacker.

Steve: Bad.

Leo: Or all bad if you're not. Brian in Wichita, Kansas heard what Steve said, but not what he meant. Listen up. He says: Steve, in Episode 164 you mentioned that routers have open TCP ports. You said they typically accept Border Gateway Protocol connections. Not really being advanced in my knowledge of routers, this is something I'd never heard of. Actually I'm going to confess I didn't know either that just a regular cheap router could accept BGP connections. Does this mean that my router I use at home has open ports? I don't use port forwarding for anything, and my router is a two-wire gateway, which doesn't support the unsafe Universal Plug and Play you've mentioned in a few episodes. So I'm safe as far as that's concerned. I assumed it doesn't have any open ports since I haven't opened any ports. Are those open ports you mentioned isolated from the firewall? In other words, do they

only allow connections to the router itself? Is my network still safe? I would hate to think the router's firewall has holes in it. What is this BGP stuff?

Steve: And I've got to apologize to Brian - and you, Leo, and any other listeners I confused by saying this. I didn't make myself clear. I meant formal, out-there-on-the-Internet routers.

Leo: Not our routers. Not our cheesy little Linksys.

Steve: Not the toys that we have.

Leo: Well, to be honest it's my fault, and I will take the blame for it because I heard that, and I should have said something. But I understood you to mean the big Cisco routers that are running the Internet. I understood that. But I should have said something to make that clear.

Steve: Right. BGP is port 179.

Leo: There's no reason for a Linksys to take BGP commands.

Steve: No. What BGP is, just to give a little information to follow up this clarification, it's what routers use - it's called Border Gateway Protocol. It's what routers use in order to share with their peers, that is, the routers they're connected to, the networks that they're able to route to. So essentially a router knows that one of its interfaces is connected to this set of networks, and a different interface connected to this set of networks, and a different interface to this set of networks. Well, it needs to share that information with the routers that it's connected to. So it uses BGP, Border Gateway Protocol, essentially to share its routing tables with its peers so that, when the peer receives a packet, it's able to figure out which one of its interfaces to send that packet out of as it hops to the next router. Well, that's, of course, determined by which networks that next router is able to send the packet onward towards its destination.

So it's very elegant. I mean, it's the height of elegance, actually, the fact that this all works as well as it does. But it does create a potential vulnerability on port 179 for formal Internet routers. Not, again, the little consumer boxes of plastic that we use at home. They've never, I mean, you could certainly load real routing software into them, many of them that run Linux, for example, and get support for BGP. But there's no purpose for it. I mean, it wouldn't do anything for your little router. But it is the core protocol that is how routers use - or how, sorry, routers talk to each other in order to share the connectivity that they all have with the routers that are adjacent.

And there have been many attacks in the past against port 179, the BGP port of commercial routers. There is, for example, there were some TCP spoofing attacks where people assumed that routers had existing BGP connections and were able to splice themselves into the BGP connection in order to, like, in order to spoof the routing tables and redirect the router traffic to other locations. So there have been some interesting exploits in the past. Those sorts of attacks have been hardened against. And something

like ShieldsUP! would immediately tell you if your BGP port of your own router at home were open and accepting connections. It would be port 179, and ShieldsUP! would tell you, oh, that's BGP port 179. So the fact that - I don't think I've ever seen anyone say wait, why does my router have BGP port open?

Leo: And if it were open, it still wouldn't know what to do with anything that came in at that port.

Steve: Yeah, no.

Leo: I mean, it'd just go, huh? You want to do what? No, can't do that.

Steve: So I did not mean little routers, I meant big iron routers.

Leo: It's not your fault, it's just there's a confusion between the two terms, frankly. They both are called "routers." You know, a little thing in the back of my head went off when you said it. The problem is, Steve, you're such an authority that I don't - I think all of us go, oh, okay, whatever you say. Whatever, yes, sir. I don't question you. And next time I'll make sure to get that clarified.

