Listener Feedback Q&A #16

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-080.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-080-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 80 for February 22, 2007: Your questions, Steve's answers, #16.

Security Now! is brought to you by Astaro, makers of the Astaro Security Gateway, on the web at www.astaro.com.

It's kind of a landmark Security Now!. I'm Leo Laporte. Steve Gibson is joining us from his fortress of security in beautiful Irvine, California. Hello, Steve.

Steve Gibson: Hello, Leo, great to be back with you.

Leo: Also joining us, the podcast dog. Can you hear him? Why is it he always starts barking right when the show begins? I've got a little terrier across the street from me, and he just drives me...

Steve: We'll probably have listeners starting to count the barks. Hey, I heard 319 barks.

Leo: Oh, are you kidding, it'll be a Wikipedia article, how many barks there were in this week's TWiT.

Steve: That's true.

Leo: We're going to do, because it's Episode 80, it's a mod-4 episode, we're going to do our questions and answers, the 16th of these. But before we do that I think we want to cover some stuff from previous episodes.

Steve: Got a whole bunch of interesting errata. First thing is, shortly after we closed down the recording last week of the really, really well-received episode on backtracking email spoofing – and in fact we've got a bunch of Q&A issues of really great questions that people had about some of that because there was a huge amount of interest in that. Anyway, shortly after we stopped recording, I was puttering around, I don't remember exactly why, but I looked at the headers of some of my email, and I realized with one of those classic kind of like "doh" events, that the header that I had been always referring to as "Received by" was just "Received:." And I was confused with something else. And it was like, oh my god, I said the...

Leo: Oh, no, I knew what you were talking about.

Steve: Well, not a single person – I thought I was just going to be held up and tarred and feathered. But no one mentioned that. Everyone was happy to have the content, so.

Leo: We weren't holding you to the literal words. It's the thing in the header that says who it was received by.

Steve: Yes, well, exactly. But as an Assembly language programmer, Leo, I've got to count my dotted I's.

Leo: You can't leave anything out, yeah. Well, it was funny because I should have said something because as we were doing the episode I was looking at my headers. I knew exactly what – I could have said – I didn't know that it mattered to you.

Steve: Nobody was confused.

Leo: Nobody was confused.

Steve: I was made self-conscious of it, so I immediately made a note in the errata for this week to confess my mistake.

Leo: There is no "by," it's just "Received."

Steve: Also, I subscribe to the SANS Security Institute newsletter. And I thought it was interesting that there's a new malicious JavaScript exploit which can alter DNS settings on routers with default passwords. And this basically relates to my continual drum-beating about scripting is bad, scripting is bad. I mean, wonderful and convenient and useful as it is, as I've said over and over, and nothing's going to fatigue me from saying this, is that you are accepting code from a remote server when your browser has scripting enabled. So I want to – there were a couple comments to flesh out this particular example of a new problem. SANS writes that malicious JavaScript placed on websites could be used to change DNS settings on home routers that are still using default passwords. Once the change has been made, the next

time the router is rebooted, the user would be redirected to spoofed, possibly malicious websites. Research indicates that about half of router owners have not changed the password from the default. Now, this is really interesting because, for example, I have not changed, I confess.

Leo: What?

Steve: My password from the default.

Leo: Steve Gibson.

Steve: I haven't.

Leo: I'm shocked.

Steve: Because I'm all closed to the outside. I deliberately said, you know, block WAN requests and kill off any WAN-side administration. So there's no way anyone can get at my router from the outside. But if I were not surfing with scripting disabled, as of course I am, what happens is, this new exploit, which is actually very clever, the script that runs in your browser goes out and tries to find your local router, then logs into it, all behind your back and without you knowing this is going on, and modifies its settings. So it's, you know, it's very clever. It's like a hack from the outside which you invited in because you've got scripting running on your browser.

Leo: Couple of things I would say there, though. First of all, this was a hack created by Symantec. It was a proof-of-concept hack. So I don't think it's in the wild.

Steve: Well, right. SANS has two editors who commented, and I thought that their comments were interesting, too. The first one says: "This is a clever attack exploiting the weak default configuration of most consumer home network products. Those consumer products tend to walk the user through a wide-open setup. And then at the end they say, if you want to turn security on, click here, but it might ruin performance or totally screw up your network."

Leo: And we don't want to support it, is what they're really saying.

Steve: Exactly. We really don't want you to, exactly. It says: "The Wi-Fi Alliance has started an effort called Wi-Fi Protected Setup. To make it easier for default configuration of Wi-Fi networks to be secure, the same thing should happen for the rest of consumer network products." Then the second editor appends to this: "Although this attack focuses on consumers, it illustrates a major trend I'm seeing," he says, "in the cases I'm investigating. Browser scripts are a scourge, with implications beyond the browsers in which they run. Recent browser script attacks that I've investigated," he says, "as well as the attack described in this article, involve a browser requesting content from a website where the attacker has posted a malicious script. The website shoots the script back in a web response that causes the browser to take some action on the infrastructure on which the browsing machine resides." That is to say, you know, inside of your network. He says: "The browser is used in effect as a remote-control sentinel inside the organization's firewall to manipulate its infrastructure, such as routers, internal applications, et cetera, controlled via browser scripts. We are seeing this in some advanced

attacks today, but watch for this vector to increase massively in the next year or so. To defend against it, you may want to disabled browser script support in your browsers associated with critical components of your network, or enable them only for important web servers that you've added to your trusted zone." So...

Leo: Well, you know, I'm not a big fan of this here. Because I have to say, first of all, all of the scripting, the Java scripting exploits we've seen are of this nature. They're not putting a trojan horse in your system, they're modifying your router and things like that. And second of all, I mean, turn off scripting and say goodbye to the 'Net. Almost every site, including our site, uses JavaScript. I mean, I tried for a long time, I used the Firefox NoScript extension. And I still do. But it causes reliability issues, and it's a big pain in the butt, and you end up approving it on most sites anyway.

Steve: Yeah, Leo, the problem is that scripting is not safe. I mean, it's just not. Yet it's so useful that it's difficult to fly without it.

Leo: I guess my question is how not safe. I mean, JavaScript is pretty sandboxed. It's not – it can't do – it can't erase a file on your hard drive, let's say.

Steve: Oh, sure it can. I mean, well, not by itself. But JavaScript, for example, on IE is able to invoke ActiveX controls, which then completely empower it to break through the sandbox and do anything it wants to.

Leo: I would submit the problem there is ActiveX, not JavaScript.

Steve: I wouldn't disagree.

Leo: I mean, that's a big problem. The other thing is, I'm surprised to hear you haven't changed the default password. We've been saying for a long time that there are a couple of basic things you should always do on a new router, one of which is change the default password, turn off Universal Plug and Play, and turn on WPA if you're on a...

Steve: Well, and in my defense, this is not a router on my network. It's on my cable modem, which is like a backup connection that I use sparingly, but...

Leo: [Indiscernible]; right?

Steve: That's true, yeah.

Leo: Yeah, I change all my passwords right away. And we do say that a lot. So I guess in this case there's a very clear fix that doesn't have anything to do with JavaScript. Your point is that JavaScript's dangerous.

