
# Listener Feedback Q&A #9

**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-048.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-048-lq.mp3

---

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

This is Security Now! with Steve Gibson, Episode 48 for July 13, 2006: Your questions, Steve's answers.

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

It's time to take a look at security once again with our security expert and the wizard of ShieldsUP and SpinRite, Steve Gibson.

**Steve Gibson:** Hey, Leo.

**Leo:** Good afternoon, Steve. How you doing?

**Steve:** Great to be back. Good.

**Leo:** So before we get into our regular – because it's Episode 48, our Mod 4 episode, so we're going to do a question-and-answer – before we get into that, a couple of items we want to cover.

**Steve:** Yeah, I did want to mention that anyone who's got their computer updating its security from Microsoft, anyone using Microsoft Windows, this the other day was the second Tuesday of the month of July. And Microsoft once again has a little packet of security updates. Nothing super critical for end-users. There's an Excel problem, you know, various obscure things. However, not accepting those and keeping your system current means that they pile up. So you

might as well do so. It will require a restart of your machine, so you might want to choose when to do this. There were a couple server-side things that were a little more of a concern, but nothing, you know, super extreme. So...

**Leo:** You know, I have mine set, and I recommend others do this if you've got Service Pack 2, it's one of the options available to you, you just open the system properties control panel and go to Automatic Updates. And I've set mine to "Automatic (recommended)." Which says, "Automatically download recommended updates for my computer and install them." However, it doesn't install them. I just noticed that, as you're talking, yeah, it downloaded them, and I have little shield in my system tray saying "Updates are ready for your computer. Click here to install."

**Steve:** Well, and I'm a little bit more of a control freak, or at least I'm a little more protective of someone else doing things to my machine. I mean, there are things that I explicitly don't want installed. And in fact you can go to Windows Update and say I don't want these things, and don't ever try to give them to me again. Just stuff, you know, I'm not using and don't need, apps I don't have installed that they're doing security updates for. I've got mine set for "Download but don't install."

**Leo:** Right.

**Steve:** Which actually may be the way yours is set?

**Leo:** No, it's not. That's what's odd. That's what I'm getting is a download and then let me know. Of course, I don't get – what it's not doing is saying what the different updates are, just giving – when I click the little shield it says – well, I guess I can do a custom install. So it is exactly as you say.

**Steve:** Right. And I do...

**Leo:** Even though I've said, well, I've just got to point out to people that, even if you have it automatic, you might want to check your system tray and see if there's that little yellow shield there because in my case the automatic doesn't do it automatically.

**Steve:** Yeah. And, you know, I do like to sort of peruse the menu, see what it's going to be doing. Just sort of I sort of know what's going on. I dislike the idea of this stuff all just happening in a completely automatic fashion.  Also...

**Leo:** But don't you always install the critical updates anyway?

**Steve:** I do. But the other thing is, unlike some people who power cycle their system constantly, so they're always rebooting, I am normally not rebooting my system unless there's a reason. I mean, I'll go months, literally, without rebooting my system. Also I generally have all kinds of stuff going on at once. I mean, I've got lots of apps running. I've got, you know, stuff spread across three screens. And so just restarting my computer at an arbitrary time is not an option for me. I need to, you know, deliberately close up loose things, save things that aren't saved, and, you know, sort of tidy up before I do a reboot. So I don't want anything telling me you must reboot now. And boy, if you install those updates and then aren't prepared

to reboot pretty soon...

**Leo:** It'll bug you.

**Steve:** It pesters you to death. It's like, okay, fine already, I'm going to reboot.

**Leo:** That's why I have not yet pressed the little yellow shield for the updates.

**Steve:** Right.

**Leo:** And I will as soon as we're done recording this podcast.

**Steve:** Well, and the other little pre-issue here is we got a lot of reports last week from people who were having trouble downloading the previous episode. It cut off after a few minutes.

**Leo:** Episode 46, I think it was.

**Steve:** I think that's two weeks ago, exactly.

**Leo:** Yeah, yeah.

**Steve:** And so I wanted to mention that, you know, that AOL is using Akamai's distribution system. And every so often – not often...