Kyle in Des Moines disagrees with the corporate IT administrators. Oh, who doesn't? Hi, Steve. Before I ask my Security Now! question, I'm happy to tell you I've listened to every Security Now! episode since Episode 1. See, he knew about the Strider HoneyMonkey. And I've enjoyed every episode. I listen to dozens of podcasts, including many TWiT netcasts. And while podcasts have come and gone from my subscription lists over the years, Security Now! has always been among my favorites. That's nice. Thank you, Kyle. I must also tell you that I'm a SpinRite owner and thankfully have never needed it in an emergency situation. Should that situation arise, however, I will be ready.

Steve: Yes, and please don't forget that you have it, Kyle.

Leo: Yes, exactly. You know, it's funny because you do kind of forget. And the other day we had a hard drive issue, and Colleen came back later and said, oh, I forgot, I should have used SpinRite. Oh, yeah. We got that. We got that somewhere. Now for his question, Kyle asks, he says: I'm an IT professional, a programmer. And I've found to be myself in disagreement with our office LAN administrator over a security issue. I bet you get a lot of these questions, you know, like answer a bet for me, Steve, who's right? Our LAN administrator believes it's possible for our public web server to be infected by spyware that is installed on the computers of customers who visit our site. I disagree. While I can see that a security hole in the web server or OS server software could possibly allow that machine to be compromised in a rare circumstance, I see no other way for a spyware-infected machine to spread its malware to a web server simply by that machine making an HTTP request to the server. I should also mention that a firewall only allows HTTP and HTTPS traffic to this web server from the outside world. The machine is hardened in all the usual ways. I just don't see any way for spyware to happen in the circumstance. Would

you care to settle our disagreement? Thanks to you and Leo, and keep up the good work. Can it happen?

Steve: Well, I agree with his assessment. There have been, once upon a time, exploits which only required an HTTP request to implement them. So it's not inconceivable that, if there were bugs in the server - IIS had some horrible ones in the early days where you could basically take the machine over just by giving a malformed URL. But he specifically says, notwithstanding a security hole in the web server or OS, is it possible that an HTTP request could do that. And so absolutely not. The HTTP request is simply a request for a page.

Now, it gets more complex because, as we know, the original spec for HTTP was the notion that web pages were static, and you were simply requesting static web pages from some remote server. That was extended with the "get" and "post" verbs, HTTP verbs, to allow you, in the guise of a request, to actually be sending information back, and that enabled the whole Web 2.0 concept that users could add comments to blogs and so forth. So you were accepting their data. But again, the idea being certainly that you're not going to allow anything malicious to be done by someone who's fundamentally anonymous and outside of your system. So I would say we never want to say "never" in security.

So to Kyle I would say it's certainly the case that, if everything is working right, that no one making HTTP or HTTPS requests would be able to harm a server. And I've never heard of any spyware that makes its goal that of trying to do so because there isn't a clear way to do that. At the same time we've seen all kinds of exploitable web-based attacks. There are lots of web-based attacks. But again, it requires some sort of vulnerability in the software or server that is receiving those requests.

Leo: Yeah, I mean, you're being generous. You're being complete. It's very, very unlikely. Like not going to happen.

Steve: Especially not random spyware infections on visitors' machines.

Leo: Yeah. Well, and most visitors are running Windows. And at least my - all my servers are running UNIX, or Linux, or some form of Linux. So that's even more unlikely; right? You'd have to write a Windows spyware program that would be HTTP aware, would be Linux aware, and then I guess it'd have to be an exploit. I mean, there's no generic way that this could happen. It'd have to be a flaw.

Steve: Well, there's certainly no invitation to allow people to accept code that you run on your machine.

Leo: The thing to understand is that you are not - the web server doesn't look at your disk, and you're in a very constrained way looking at the web server's disk. It's not like you guys have write access to each other's disks.

Steve: Right.