Steve: Right, right. Okay. I was about to say "Moving on," as Kevin always says on Diggnation. I've been spending way too much time listening to him and Alex. We talked weeks ago about what muslix64 and the postings over on Doom9 about decrypting HD-DVDs. I'm only going to,

in this errata section, mention that those guys have continued to go crazy. There's an amazing thread over on the Doom9 site. We will have a link to it in today's show notes if anyone's curious in continuing to follow that. They're uncovering process keys and volume IDs and basically just really pulling HD-DVD encryption apart and developing a broad base of understanding all kinds of additionally clever hacks for ways to decrypt HD-DVD content.

Leo: Kind of reminds me of a feeding frenzy. It's like people are just dying to crack hi-def DVDs.

Steve: Probably just because it's there.

Leo: Yeah, well, it makes them mad that, you know, they came back with, quote, a better – that's the problem with any kind of copy protection. It just raises, riles up the masses. And they get upset, and they get – and they're going to crack it.

Steve: Also the other thing that we were swamped with this week were people asking about basically my reaction to and wanting me to talk about the Google white paper that they – I believe it was a USENIX – it was just recently released, in the last week, essentially. Google did a huge study on hard drive reliability based on the hundreds of thousands of hard drives that Google has in its massive Internet services and Internet indexing system. And it was really funny, I mean, I was going to talk about it from the beginning. I made a note in the errata a few days ago, I ought to mention this because one of the most shocking things about this was that hard drive reliability is much worse than even I thought. Remember a couple weeks ago I mentioned that I was surprised by a Seagate specification that talked about the reliability being 0.34 percent per year, meaning one in 300 drives. It turns out that, after a couple years, Google is seeing more like 8 percent failure per year.

But anyway, so many people wrote and said, Steve, you know, we know you're really supposed to only be talking about Internet security, or PC security. But the reliability and lack of failure of our hard drives is kind of important, too; right? Couldn't you talk about it? And I got so much email from people that – and since I obviously have a lot of knowledge about this stuff through my experience with SpinRite, next week we're going to talk about this Google paper. Again, we will post a link in this week's show notes to the PDF, so anyone who's interested can grab it and read it beforehand. Or you can just put something like "hard drive reliability" into Google's own search engine, and it'll find that link for you also. Anyway, so many people wanted us to talk about it that I thought, okay, I'm not going to turn into a SpinRite commercial, I'll just – we'll talk about what Google found because it's really interesting. And they did a beautiful study using this large sample base of hard drives that they've got.

Leo: Well, I'm all for making it a show. I mean, I think this is your area of expertise, why shouldn't we?

Steve: I also wanted to mention, and I picked this up from one of the other podcasts, Leo, where you, and maybe it was Amber, I'm not sure, mentioned that Google has now gone public. In other words, you no longer need someone else to invite you.

Leo: Free, yes.

Steve: Because I remember that one of our Q&As a couple sessions ago someone asked, hey, you know, I'd like to use Google and Gmail, but how do I get invited? So I wanted to just make

a note that it's no longer necessary to be invited. Anybody can get a Google Mail account and take advantage of Google Mail.

Leo: And well worth it, I might add.

Steve: I also noted that ShieldsUP crossed 48 million downloads. There was a big milestone for ShieldsUP.

Leo: We've got to party for 50 million.

Steve: Yeah, that's a serious number, yeah.

Leo: No kidding.

Steve: And finally, I've had a bunch of people that I've never mentioned before tell us that they really enjoyed the Peter Hamilton book that – there was "Pandora's Star," and then there was...

Leo: Well, the one you recommended to me was "Fallen Dragon."

Steve: "Pandora's Star"...

Leo: "Pandora's Star" and "Judas [Unchained]" is the...

Steve: "Judas [Unchained]," exactly. And in fact several people said they liked "Pandora's Star" so much they were now waiting for "Judas [Unchained]" to come out in paperback, and they couldn't...

Leo: I bought the hardcover.

Steve: I did, too.

Leo: I couldn't wait.

Steve: The thing I wanted to mention was, I just finished my third reading of "Fallen Dragon."

Leo: You know, I have to say, having now read "Pandora's Star" and halfway through "Judas [Unchained]," and "Fallen Dragon" is the best of the three. It is wonderful.

Steve: Yes. So I wanted to tell our listeners that if they liked "Pandora's Star," I mean, I read it already twice. It had been long enough now that I'd sort of forgotten the details. And I just

really enjoyed reading it again. So I will recommend "Fallen Dragon" by Peter F. Hamilton as a great, you know, it's not one of those things that you want to hurry through. All of his stuff is sort of literary in nature, sort of Dickensian, I guess you would say.

Leo: Yeah, it is, yeah.

Steve: You know, with lots of character development. But just a really good...

Leo: It's great, I mean, it's a full-fledged universe. And if you like sci-fi – I had never heard of Peter Hamilton till you told me about him. And I remember I was in Toronto, and I bought "Fallen Dragon" that day. And then I got it in the mail from you the next day.

Steve: That's right.

Leo: But you know it didn't go to waste. I lent it to a friend who's a sci-fi fan who loved it. We're turning people onto Peter Hamilton right and left. And then I did just – I'm in the middle "Judas [Unchained]." But, you know, it's funny because you'll get "Pandora's Star," and that in itself is 900 pages.

Steve: It's a read.

Leo: Yeah. And you figure, well, hey, at least it's going to end. It doesn't end. At the end of the 900 pages you've got one character going over a waterfall, you've got the aliens approaching, and for more, read "Judas...." It's like you're only halfway through.

Steve: Oh, it is, well, in fact, his bigger volume – that's not even the biggest one. I'm blanking on the name. His famous was a three-volume hardback that became six paperbacks. The "Night's Dawn" trilogy is the name of it.

Leo: Is that worth reading? Do you recommend that?

Steve: If it weren't so darn long. I mean, it's really interesting, some fantastic sci-fi ideas, real neat stuff. But boy, does it go on and on and on. So...

Leo: Well, I'll bring it to a desert island the next time I'm shipwrecked.

Steve: For people who have time it's, you know, like for non-TV watchers who really enjoy reading, I don't think there's anything better than Peter Hamilton.

Leo: I agree. Hey, you said we aren't going to do an ad for SpinRite. But yes we are. Let's do an ad for SpinRite, everybody's favorite hard drive maintenance utility and disk recovery. And if you're one of those 8 percent a year – that really brings it home, doesn't it.

Steve: It really does.

Leo: SpinRite is Steve's bread and butter. It isn't really an ad because we just do it because we love it. But I do want to make sure people know about it. You can read the testimonials at SpinRite.info, or buy a copy, as everyone should, at GRC.com.

Steve: While I was browsing through these questions in order to pull 12 together, which we're about to start on, I ran across a really neat one where the guy used the Security Now! web page to send his testimonial to me, rather than sending email to, you know, sales or support at GRC.com. And I'll just read the first line. He just said: "SpinRite saves a business." So I got a kick out of that, and I'll share that with our listeners at some future time.

Leo: Or just go to SpinRite.info. Let's get to our questions. We've got 12 good ones, as usual. It's Episode 80, divisible by four, so that means it's time to read a question for Steve. And you said that our episode on spam stimulated some email and some questions, and you're absolutely right. We got quite a few. Let's start with the first one from Chris Mckamie of Taft, California. He says: In episode 79 you talked about spoofing email headers. I wanted to know, how do I look at or find the headers in the email I get? We didn't mention that.