**Leo:** Well, I can give you – let me give you – I can give you the inside story on that. I screwed up. So it doesn't happen if I do it right. That was a case where I'd uploaded it to CacheFly or some other server – oh, no, I know, I uploaded it to your server, and then I did a cross copy. And instead of copying it to a staging folder, which I usually do, I copied it to the main folder. And what happens is, when it's Thursday, and people know that Security Now's about to be posted, they start hitting the server, knowing that – because we use a consistent filename and convention – knowing that that's going to show up. If you hit the server as I'm uploading, your particular Akamai caching server, and there are many, will cache a partial version of the file. So those people who share that server with you will from then on, until Akamai updates, get a partial file.

**Steve:** So one guy zaps it for everybody else.

**Leo:** Well, not everybody else. And that's why what you probably saw, and certainly I saw, was it was a small percentage. It wasn't by any means everybody who downloaded it.

**Steve:** No.

**Leo:** And there were couple of different file lengths. And that's because two different caching servers had two different partials. So the trick on that – what I did was, as soon as I figured out that that had happened, is I reseeded the server with a new copy. Akamai makes a hash, as far as I can tell, makes a hash of the file. When the hash changes even the slightest bit, it says, oh, there's a new file, and it will reseed its servers. So I did that. Now, if you've already, you know, if you're using iTunes and it's already downloaded a partial copy, it won't know to get the full copy. So unfortunately what you have to do then is remove Security Now! from your podcasts and resubscribe.

**Steve:** Ooh.

**Leo:** And at that point it'll go, oh, and it'll download a full new version. But as you said, you can always get a 16KB version from you. Because you don't use Akamai.

**Steve:** That's true. And so what I wanted to say was for people who are frustrated that they, for whatever reason at the moment, they can't get the full one, that the lower quality quarter-size is always available from GRC, just as a fallback. So I just wanted to reiterate that for those people. And I'm glad for the explanation. You'll do it the right way from now on?

**Leo:** Well, I can't promise always. I mean, I knew that this would happen. I figured this out the last, you know, when we first started using Akamai this was happening a lot.

**Steve:** Right.

**Leo:** And I finally figured out, oh, duh, and this makes sense. The way I do it is upload it to a hidden folder so that people can't download a partial. And then I copy it from the hidden folder to the visible folder. And presumably that copy happens so quickly, or quickly enough, that nobody is able to kind of start to download a small partial version of the...

**Steve:** Boy, it sort of sounds like an imperfect distribution system, though.

**Leo:** I think it's the way caching servers work. So, yeah, I can't promise it won't happen again because I'm stupid, and occasionally, as I did last on 46, I make mistakes. But I do know what causes it, I believe, and I know how to avoid it. And so assuming no operator error it shouldn't happen. Shouldn't happen very often, let's put it that way. And I'm sufficiently chastened, and I'll pay attention. I just wasn't paying attention when I transferred it over. I think it was I had forgotten to upload it, and it was a late night, and agh, I'll just get it over here. I got it from you, Steve, and then...

**Steve:** I think you were in Toronto, weren't you?

**Leo:** Yeah, I think that's what happened. I was in Canada. So apologies for that. But when that happens, yeah, that's what – you can either get a 16KB version from Steve or wait a little while. And once I've seen enough screams of pain, I'll fix the thing.

**Steve:** Cool. Well, we've got some great questions.

**Leo:** Oh, I'm excited. I've got my list in front of me. Shall I be your reader again?

**Steve:** Sounds great.

**Leo:** Starting with Fred from Mountain Home, Arkansas, who says he's becoming increasingly comfortable about using online banking. The same is true for...

**Steve:** Actually uncomfortable.

**Leo:** I'm sorry, uncomfortable. I'm becoming increasingly comfortable. He's more uncomfortable. Same is true for paying for online purchases by PayPal, removing funds directly from my checking account. Are these practices safe? My bank account uses a 10-character password. I use KeePass Password Safe to copy and paste the password into the login page for my bank. Does this copy-and-paste procedure eliminate keyloggers – boy, this guy is really paranoid – from being able to read the password?

**Steve:** Well, I liked his question because he says he's becoming increasingly uncomfortable about using online banking. What I think is happening is he's becoming increasingly aware of security...