Leo: Eric in Los Angeles has seen a change in malware behavior. Oh, it's changing all the time, I'll tell you. Sometimes, he says, we want to test using ShieldsUP!, even when the connection is proxied. I work for a large company, and I'm trying to check the security of a new proxy server; but ShieldsUP! detects the proxying, refuses to proceed because the connection is proxied. That's just how you do it. I deliberately installed an SSL proxy because more malicious sites are using SSL to avoid the antimalware filters. Maybe you could adjust ShieldsUP! and allow for that, SSL proxies, and mention it on the show? What do you say?

Steve: Well, I don't know how to do that. That is, without causing ShieldsUP! to do the wrong thing. Many ISPs are using proxies, so-called "transparent proxies," where their users are actually having their requests intercepted by the proxy, and then that proxy reissues the request to the web server and then receives the result, caches it, and then sends the response back to the user. The advantage of that to the ISP is that, if a different user visits the same site, many of the site's components, all the little widgets and pictures and UI fluff that pages have that take a lot of time to fetch remotely, they'll be available on the local proxy because somebody else went to the site and pulled the same collection of debris, visual stuff, through the proxy. It kept them in the meantime.

So the advantage is the ISP is able to improve the performance, the apparent web performance, for that second customer and any others who use those before they expire from the proxy's cache; and the ISP is able to limit and reduce the bandwidth that it needs to pay for out onto the Internet, basically keeping a lot of that traffic in its own network, for which is does not have to pay because it's just using its own infrastructure. So there's a tremendous benefit and, that is to say, economic motivation for ISPs to use proxying.

So if one of that ISP's customers wants to use ShieldsUP! to test their computer, I need to, and I do, detect that the proxy is an intermediary. Well, in seeing that it's an intermediary, I'm unable to affirmatively get that customer's actual IP address. It could be a public IP address that is publicly routable, so that if I knew it, I could send probe - ShieldsUP! probe packets to it to check that machine for open ports. But I cannot detect reliably what the connection is.

Now, he says he's installed an SSL proxy. And I've been very careful in the design of ShieldsUP! even to detect that, so that I know that, even though it's an SSL connection, it still is being proxied, and I just go no further. I bring up an intercept page that says, hey, I'm very sorry, there's a proxy in between you and us. I cannot reliably determine your IP, and I just refuse to go any further. I mean, I couldn't, for example, allow them to type in their IP because then bad guys could immediately use that to go test - to have me test the security of random machines on the Internet. Well, go check the NSA. No thank you. I don't want GRC to be probing the NSA. So there's just - there's no way around that. There's nothing I can do that is safe. And so I just simply stop.

Leo: Yeah. I think that's completely reasonable. You can't do everything. You can't please everybody. All right. Get ready. Get a cup of coffee. This is a long one. Another quad venti latte, whatever it is you drink. Ben Isenhour of Lexington, Kentucky, fed them a knuckle. He fed them a knuckle, whatever that means. Let's find out. Hey, guys. I want to mention a few thoughts on security questions. As a business intelligence professional I always cringe when people solicit and record my personal information. BI, Business Intelligence, is a very interesting discipline which

collates information and allows businesses to make intelligent choices about - predictions about their future business and so forth. But it probably does make you a little sensitive to the idea of getting your information harvested.

My family thought I was paranoid when at Disney World I refused to let the ticket-taking machine take my fingerprint. What? I don't blame him. Geez. This was supposed to control passing a multiday ticket to another individual. So instead I put the crease of my knuckle on the machine. It worked. I have no idea what type of security practices they use in their corporate IT department, and I don't want a future biometric identifier escaping to the highest bidder. You know, he's absolutely right. I think it's unconscionable.

Steve: Yeah, here, take this, Pluto. You know?

Leo: It's none of their business. Disney taking my fingerprints? Give me a break. And I feel the same way about security questions. It seems like this is the new popular thing to do because I keep getting grilled by my online accounts to provide answers to very personal questions. I tell them a lie. I just say lie. As long as you can remember the lie. Sometimes up to as many as eight questions. This is just...