Steve: No, it was a great question from a practical standpoint. I'm a Eudora user. And it's funny because the button is labeled "Blah blah blah" on Eudora. As in yeah, yeah, yeah, you know, this is stuff you really don't need to see. And so you click the "Blah blah blah" button on Eudora, and it will show them to you.

But of course probably the majority of people are using Outlook or Outlook Express. And so I did a little – I just fired up one of my XP machines where, you know, you get Outlook whether you want it or not, thank you Microsoft. And if you right-click in the upper pane where it shows you your list of email, and go to the bottom of the pop-up menu, the context menu that comes up, there's a Properties item. Under Properties you can choose the Details tab. And what it shows you are the headers, and there's also a button there that says "Message Source." And if you click that, it opens another window with the entire source of the message, although probably just looking at the headers is enough. And so right-click on the item, go to Properties, and then go to Details, and that's where you'll find your headers in Outlook Express.

Most, I would say all email clients, without knowing that it's all, but it's probably all, do give users some way to find your headers. You might check under their Help or just browse through, you know, like the Edit menu or the Messages menu or whatever menu they've got to see if there's something about, you know, show all headers, or show header, something like that. It's probably discoverable by anyone who wants to track that down.

Leo: Oh, yeah, it's always there somewhere, absolutely. Joe McDaniel of Tallahassee, Florida, Go Seminoles – did you add that, or did Joe say it?

Steve: No, he did.

Leo: Right. They love their football in Tallahassee.

Steve: I don't know who's in Florida, Leo. I don't even know what sport that is.

Leo: Well, probably this time of year it's basketball. Anyway, he says: I was looking at your password page – that's GRC.com/passwords, or password, or password.html, or any variety thereof...

Steve: Yeah, you've got this nailed.

Leo: And a question occurred to me. I use it all the time, that's why, Steve. If a strong hashing algorithm is being used to hash the user's input down to a cryptographic key – okay, I'm going to have you explain that. But is there any advantage or disadvantage to using a password that contains the full printable ASCII character set, versus one that only contains hex characters? In other words, is a password containing the full ASCII set more robust than a password containing only the hex set? He wants to know if dollar signs and asterisks and all the other stuff...

Steve: Right. This was a great question because it gives people a way to think about – it's one of the ways I've always approached problems. So first of all, to back up a little bit, he's talking about how on GRC's passwords page you get this, you know, many different forms of strings. I have a full ASCII characters. I have one that's zero through nine to A through Zs. And then there's one which is only hex. And so presumably maybe things are able to digest the hex strings more easily or with less confusion. He's wondering would there be a reason to use the more complex, fuller character set strings.

The way I often solve problems is I, when I'm trying to get my handle on the answer to that sort of question, is I'll take it to one extreme or another. And so for example the idea being that, if you're using hex characters, then you only have four bits of data for every character you're using because you've got 0 through 9 and A through F, a total of 16 characters in the alphabet. And so you've got four bits of data per character. Well, if you use full ASCII, you've got seven bits for printable ASCII. So you've got many more bits per character that you're then stringing together. Well, so going to hex means fewer bits per character. Well, let's take that further and say that we only used 0 and 1 characters – that is, essentially, one bit per character – and we had passwords that are 64 characters long. Well, that means that we've essentially reduced the space of possible passwords to 2 to the 64th, which are then being hashed down to the cryptographic key.

So by taking it to an extreme, you can see that having a smaller alphabet does limit the number of possible passwords that you could be using, and therefore you could argue that it's less secure because, if somebody was going to brute-force it, they'd have a smaller world from which they were trying to brute-force your password. And somebody might, for example, glance at your password in the clear and notice that it's all ones and zeroes and go, oh, that's easier for me to crack than something that looks like your typewriter broke. Which is, you know, what the full character set looks like. So the answer is yes, having and using the full ASCII character set that gives you seven bits of randomness per character, so it's 7 times 64, that's going to end up giving you more security than using more limited character-set passwords.

Leo: I guess the question is does it make a difference. I mean, 2 to the 64th is still quite a few.

Steve: Yeah. It doesn't really matter.

Leo: Okay. Great. Thanks a lot. I mean, it's one thing to say binary. But, I mean, it's not like your only two choices. You do have a few more choices.

Steve: That's absolutely true.

Leo: He also says if it's doing hashing. Isn't that even processing it more and making it completely random? I mean, you can't do a brute-force on a hash, can you?

Steve: Well, no. In fact, what's happening is we're taking the large, you know, for example...

Leo: It's almost just a seed; right?

Steve: Well, WPA, for example, which is what most people are using these passwords for, WPA takes what you give it and hashes it down to a 256-bit key. So you want to give it more bits than it's going to be hashed down to, rather than fewer bits. So if you did give it 64 characters of 4 bits, well, now you're only giving it 256 bits, which are being hashed to 256. So basically that would be creating a mapping. There could be some overlap between the input and output. So it's better just to give the hashing algorithm as rich a source of entropy as you can.

Leo: Now, that makes sense. Brian "Da Hammer" – he probably prefers "DeHamer" – of San Diego has his thinking cap on. He says: I just listened to the Spoofed Spam Email episode, last episode, where you discussed how faked "Received" headers are used to try and mask the true identity of the mail's origin. You described how this sort of spoofing doesn't hold up under close scrutiny due to the fact that each SMTP server will use the IP address of the server it receives a message from. However, what if a spam bot were to initially add two "Received" headers to the outgoing message? The first header would make it look like the message originated from GRC.com; the second header would make it look like GRC.com had passed it on to some other SMTP server. The second spoofed header would of course include the actual IP address of the GRC.com SMTP server. Wouldn't this make it almost impossible to tell that the message did not originate from GRC.com? In other words, spoof the "Received" twice.

Steve: In fact, many, many people, to the great credit of our listeners, many people – and that's why I've included several questions that relate to this in today's Q&A – really thought this through and came up with some clever, wait a minute, how about this sort of ideas. And this is a perfect example.

Leo: I love it.

Steve: I was saying that, you know, that a single header would be spoofed. So they said, hey, what if you spoofed two? Because then, if someone was smart enough to look at the first one and go, well, that might not be right, they might not be smart enough to look at the second one. So the answer here would be to scrutinize the chain. Remember that we were talking about how email moves from SMTP server to SMTP server. And basically the "Received" headers – I'm not calling them "Received By" anymore – the "Received" headers would show you a chain which at some point is probably broken. And it's the break in the chain that would tell you, wait a minute, here's where this thing really came from.

Leo: So...

Steve: So he's right that adding a second spoofed header or more could make the task of figuring out where it came from more difficult. But if a discontinuity were found in this chain of received headers, that would be a key that something wasn't right somewhere.

Leo: Ah. So you'd still know something was wrong.

Steve: Probably yes.

Leo: But that's the point is that this is to fool people who aren't really that savvy.

Steve: Exactly. And it would be really good if people or servers really scrutinized the whole chain and said, wait a minute, this is fishy.

Leo: They don't, obviously.

Steve: Exactly.

Leo: Thom in Cortland, New York raised his antenna and asked: I have a question about Wi-Fi. Recently I took my university laptop home. Instead of a presentation, I started the system. Instead of a presentation.