Steve: Where was I born? Gosh, I can't remember. No.

Leo: I just lie. You've just got to write down your lies, otherwise you forget them. This is just another way for more of my personal sensitive information to get out of my control. I don't think anybody really cares the name of your first pet, but okay. I really feel they're missing the boat on security, and in fact making us all less secure as time goes on. We know that. Those personal security questions are worthless, as we know from - you know they've now indicted the son of the Tennessee legislator that cracked Sarah Palin's Yahoo! email using security questions.

Steve: Yeah.

Leo: Is this something that Sarbanes-Oxley is suggesting they do? Or is it just their feeble attempt at multifactor authentication? YubiKeys for everyone. That's Ben saying that, but I agree. One last comment. Here's another interesting piece of info on Wells Fargo. The username for my Wells Fargo online account was initially set by them to my Social Security number. You know, my Fidelity account does that. But now recently Fidelity's given me strong encouragement to change it. Thank you.

Steve: Yeah.

Leo: I called in, was instructed to use an online screen to change this. I changed it and was feeling better. Then one day months later I was in automatic mode, went to the splash page, typed in my Social Security number and password by accident instead of this username. Well, imagine my surprise, it worked. It worked. Are old

usernames kept active? If so, as so it seems, for how long? Hmm, maybe you guys could get someone from Wells to come onto the show and grill them. Thanks, guys. Keep up the good work. What a great question. Or actually statements, really.

Steve: Yeah. Ben had a lot of interesting stuff to share. And of course - and our listeners, I'm sure, agree with the sentiment of, you know, the idea, again, that it's a matter of risk-reward. You know, our fingerprints are valuable. I mean, they're valuable identity information. And the idea of giving them to Mickey Mouse, or a Mickey Mouse ticket-taking machine at Disney World, because they want to prevent you from passing a multiday ticket around, I mean, there's no way the...

Leo: Sorry.

Steve: Yeah, that a machine is going to get one of our listeners' fingerprints.

Leo: And what he's worried about is now there's, in a database at Disney, a record that combines his name and other information, maybe credit card information, with his fingerprint.

Steve: And somebody will sell it in a camera on eBay, along with his Social Security number.

Leo: Maybe Disney will sell it. I mean, look at what the privacy - did they give you a privacy statement when they asked for it? I bet you they didn't. Unbelievable.

Steve: Yeah, yeah. And so anyway...

Leo: Very cavalier.

Steve: ...there's a tip, folks, use your knuckle. When it's not a situation you actually want authentication, you don't want to provide any useful information. You're somewhere where it's like, okay, I don't want to give this. I do not...

Leo: So he just flipped his finger over.

Steve: I don't want to ever authenticate again. Yeah, I think he just gave it the back of his finger instead of the meat of his finger. Which...

Leo: It worked.

Steve: Lot of sense, yeah. They're just not going to know.

Leo: [In Mickey's voice] We got it. Ho-ho. Have fun! Brian Hoort at Michigan State University wanted a bit about OQO. That's that new little teeny-weeny - did you get an OQO? You did, didn't you.

Steve: Yup.

Leo: Yeah. Steve, during the Google Chrome episode you briefly mentioned you have an OQO PC. I've been interested in them for years, ever since they were announced prior to release - yeah, they announced it and, like, three years later released it - but I've never had the opportunity to see one. Would you comment a bit on yours? I'd love to hear your review. P.S.: Security Now! is my favorite programming across all media - radio, television, film. I listen to it while running and doing yard and housework on the weekends. I love the show. Please don't change a thing. P.S.S. - I think it's P.P.S. But anyway, SpinRite is fantastic and worth every penny. We use it often here at the office, where it has become one of my must-have troubleshooting tools. All right. It saves drives more often than it should have to.

By the way, I've been wondering, what is so wrong with the operating systems that they corrupt their file systems? Is it indeed buggy OS code that's responsible? Is it a failure in the drive firmware? I haven't noticed any correlation between OS lockups and the need for SpinRite. The drives just seem to fail randomly. That's actually a very good question.

Steve: Yeah. I think the systems are - it's like chaos, I mean, formally math chaos, where you're unable to predict the weather because there are just so many variables all interacting at once. I mean, there's temperature. There's power supply power; there's vibration; there's OS activity; I mean, just so many variables. And there just isn't enough margin built into the reliability of all of those things. The computers are not so cold that they can afford to warm up. The power supplies aren't so extra powered that they can afford not to have enough. The drives don't have so much more tolerance in them that they're going to be reliable just themselves, and on and on and on. And so when everything is sort of near the edge, it doesn't take much to have the stars align, so to speak, and cause there to be a problem. I'm in love with my little OQO, Leo.

Leo: Really.

Steve: I really am. The only complaint I have is that the battery life is about three hours.

Leo: Well, that's not bad.

Steve: Well, no, it's not bad. It's like many laptops. And in fact it's a little bit more, probably about maybe, yeah, about three hours. The reason I have it is that I wanted true connectivity. My goal was that I wanted to be able, wherever I was, if something comes up, I wanted to be able to answer the question, check Wikipedia, or check for something. I mean, just I'm - when I'm sitting here at home, I've become so used to the persistent, wired-into-the-network mode. And when I'm in a restaurant or I'm at

Starbucks or I'm anywhere, it's like, hmm, gee, I wish I had a 'Net connection. And I've got my little Treo, that kind of has a web browser but not really. And I didn't go with an iPhone because they were on the wrong network. By the way, did you see that Verizon now has a very, I mean, a touchscreen BlackBerry, the...

Leo: It's the Storm.

Steve: ...Storm, which people are saying apparently has a tactile touch where it taps back on your finger when it recognizes a touch.

Leo: They call it a "haptic response," yeah.

Steve: Right, a haptic keyboard that really - that makes a big difference there. So...

Leo: You know, nobody's - it hasn't come out yet. But believe me, I'll get one the minute it comes out. I'm very intrigued by it.

Steve: Well, and it's on my network; it's on Verizon. I specifically left Cingular back when they were Cingular and before AT&T acquired them because Edge just doesn't do it. You need EVDO. It's dramatically higher speed. And they were talking about UMTS that was going to be the Cingular Edge-related technology that was going to give us the same kind of performance. But it was already on Verizon and still not available here. So I'm glad I made the move to Verizon. So anyway, but I just wanted a real Windows machine. So this is a - it's 800x480 resolution. It has a Wacom active tablet as part of the screen. Although I've got to say in my experience I have a couple tablet machines with much larger screens. There you really do need real estate for, like, scrawling and writing and taking notes and things. So it's really not very useful to have, in my experience, that Wacom. Although I think they've done it more for vertical, like vertical OEM applications is why it's there, for, like, specific custom purposes.

But I really do like it. It has answered the need I have felt for quite a while of just having it always with me. It's small enough that literally it can be in my pocket. And it lets me be on the 'Net because it has a Verizon wireless WAN built in, a little antenna you can pull out if you need greater signal strength. And so wherever I am, I don't even need WiFi, I'm able to get on the 'Net at completely acceptable speeds. And I'm using Firefox. Firefox has a nice feature. I'm, by the way, Leo, falling in love with Firefox.

Leo: Oh, good.

Steve: It's all I use now. I absolutely - everything I wanted to do, it turns out it has a way. One of the things is that the fonts on many web pages are a little small for this small screen. Firefox, under advanced font-rendering features, you're able to say minimum font size. And I set it to, like, 16. And now I don't have to constantly be scaling up all the screens. They're just the right size for it. I mean, it's just - so anyway, it's a perfect solution.

I did want to mention that the OQO Model 1 was pretty much horrible. It was a

Transmeta chip. It had a horrible sort of like flat plastic keyboard. The Model 2, which is now what they're selling, is far nicer, a real BlackBerry or Treo-style click-click-click keyboard. And anyway, it's just a tremendous little machine, if it works. I mean, it's expensive. It's pricey. But it finally answered my need for literally always having a connected Windows, real Windows machine with me.