Steve: Ahead of.

Leo: I'm sorry, I was misreading it. Ahead of a presentation. I started the system. The laptop is wireless. It automatically connected to an open hotspot in my building titled – oh, this is the Free Public Wi-Fi question. I like this one. I am certain – he got something, a hotspot saying "Free Public Wi-Fi." I'm certain there's no free public Wi-Fi and recognize this to be likely a scam hotspot. But I noticed it quickly; I shut the system down immediately. Even though I shut it down so quickly, and there's not any personal data on my machine, am I at any risk? Is it likely the hotspot was even able to do anything nefarious considering the quick shutdown? Do you know the answer to this, Steve? Because I do.

Steve: Go for it.

Leo: This is actually not a scam. When I first got this question on the radio, I said what you probably were planning on saying, which is it probably is a scam because there's no free public Wi-Fi. It's actually a bug in Windows. Did you know this?

Steve: No.

Leo: Yeah, it has to do with...

Steve: And I've also seen that, so I was wondering, isn't that a coincidence.

Leo: [Indiscernible]. Because if you've ever logged into an access point called "Free Public Wi-Fi," it has something to do with infrastructure Wi-Fi. I don't know the exact details, and I'll find the reference.

Steve: Ah, right. I know, where you're going machine to machine instead of machine to an access point.

Leo: And this is actually, in a way, spreading like a virus because one guy apparently did it; right? And then other people saw it and joined it. There's nothing there. You can't get any Internet access from it. So they forgot it. But it persists. And it's in your system.

Steve: It crept into your registry somewhere.

Leo: It shows up on other systems, and it has now spread across the land, and there are quite a few places where you will get online and see something called "Free Public Wi-Fi." It is a Windows machine in infrastructure mode that at one point logged onto another Windows machine in an infrastructure mode with "Free Public Wi-Fi" as one of the hotpots it had been to. I'll find the article because I did some research when this person asked me the question on the radio. And I answered it as you I'm sure were going to answer, which is, yeah, it's probably not a good idea to join such a thing.

Steve: Right. We can say a little bit more, although I think that's very cool news.

Leo: It's fascinating, yeah.

Steve: We can say a little bit more about this issue in general. That is, for example, if you were to find that your machine had automatically connected to a hotspot that you were suspicious of, first of all, answering this question, it's probably not the case that just the act of connecting could be a problem, so long as you've got – probably you're using Windows XP or maybe Vista. But one way or another you probably have a personal firewall on. So as we know, today's exploits generally are people going out through a firewall asking bad stuff to come back in, or in the case of a Wi-Fi, people, for example, doing nonencrypted email log-on where their username and password are going out in the clear. In both instances, the user is doing something with the computer, and somebody deliberately running a malicious Wi-Fi hotspot will be monitoring that. However, with the firewall up, just the act of connecting to even something malicious, as long as you recognize it and shut down, it's probably not – there's no opportunity for you to be infected in the normal case.

Leo: Yeah. And to add to this, it is a bad idea in general to log into ad hoc networks.

Steve: Right.

Leo: There's a difference, and you can tell, whether it's an infrastructure network or an ad hoc network. And an infrastructure one is with a Wi-Fi base station and, you know, like at the coffee shop and so forth.

Steve: The normal style.

Leo: Ad hoc is coming off of somebody's computer. So you really probably don't want to join a network on somebody's computer. That would just be very trusting. The reference I'm going to put on the website is from Dwight Silverman's tech blog in the Houston Chronicle. That's where I saw it first. It's a great story, and it just – it's a bug in Windows. Microsoft says they plan to fix it at some point in the next XP Service Pack. But who knows when that's going to be. Unknown whether it's in Vista.

Steve: Interesting.

Leo: Isn't that weird?

Steve: Yeah, that's cool.

Leo: Chris Detrick of Westerville, Ohio has been watching some TV. He writes: I was watching TV the other night. I happened to see the new Mac commercial – oh, I love that.

Steve: I do, too.

Leo: It makes fun of the Vista User Account Control. In the commercial the Windows machine says, I'm receiving a connection request, accept or deny. And it goes on and on. In fact, I'll find it while you answer the question. While I found the commercial quite funny, after thinking about it for a moment, the Mac does the same thing. I own both a PC and a Mac, and I've used both Vista and OS X quite a bit. They both require approval for system-wide changes, installing software and other similar actions. If anything, the UAC in Vista is more relaxed because it doesn't require a password. It only seems more aggressive because the desktop fades away, and the UAC window is running in the foreground. But I believe the effect is the same. I'm wondering what your opinion is on the subject. Is one better than the other? You know, Bill Gates was incensed about this commercial.

Steve: No kidding.

Leo: Yeah, he said they lie. How are they allowed to lie like this? Because he makes the same point. We're just copying what Apple's been doing for years.

Steve: Okay, now, you know better than I do, Leo, because you have a lot more Mac experience than I. But doesn't the Vista UAC, the User Account Control, seem more in your face, seem to actually require more yes, this is what I want to do acknowledgements than the Mac?

Leo: It does. And I think some of that ad comes from the early beta versions of Vista, when it would do it all the – it was so bad, it would take four clicks to get anything done. The actual shipping version of Vista is much less intrusive. I don't mind it. And again, as a Mac user, I'm used to some of this already. I don't think it's – if it's more, it's not much more. And it's not more intrusive. Now, if it's not asking you for a password, that means you're running as an administrator. And you might want to rethink that strategy. You might want to start using a limited user account on Vista.

Steve: On the other hand, if you do that, then you will be asked for a password every time one of those pops up.

Leo: Well, not every time, but every time there's a system change. There are certain rules about when it needs a password, when it needs to be an administrator to make the change. Frankly, that's what the Mac does, as well. And I think that's good. Let me play, as long as we're talking...

MAC: Hello, I'm a Mac.

OS: Mac has issued a salutation. Cancel or allow?

PC: Allow. And I'm a PC.

OS: You're returning Mac's salutation. Cancel or allow?

PC: Allow.

MAC: Okay, what gives?

OS: Mac is asking a question. Cancel or allow?

PC: Allow. He's part of Vista, my new operating system. PCs have a lot of security problems, so he asked me to authorize pretty much anything I do.

OS: You are pointing out Vista's flaws. Cancel or allow?

PC: Allow. I could turn him off, but then he wouldn't give me any warnings at all, and that would defeat the purpose, so...

OS: You are coming to a sad realization. Cancel or allow?

PC: Allow.

Leo: You've got to, I mean, it's a funny ad, but I think Chris is absolutely right. This is something Apple's done for quite some time. And frankly, I'd say it's the number one reason that Macs have been more secure than XPs.

Steve: Well, yes. And in all seriousness...

Leo: Kind of disingenuous of Apple to make this ad, frankly.

Steve: Right. In all seriousness, one of the things I have said, and we will do, there's much more technology going on in User Account Control than just that dialogue. So we will be doing a Security Now! episode about the Vista's User Account Control and really going on there because there's a lot of technology there. But, you know, this comes back to the fundamental security rule that we talk about often, which is there isn't a way to achieve security in the existing structure without requiring some tradeoff with user convenience. The browser scripting problem that you and I always tussle with is exactly that. As you know, you want scripting on. Turning it off and selectively enabling it is a pain in the butt. But it's the only way to browse without scripting on while still being able to go to some sites. And similarly, the problem with Windows is that malware has access to everything going on in your system. And of course this dialogue popping up is to give you some opportunity to intercept malware doing things behind your back, to make sure that this is something you're really asking to be done. And there's no way to do that without asking you in the current model, unfortunately.