Leo: And you just carry it with you all the time in your pocket? Or where do you put it?

Steve: No, I've got a little - actually I found a Franklin Day Planner, sort of a big one, that actually had handles. And I hacksawed out the rings inside to create sort of just an empty binder area. And so I have a whole little kit of stuff. I've got my Bose QuietComfort 3 headphones and an iPod with background music and the little Windows machine. And sort of it's like a little tiny little briefcase that I do - oh, and the Kindle fits in a side pocket perfectly.

Leo: But Steve, you don't travel. You don't go anywhere.

Steve: I go to Starbucks. And I do that...

Leo: It's a Starbucks kit.

Steve: I got to Starbucks, and I eat out. So I'm reading my Kindle. And if something comes up, I can order a book or check the 'Net or...

Leo: You're ready.

Steve: Yeah, I'm ready.

Leo: I love it. Jimmy Retzlaff in San Jose, California wonders whether banks will ever learn. I have a credit card from JPMorgan Chase, and they just sent me an email trying to sell me services. As far as I could tell it's actually genuine. The links - oh, man, an email.

Steve: I know.

Leo: The links all go to Chase domains. The headers are not suspicious. It has my full name and the last four digits of my credit card, et cetera. I've asked them not to send me such emails, but that's another story. At the bottom of the email is a section entitled "Email Security Information" with the following text:

"If you are concerned about the authenticity of this message, please click here." Isn't that great?

Steve: Oh. Yeah.

Leo: "Or call the phone number on the back of your credit card and reference the Chase Library Code: ALLSTANDINSR1001. If you would like to learn more about e-mail security or want to report a suspicious e-mail, click here." Huh? Well, I'm concerned about email security. And because of that, the last thing I want to do is click a link in an email that seems to be from a bank. If I were putting together a phishing scam, I can't imagine better text to put in the bait note. You are right on, Jimmy.

Steve: Yeah. I mean, it's just nuts. It's like...

Leo: Unbelievable.

Steve: Yeah. And as he says, when will banks ever learn? As our listeners know, the only acceptable thing to say, if you are concerned about the authenticity of this message, call the phone number on the back of your credit card. Period. I mean, you cannot click on a link in a message whose authenticity you are questioning.

Leo: I guess if they said "Type www.jpmorganchase.com into your browser," that'd be okay. But click a link, give me a break.

Steve: Yeah. Please press the back of your knuckle onto your credit card. Oh, just nuts.

Leo: What a world, what a world. All right, Steverino. The Troubling Question, or actually News, of the Week. This is from James Hudson, another British listener. We have a lot of listeners in the UK. I think that's fantastic.

Steve: Oh, maybe we have a lot of people who are writing to us from the UK.

Leo: Well, they're more literate there, you know.

Steve: Oh, careful.

Leo: Oh, boy.

Steve: You're going to get in trouble.

Leo: Hi, Steve. I opened Outlook 2007 for the first time in a few days today. Since I use Gmail, I find myself going to the website more often than not because I don't particularly like Outlook's IMAP handling. And I found there were a whole load of

undeliverable emails in my box. After searching around, there were a load of items relating to them in my Sent Mail folder with spam-type subjects. Uh-oh. Looking at the headers, I found they had actually been sent from my PC. So I did the obvious thing and did a virus and a spyware scan, which revealed a couple of tracking cookies in IE and Opera, but not my default browser, Firefox. Nothing that'd be sending spam. Looking at the sent items again, I noticed each one began with "Not Read." Now, doing a Google search revealed that they were Outlook read receipts that get sent when you delete an email without reading it.

Steve: Yeah.

Leo: I hate those read receipts. I hate those. It's one of the reasons I don't use Outlook. Further research reveals that Outlook has a bug in it that means if an email requesting a read receipt is deleted without being read, then it will send Not Read receipts automatically, no matter what setting you're actually using. Now, this is good to know because I turn off the email receipts feature in Outlook. I don't want to be doing this.

Steve: Set it to never. Never, you know...

Leo: Set it to never. But here's the point. In my case this is using Outlook 2007 with IMAP, which Microsoft has admitted is a bug on the MSDN forums, the Development Network forums, last year. He gives us a couple of links. But there are some reports of it happening with Outlook 2003 and POP3, as well. I'm not sure how much of an issue this is, as it only sends it out if you don't read the email asking for a read receipt. But it does generate a lot of extra and completely unnecessary traffic. So is a spammer using this facility to...

Steve: Well, it verifies your existence. I mean, that's why it's so troubling is that many of us can look in our inbox and easily identify spam. And we know that it's dangerous to open them. We don't want to preview them. So as we know, the safest thing we can do is just delete them, and that's what we tell people. Delete them without reading them. Delete them without opening. If it's clear to you that this is not something from your mom, or something completely random. Well, now we learn that even if you've configured Outlook to never send acknowledgement of read receipts, a bug in Outlook causes it to do exactly that. It sends, well, if the email had a valid From address.

Now, the reason he found these in his Unsent folder is that they were bogus addresses, so his Outlook was unable to send them to where they were trying to go. But that just - all that did was reveal the bug. You could certainly imagine spammers sending from some sort of a, you know, spoofing the source as being some sort of a catching server somewhere. And it would catch any email that someone had deleted without reading and, in the process, confirm that they were real, that here was a valid, useful address. I mean, this is a serious privacy concern.

Leo: So what happened was a spammer sent him email with a read receipt. Even though he had said never, this bug that crops up in IMAP and POP3 in some cases

sent the read receipt. The spammer - now, the problem is, don't spammers - they don't use real return addresses. So maybe this isn't...

Steve: Agreed. It's not the end of the world. But it's certainly the case that I just - I wanted our listeners to know that any mail they get which requests a read receipt, which they delete without opening, will result in Outlook sending a Not Read response, even though you've configured it not to do so.

Leo: I have to say I just hate it because, I mean, for security reasons; but I also hate it because people will then send me a note saying, "You deleted my message."

Steve: Well, exactly. I mean, it could be someone you know who's annoying. And it's like, oh, not this guy.

Leo: Well, it's not even that. I have three or four - I use IMAP. I have three or four systems. If I delete it on one system, maybe I already read it and responded, whatever.

Steve: Ah, very good point. So, right, you would have duplicate copies pulled onto different clients. And, oh, I already read that somewhere else.

Leo: Yeah. So I've actually had that happen from fans who say you didn't read my message? No, I read it. I just deleted a copy of it. [Sighing] Steve, a great 12 questions. If people want to ask you a question, how do they go about doing that?

Steve: GRC.com/feedback.

Leo: Okay.

Steve: And I read as many as I can. I love hearing from our listeners. We get so much good - I get ideas for shows. I get news. People are writing in - as I expected, by the way, many people had picked up on the Sockstress issue on Wednesday and Thursday before - after we'd recorded it, knowing that this was going to be a big deal. And so our listeners are on top of things. And I really appreciate them letting me know what they see going on so that we have a chance to make sure everyone knows. So GRC.com/feedback.

Leo: And of course when you're at GRC.com don't forget the great SpinRite. You can get it there. And it is the disk recovery and maintenance utility, a must-have for everybody who has - if you've got a hard drive, you need SpinRite. Just that's all there is to it. Also some great free stuff, like we mentioned ShieldsUP!. I love Wizmo. It's a little tool that's just really fun. All of that for free at GRC.com. 16KB versions of the show, too, for your friends who are on dialup still. Transcripts, a lot

of people like to read along as they listen. It's all there. GRC.com. Steve, we will talk again next week.

Steve: See you in a week, Leo.

Leo: Bye bye.

Steve: Bye.