Leo: Well, and I don't find it onerous. I don't find it a pain in the butt. It's a reminder to me that I'm doing something that's modifying the system. I think it should ask.

Steve: And we got a good commercial out of it, anyway.

Leo: It's hysterical. I mean, it's a very funny commercial. Tony Chen of Orlando, Florida wants to know how to best perform long-term data archiving. He writes: I work as a UNIX sysadmin, where security is only one aspect of my job. It's been noted many times in the podcast that, due to the high areal density – that's the number of bits per square centimeter – of today's hard drives, they're constantly throwing out errors that the drive controller or firmware is correcting on the fly. That's that ECC you talk about, Steve.

Steve: Right.

Leo: Do these errors only occur when the drive is powered on? In other words, if a drive is sitting unused, outside a computer in a box, for instance, is the drive protected from degradation – maybe, but not humiliation – assuming it's stored properly, without the firmware to constantly correct it? I have automated backups to hard drives – oh, he's worried about archiving things off-drive for fear they won't be recoverable on a dead DVD-R. At least with a hard drive I can use Disk Warrior on OS X or SpinRite to recover some, if not all, of my data. So I guess what he's saying is his method of backing up is to – and I know people who do this – back it up on a hard drive, put it in the shrink wrap, and stick it in the corner. Is that better than putting it on a DVD?

Steve: Okay. Well, this brings up a couple issues. There was some interesting dialogue on TWiT a few weeks ago that I caught that I thought was interesting because all the studies that I have read have said that recordable media is more reliable, for example, than pressed media. That is, that pressed media, for example traditional DVDs, you know, regular consumer video DVDs, have been shown to have a long-term oxidization problem with their surface, which due to the different process used in recordable either CD-R or DVD-R that those photo-based technologies do not have. And so properly stored, everything that I've seen says that recordable media has more reliability over the long-term than non-recordable media, just from a disk, you know, an optical disk standpoint.

Leo: We said it the other way around. I probably...

Steve: Exactly. Okay, so now in terms of long-term hard drive storage, the answer is, once upon a time there was a problem with something called "stiction." Stiction was a phenomenon where the head would essentially molecular bond to the hard disk surface because they were both so smooth. The head was so smooth and the disk platter was so smooth that, if you turned the drive off over a long period of time, you would end up with the head literally just sort of welding itself to the hard drive. And in fact, that's one of the reason that heads park on the inner area of the disk. They go to the inside spindle for two reasons. One is that if the disk is bounced by mistake, you can imagine the platter edges are going to be oscillating wildly out on their perimeter. But because they're anchored at the spindle in the middle, there's not much motion in there. So you'd much rather have your head on the inside near the spindle mounting, where there's not going to be much oscillation of the drive is bounced, than sitting way out on the outer edge flying back and forth and up and down.

Secondly, from a mechanical standpoint, the starting torque of the spindle motor needs to be much less to make sure that any stiction is broken to get the head loose of the platter as the drive starts up. You can also imagine from just a standard mechanics of physics standpoint, if the head is way out on the perimeter, and if that's where it's sitting, it's got a much greater mechanical advantage to keep the disk from starting up in order to overcome the friction between the head and the disk platter. Thus heads are now parked on the inside of the drive as a much better place for them to stay. I haven't heard of stiction being a problem for a long time. Back when SpinRite was young and people had this problem, the common wisdom was you powered your computer up, and then you bonked the hard drive with the handle of a screwdriver, literally to give it a little bit of a knock in order to shake the head loose from the platter. And many times that worked for people. But with current technology and new drive and head lubricants, that's just not a problem anymore.

Leo: But I still hit my hard drive with a screwdriver...

Steve: Just bonk it a bonk, bonk it with a screwdriver.

Leo: I used to love that.

Steve: Keep those bits from getting too set in place. It works.

Leo: It does work.

Steve: Back in the old days it did, yeah.

Leo: And it was the best thing to tell people, oh, you got stiction, hit it. And people go, what?

Steve: So this is one of the advantages of using, for example, a current technology removable Firewire or USB drive is you are able to plug it in, do some backups, and just stick that drive on the shelf. To answer his question, you know, in the reasonable term, like 10, 20 years, there's no reason that a drive would age in a way that would cause it to become less reliable when it's just sitting on the shelf doing nothing. You want to, you know, make sure you don't subject it

to temperature extremes for long-term storage. But the drive would be much more stable that way than if they were in constant use where, as Google has shown, and we will be talking about next week, reliability really does start to suffer after a couple years.

Leo: So sitting there, well, cosmic rays hit it and stuff. But it's not going to – you're not going to lose enough data to make it something to worry about. In fact, I think Alex Lindsay tells a story of somebody who, you know, hard drives have gotten so cheap, the video editor people who will just when the project's edited and done, they'll just take that drive, wrap it up, put it aside. That's the backup.

Eden Li of Beijing, China by way of Arizona has been paying attention. He asks: I am a long-time listener, first-time emailer. I was listening to your latest episode about determining the origins of spam by using the "Received" email headers. You mentioned that it was possible to determine that someone had faked a message from GRC.com by putting in the domain name for the first received header. However, because the IP address of the connecting machine was included, it was possible to prove your domain did not relay the message. The domain name doesn't match the number. Since the sending agent can add whatever headers it wants, can't it just fake a path with valid IP addresses from GRC to itself before sending it out? This is kind of similar to the earlier question. The path up to the malicious sender would look entirely valid since it controls the headers. In this case, it would appear that GRC had sent the message to the sender, and the sender was just doing its duty by relaying it to the next server. So kind of like a man-in-the-middle attack. You put GRC first, and then put yourself in there underneath.

Steve: Yes.

Leo: Is there any recourse against this? Am I missing something?

Steve: He is not missing anything.

Leo: Oh, dear.

Steve: And this one wins the prize for the cleverest hack around the way that we talked about last week of backtracking and determining a spoofed email. Now, one thing that the vulnerability, of course – so first of all, just to quickly summarize this, rather than having a header saying that this originated from GRC.com that sort of just is there, and then the second header in a row would be added by the SMTP server that the spammer connected to, so that server would be able to see the connection from the spammer, he's saying intelligently fake the very first spoofed email header so that it says GRC.com connected to their IP address so that the chain I was talking about when I answered a related question earlier, the chain of header would be unbroken, and it would literally look like GRC had used some random consumer's machine somewhere else as an email relay, and that it was then relaying it. And sure enough, that works. Of course the giveaway is, if you looked at the IP range or did a reverse DNS lookup on the IP, you would see, you know, some Comcast.net IP, or, you know, basically a consumer-looking IP rather than a regular server. So...

Leo: And you would never send to that [indiscernible] to relay your mail.

Steve: It is a clever hack for getting around this notion of an obviously spoofed "Received" header.

Leo: But I'll tell you it probably never happens because nobody looks that closely anyway. There's no reason for a spammer to worry about it.

Steve: Exactly.

Leo: Why go to the effort? I mean, if we all started looking at our headers, I guess. But we don't. And it doesn't matter. The spam still gets through.

Jeff P. of St. Louis, Missouri has learned to become more cautious on the 'Net. He says: Listening to Security Now!, I've also learned to question things when they don't seem right, which is what I'm doing now. I recently logged onto my work laptop at home, connected by VPN, to do some work. I opened an internal web location, http://something, that has always worked in this manner. To my surprise, I get a Charter Communications search page – that's his Internet service provider – that says "Page Not Found." I tried again, then another internal site, another, no luck. In the past week I noticed on my home machines that 404 errors had been redirected to Charter. I thought they were probably starting to proxy web traffic, or had been doing it but just now made it easy to see. However, when on the VPN I thought all traffic would be tunneled to the corporate server. Why is Charter seeing it? I still get that Charter-served error page. Even when he's on VPN. How is that possible? I did go to GRC with and without the VPN connected. It did report two different IP addresses. From what I know, I can't explain what is happening. I can't either. This is interesting. Can Charter proxy HTTP traffic when it's connected as VPN? Any ideas?

Steve: Well, we know that it cannot. That is to say that, if you're trying to go through a VPN, or your traffic is going through a VPN to your remote destination, that no one can proxy that connection in the way that it's happening. That is, you ought to be able to get a connection into your corporate network, and it would be an encrypted tunnel that is impenetrable. Listening to this, the two things I would ask Jeff to check would be his browser's proxy settings. There have been various attacks where proxying was turned on. So if his browser is not just using his raw Internet connection, but is for some reason running through a proxy, it might be possible, depending upon the way his VPN is configured, for him to still be using his raw, non-VPN traffic. And when he thinks he's going to his internal site, he's actually going to the external Internet; and then Charter is saying, wait a minute, we are now unable to find this, and giving him an error message. So the first thing I would look at is the proxy settings on his browser. Secondly, check your hosts file. This, you know, this happened in the last week or so, and it's new behavior. It's possible that something may have crept into the hosts file and is messing up the DNS resolution on his system, basically trying to intercept traffic and route it somewhere else.

Leo: There's one other thing I would mention. A number of – I've seen this on a number of VPN clients, and I'm actually just launching a couple now, just to see – allow you to route web traffic not through the VPN.

Steve: Deliberately. Yes.

Leo: Deliberately. So in fact you'll see a checkmark in the settings that says something like route all traffic through VPN, and that I would suspect as being unchecked, or it would say do not route web traffic through the VPN. Clearly it's not being routed through the VPN.

Steve: Right.

Leo: So but I think that it's not necessarily something nefarious. I think that that is a setting in some VPN clients.

Steve: Cool. Then certainly look for that, Jeff.

Leo: And that makes sense because, if you're surfing the 'Net, and you just want to surf the 'Net as opposed to using the corporate network for transferring files, you might not want to go through the overhead of going all the way to your office.

Steve: Well, and moreover they might, by policy, they might not want you to use that much bandwidth over their VPN server. They might want to reserve that for email and just, you know, internal corporate stuff.

Leo: Yeah, excellent point. Let me just – I'm just going to look in my – I have a Cisco VPN client running here. And I'm just, out of curiosity – I'm pretty sure I've seen this as a choice in a lot of...

Steve: I absolutely have. When you talk about that, the name on that setting, it seems very familiar. And I'll bet Jeff's going to go, oh, that's exactly what it was.

Leo: Of course. Now, he raises another interesting thing which I think we should talk about, and you're seeing this more and more. When you use a company's DNS servers, it's possible to configure them to handle a 404.

Steve: And in fact, of course, we know that, what was it, VeriSign that got famously in trouble because they ended up redirecting unknown domains to their own, you know, commercial advertising page, essentially. And, you know, the whole Internet community got very upset with that.

Leo: Yeah. That was wrong because they're just a registrar. They're not your service provider. The same thing happens with OpenDNS. If you use – in fact, I think that's how OpenDNS monetizes itself. If you use, instead of your ISP's DNS server, if you were to use OpenDNS, one of the advantages of it is OpenDNS does a very good job of trying to guess what you were trying to say when you type in an address so it can fix a mistyped address. If you typed .CMO instead of .COM, it'll fix it. But also if you type a 404 it'll give you actually a very useful search page. So it can be a positive. But that's important to know, that many ISPs are starting to do that. And they do...

Steve: And they have that power.

Leo: Yeah. And they do it because they put ads on there; right? That's what VeriSign was all about.

Ted Hosmann of Monterey, California has been thinking about email spoofing. Who hasn't? I have a question regarding the spam backtracking episode. You had mentioned that in some cases the received headers in email can be many, based on how they're routed to

the final destination. Is there a chance that more than one email header could be spoofed? Say 10 or 20 or 30 headers before you'll find the source of the zombie machine? Wouldn't it make sense to throw in a number of headers from which spam could be coming from, in a red herring-style attack?

Steve: Yeah, I mean, there's another creative approach is to literally bury yourself in a whole bunch of spam headers. Now, again, it's the last header in that group that is going to be the real connection point to the email as it moves to its first SMTP server because the destination SMTP server will add the top received header to the existing pile of them. But and so there's no way for a spammer to bury its received header among others. But it certainly could confuse the issue just by having a bunch of them. And, you know, just make some poor recipient's eyes cross when he thinks he's going to be tricky and look at the email headers and just sees this whole nest of them.

Leo: There's two morals to take away from this. One is that SMTP is hideously insecure. I mean, it does nothing to validate headers. Zero, right? It just, I mean, you could do anything you want. But I think also spammers don't even bother because, so what? So what if you figure out it's spoofed?

Steve: Well, yes. And in fact, we're probably guilty of being so much overboard on the technology relative to the value of the content of the junk that's being sent out. I mean, it's phenomenally amazing to me that spam still works. These spammers are making reportedly an amazing amount of money just by spewing all this junk out on the 'Net. It's like, who clicks on that stuff?

Leo: Well, in the last analysis, that might be the best way to fight spam is not to do it. Don't buy anything from them.

Steve: Unfortunately, it is so cheap to send this stuff, especially now that we've got spam bot networks.

Leo: Yeah, they're free. Well, you have to pay some hacker, but that's about it. Yeah, in fact, don't even – I think a lot of people say, oh, well, let me click – I can click the unsubscribe link; right? No.

Steve: All that does is validate that you exist.

Leo: Don't click anything. Todd Zervas of Riverside, California has crucial files he wants to protect. I'm very tempted, he says, to use one of these online backup sites to backup the files from the computers at my house. Making DVDs at regular intervals is a pain, and I have to move them offsite in case the house catches fire and melts all my pretty backup disks. Well, he's got the right idea. I think offsite backup is really important. He says: I do fear these online backup sites. How do I protect my data from them? If I back up all my computer's data to these online sites, can't the operators look at my files? I've got financial information on one of my PCs. My daughter is writing novels on another. Can I trust the encryption on these online backup sites? Can I add my own encryption and not give them the keys?

Steve: Isn't that a great question? I love the question.

Leo: You've got to trust some of these guys, though; right?

Steve: Well, I would add one more thing, and that is, if they have the ability to decrypt, then they could be legally compelled to do so. And I ought to tell you, I actually have the domains cryptkeepers.com because I've been interested in the idea of providing secure online backup services. Not like I have – believe me, Leo, it's not on the list of things I'm going to do. But I grabbed those domains a long time ago because I liked this idea. I don't know if it ever makes economic sense with hard drives becoming so cheap and with things like BitLocker and, of course, TrueCrypt being available. But it is an interesting alternative.

So to answer Todd's question, again, it's not so much that I would worry about malicious intent on the part of the backup sites. But we know, especially with everything that's been going on with the Patriot Act in the U.S. and concerns about what subpoenas could be issued, I would be very concerned about even a reputable company being subject to a subpoena and compelled to decrypt somebody's offsite backup under court order. So what you want is you want a technology where you are absolutely sure you're giving data to people, and there's no way they can use it, that the data is just an opaque blob.

And thinking about the answer to this question, I thought, hey, you know, TrueCrypt would be a perfect solution. You configure TrueCrypt in that mode where it uses a file as its container, and then you mount that as a drive. You see a drive from your perspective. On the outside, the operating system just sees a file. That one file is then what you keep synchronized with your offsite backup facility. Anytime something happened, you would be able to get that file back and apply your TrueCrypt passwords in order to decrypt it. Only you in the world would be able to make any sense of that file. And it's just a cool solution because basically you see a drive letter. Everything that you're storing there is encrypted in that container, always stored, as we know the way TrueCrypt operates, always stored in an encrypted fashion, never in the clear. And so that one file, anybody can have it, essentially.

And in fact, as we talked about a long time ago when we were talking about TrueCrypt, it's exactly what I do for my precious cargo. I burn TrueCrypt containers onto DVDs and send them to my attorney and to my mom. And I just know that that stuff is out of here, I can always get it if I need to. But if their home or office was burglarized, I've lost nothing because nobody else can make use of that data.

Leo: I've become a kind of fan of these offsite backups. In fact, Carbonite, which is one of them, a really good one – it's a sponsor of the radio show, so a disclaimer there. But the reason I like that is because it's in the background, it's doing it all the time, it's always backing up. And I'm backing up podcasts, I mean, I use it because I backup our raw tracks of the shows and the podcasts so that nothing gets lost. You know I've lost some podcasts, and it's been very embarrassing. So this has been a boon. I don't care if somebody, some government agency wants to hear the podcasts, so I'm not going to encrypt that. But it makes sense to do that for your financial data.

I am looking at Carbonite's privacy statement. And certainly before you use any of these services you should read their privacy statements and make sure they're verified by an independent third party. But one of the things they say is – they do encrypt, they encrypt with SSL, and they have all the security on the facility and stuff. But it says Carbonite will not share your encrypted files with any third party unless such action is necessary to comply with a government or court order legally compelling us to do so. So I'm sure every one of these services has a statement like that.

Steve: Right. And so from the consumer watch standpoint, you simply give them a pre-encrypted blob.

Leo: And there's nothing they can do.

Steve: And nobody can ever do anything with it.

Leo: They'll comply with the court order. They'll give the court order the unencrypted data, removing their encryption. But the court can't do anything. Now, the court can come to you. That's between you and the court.

Steve: Exactly.

Leo: Mark Frautschi in Rockville, Maryland is wondering, how do spammers get our email addresses in the first place? Could it be that our friends, families, and colleagues are unwittingly turning our addresses over to the spammers and malware writers?

Steve: You know, we had never really talked about this. And I know that a lot of people probably have a sense for how this happens. I have some experience myself running an SMTP server, which I'll share. But one of the things that makes me cringe, Leo, more than anything else is when someone I know who has my cherished and protected email address uses a website to send me something. You know, like, oh, your Valentine has sent you a Valentine's Day card. Oooohhh, Mom, don't do that. Because there they're deliberately giving your email address to some website that now has it, and lord knows what they're going to do with it from there.

Leo: Yeah. And it may be Evites, fine, I'm sure it is, or Egreetings. But who knows? And Mom didn't check. Mom didn't read that privacy policy. So it is true, that's probably one way.

Steve: And certainly websites are being scraped. The other thing we knew from the way GRC's email addresses used to be getting out was we would be getting lots of spam to our support and sales email accounts, even though those were scrape-protected on our website. What was happening was, people who had written to GRC Support or Sales, they would get some malware on their machine that was scanning their machine for other email addresses and combining them into a master database. So it's certainly the case that malware, one of the things malware has done is mail stuff out. Well, we know many viruses like this in the old days that you'd get your machine infected with, and they would send themselves to everybody in your address book. So that same kind of scenario is now being used by spammers.

The one thing I would add, as an operator of an SMTP server, I've watched the traffic on my server carefully. And it's phenomenal, I mean, doing packet captures, you will see remote SMTP spam servers connect to any other SMTP server, TWiT or GRC or Amazon or anybody, and just sit there guessing. I mean, literally alphabetical dictionary lists of names, they just go, you know, albert@grc.com, andrew@grc.com, anyone@grc.com. And they just go through the list. So the one thing I would say to anybody is, if you ever change your email address again, change it to something unguessable, not something that's probably going to be in someone's dictionary somewhere, because back when I was just steve@grc.com it was not pretty, Leo.

Leo: No. I get mail to leoshair@leoville.com. Oh, no, I get several thousand messages a day.

Steve: Leoshair is an account that you have? Or do you have a wildcard account, so anything...

Leo: Well, what happens is the mail service that filters my spam, I use a commercial company called MailRoute, knows what my canonical, my good addresses are, and only forwards me mail to those good addresses. But it gives me statistics on all the other addresses.

Steve: I see.

Leo: And besides Ann and Sally, I get some to Leoshair. My thinking is somebody who didn't like me added it to some mailing list somewhere, and I've been getting stuff. So maybe it was my mom. I don't know, Mark.

Jim Etherington of Ottawa, Canada writes – is this our last one? No, we're going to have...

Steve: This is #12, but we have a closing comment.

Leo: Closing comment. I've been a longtime fan of GRC from reading the interesting security articles and trying the freeware programs to using ShieldsUP. I've listened to every single Security Now! podcast from the very beginning, and I have bought and used SpinRite 6 to keep my hard drives working. This is the kind of guy we like. My question is, how good is port knocking as a security measure to selectively allow connections to closed ports in a router? Oh, port knocking, this is something I'm starting to hear about now a little bit. Thanks for doing a great job with the podcast, and look forward to hearing more. What is port knocking?

Steve: I am so high about port knocking. It is a very, very cool solution. It's had some bad press because people are attacking weak implementations of port knocking, saying that port knocking is bad. But it's just very clever. And in fact, I've been considering using it in some fashion for my own security because it can be really bulletproof. The idea is you've got a router with all ports closed. If you want, from the outside, say that you're a telecommuter, and you want access into your network, with port knocking you would send a pattern of packets, a sequence of packets to specific ports on your router.

Leo: A sequence known only to you.

Steve: Known only to you.

Leo: It's like da, da da da da da, da da [to the tune of "Shave and a Haircut"].

Steve: Exactly. And it could be as long and arbitrary as you want it to be. Of course, as we know, the longer the sequence, the less chance anyone's going to be able to guess it, blah blah blah. And with 65,535 possible ports, that is to say, essentially 16 bits per port, if you just did

four packets, now you're at 64 bits. And if you do eight packets, you're at 128 bits of equivalent security. The idea being that the router notices the sequence of packets that it's rejecting. Remember, you have your router running in full stealth mode so it's blocking these packets. But even though it's blocking them, it can still sense them. It's still receiving them. So you send a specific sequence of packets from your IP, your router or some software running behind it, maybe for example reading the router log in real-time, it sees a given IP has sent a matching sequence of packets, and it says, oh, this must be somebody who knows the secret knock. And so it then alters the router's configuration on the fly to allow inbound traffic from that IP. I mean, there's just nothing wrong with this, Leo. It's a tremendously clever idea.

Leo: I'll have to ask, I bet it's the kind of thing Astaro's Security Gateway would implement.

Steve: That'd be – oh.

Leo: They're very up-to-date. You know, the new 7 came out. This is a commercial, by the way, right now. Just for those of you who are...

Steve: That was a smooth segue, Leo.

Leo: ...a little puzzled by the segue. No, it just made me think of this, and I thought, well, let's do the – Astaro, of course, has been a sponsor of the show for quite some time now. Last 30 episodes, something like that. And through the rest of this year. And we're just really thrilled. They did just update to Release 7. And I'm wondering – of their Astaro Security Gateway software. And I'm wondering if that's – they added so much new stuff, including – and we talked about this before – transparent email encryption and decryption. So it happens at the gateway, so users don't even have to worry about it. It's automatic. Digital signatures. It's based on S/MIME or OpenPGP. Really slick. There's also a secure remote access via SSL built into its VPN, so that's much easier. SSL VPN is so slick. So now IPSec, L2TP over IPSec, PPTP tunneling with SSL, you've got it all. In fact, I think it's the only UTM appliance on the market with this kind of flexibility and VPN and remote access. You can cluster them so you can scale. I mean, it just goes on and on.

Try it, v7, in your business. A free trial is waiting for you. All you have to do is go to Astaro.com. And now the best news for you noncommercial users. And I'm really pleased to hear this. They continue to offer the home use package absolutely free; but now with v7, not only do you get the software absolutely free, you get the base license, all subscriptions, and Astaro up-to-date, the complete set, free. They used to charge 79 euros a year for this. It's now absolutely free. That's limited to 10 IPs, 10 users, and a thousand concurrent connections. But basically this is such a good deal. Try Astaro v7. And if you're a noncommercial user, you've got to check it out. Astaro.com, or you can call Astaro and get a free trial of the Astaro Security Gateway appliance in your business: 877-4AS-TARO. And I will ask them about port knocking because I wouldn't be surprised.

Anyway, port knocking. Is that something that's brand new? Because I hadn't heard about that till very lately.

Steve: It's been around for the last year or so. You need something on the client side to generate the knocks. And that's kind of tricky because you'd really like to use, well, you'd like to use some sort of raw packet technology. But there are ways to do that with UDP packets, where you're able to aim them at your IP on the fly. And so you need something at each end. I just think it's very clever. I mean, what you could also do is just use a single packet containing

your own, basically, a password as long as you want. And just when the router sees any packet, it takes a look at it, and it checks to see whether it's your magic authentication let-me-in packet, and then takes off from there.

Leo: That might be a little more risky because couldn't somebody intercept and catch that?

Steve: They could, except that then what you would do is to take care of that you would base the detailed – actually I've given this some thought, Leo, because I really find it...

Leo: You're going to implement it, aren't you. You're going to do it.

Steve: I'm thinking about it. What you would do is you would base the content of the packet on your current source IP and port.

Leo: Perfect. So you [indiscernible] IP address or something.

Steve: Exactly, and nobody could spoof – no one could get a connection that spoofed your IP because they would be at a different IP and would never be able to – so basically it would really allow traffic only from that remote IP.

Leo: Unbelievable.

Steve: It's kind of a cool idea.

Leo: A closing comment from an avid listener, Paul Fisher, with Hughes in Illinois. He says: In Episode 79 of Security Now! you mentioned that two thirds of all mail on the Internet is spam. That's what we're hearing from a lot of different sources, including some of the companies that actually filter the stuff. He says: You probably have access to better information than I do. But I thought you might be interested in another data point. The mail servers at our company – Hughes is huge, 35,000+ worldwide users – has a spam rate of, get this, 99.7 percent. That is 99.7 percent of all email bound toward our email servers is identified and blocked as spam. We get very few false positives and very little complaints from users receiving unwanted spam in their mailboxes. It seems unbelievable, but that's what our servers have seen for some time. I really enjoy the podcast, hope you continue it. I'm also happy that you can force Leo to do one every week. He knows that I'm a lazybones. It's the one TWiT podcast I can always rely on. Fortunately, I have no stories for you about SpinRite. I own a copy, but I've never had to use it. I'm hoping just owning it will keep my drive from going bad. Yes, you've terrorized your drive.

Steve: That's right. If the drive knows you've got SpinRite nearby, it'll say, okay, I give up, I give up, I'll keep working.

Leo: I'll keep working. Actually his number matches my number. But I think when they say 70 percent of all mail worldwide is spam, that's a conservative number.

Steve: I think so, too.

Leo: I'm much more like 99.8 or 9 percent. In fact, I think it's 99.9 percent. It's awful. But fortunately, the good news is there are good technological solutions that get rid of – good filtering solutions that get rid of most of the spam. Most spammers don't make much of an attempt to hide their spam. And those that do get through, but...

Steve: Yeah, because they're just on the shotgun profile. It's just like we're going to send this to as many gazillion people as possible, and a few of them will get through, and that's going to pay the bills, so.

Leo: They don't care. Sons of guns. Well, Steve, we've come to the end of another mod-4 episode. Boy, this was fun. So much interesting material.

Steve: Yeah, we've had some great feedback from the Q&A people. And I ran across a number of pieces of email that I sort of like as this closing comment thing, so that we may add that as a feature to Security Now!, just so we can begin to, when they're relevant and interesting, include some of our listeners' comments.

Leo: We do listener mail in every episode of the Daily Giz Whiz, and it's in many ways one of the most fun parts of the show.

Steve: I think it's a great idea. And next week we're going to talk about this amazing study that Google did on what they discovered about the reliability of their monster hard drive farms that they've got that runs Google, by popular request. I was going to make a little mention of it this week, but so many people wanted more coverage that I said, okay, let's do that.

Leo: Oh, absolutely. I mean, this is an area where you're probably one of the top experts in the world. So it's absolutely something we should talk about.

Steve, we've come to the end of a fabulous episode. But it's not over. You can get more at GRC.com, that's Steve's site where you'll find SpinRite, but also all of his great free security utilities, now the 48-million-user – there are not many people who could say that anything they've done have 48 million users.

Steve: And I have to say that's a really clean number. I go to the trouble on ShieldsUP of maintaining a list of the most recent 4K, that is 4,096 users. So anyone who's poking around the site, using ShieldsUP for several hours, half a day, I am not multiply counting them. If their IP is already in the list, I just move it to the top of the list and only – so basically I'm only counting people who are gone long enough to have their IP fall off the most recent 4,096 users. So it's a really good number. And you're right, 48 million is very cool.

Leo: That's just amazing. Also we've got 16KB versions for the bandwidth-challenged, and Elaine's great transcriptions so you can read along. In fact, we'd better wrap this up so Elaine will...