**Leo:** Right.

**Steve:** ...and of the problems. I mean, he's obviously now sort of aware of what a keylogger can do. He's just – he's uncomfortable with, you know, the possibility of bad things happening. So I'm glad because this is showing an awareness of, you know, what can go wrong. At the same time, obviously he wants the convenience that online banking provides. He also mentioned PayPal removing funds directly from his checking account, which brought up something that I wanted to mention about PayPal.

I use PayPal, and we've talked about it on the show several times, because it allows you to transfer funds to a third-party site without giving them all of your credit card or personal details. They get the money, but PayPal does the transfer. So, I mean, PayPal could screw up, too, potentially. But the idea is you'd like, from a security standpoint, you'd like to minimize the dissemination of your personally sensitive information where possible.

The problem I have is that – and maybe you know why, Leo – PayPal really wants a bank transfer and not a credit card transaction. I have both. They required that I give them my bank account information in order to, like, verify me, whatever that meant, even though I was using my credit card for a long time before then. So now they have both. But I have to manually go in every single time and change this, you know, it's like, from don't debit my checking account, I want you to take it from my credit card. The reason I want that control is, if something happened, I was buying something online, for example from eBay, and I didn't receive it, or it wasn't in the condition that they've promised, I want the control of being able to challenge that charge on my card. If PayPal has sucked it out of my bank account, that money is gone. And so I prefer manually forcing PayPal to pull the money from my credit card because then, you know, I still have that money because my credit card will stand on my side, and I'm able to, you know, deny that charge. Maybe that's why PayPal doesn't like the idea is that, you know, it's just more complex for them.

**Leo:** I think it's mostly cost. I can address this a little bit because we use PayPal for donations. In fact, it's the only way you can donate to TWiT.tv. And that's easy for us. We don't have to take credit cards, and I think it's more secure. Some people are nervous about PayPal. In its early days PayPal was kind of notorious for some security issues. But they've been owned by eBay for a couple of years now. I think they're very secure, and I think you do have the advantage of only one provider knowing your credit card number.

**Steve:** Yup.

**Leo:** For instance, when you pay us at TWiT.tv, I don't see any of your financial information. The money is transferred into a PayPal account. I think the reason PayPal would prefer to use checking – and I agree with you, Steve, I don't use – when I want to make payments I'd far prefer to use a credit card than my checking account. But of course the fees are lower if they use your checking account, and I think that's probably the primary reason.

**Steve:** Ah, okay, right.

**Leo:** As you know, credit card charges are very expensive. So for whatever reason, they do really encourage you. You can, in fact, you can actually run up against a limit of credit card charges. Eventually they'll say sorry, you can't do this anymore, you have to become a verified PayPal account with your checking account. The way they do it's kind of interesting. You give them the checking account number, and they transfer two small, penny-sized deposits into your account. And then you check your statement and see if they came in, and verify by saying what size, you know, how big the deposit was – you know, it's two or three cents, a nickel, 12 cents – and by doing so verify that you in fact own the account, you have access to it. And it seems to work pretty well. I haven't heard of any large problems with PayPal.

They do have a dispute resolution service that I think works quite well. In fact, a number of people have been using it. I get disputes every few days from people who forgot that they signed up for a yearly subscription to TWiT. And when their yearly $20 is deducted from their PayPal account, they go, hey, I didn't get my money. What's going on? And so I'll give them the money back, and that works very well. I mean, it's a very straightforward dispute resolution system. It doesn't have the force of law that the credit card dispute resolution system has, though. So I think you're probably right that when you're buying something it probably is better to use a credit card. Now, I don't know, since the credit card is with PayPal, the payment goes through PayPal, I'm not sure if you have the same protections against the merchant.

**Steve:** Well, on my side, I mean, it's my normal regular credit card that's registered with PayPal.

**Leo:** So you could still say to the credit card, don't pay this.

**Steve:** Oh, absolutely. Absolutely.

**Leo:** Yeah, yeah, even though it's through PayPal, not to the merchant directly.

**Steve:** Correct.

**Leo:** Yeah. I think PayPal's dispute resolution works pretty well. But I don't like anybody putting – having access to my checking account. I have to say, having – I had to do that to become a – to be able to accept PayPal. And I haven't had any problems, nor have I heard of any significant problems.

**Steve:** And it's funny, from a consumer standpoint, on my side, when you mentioned that credit card limit, that is what finally induced me to give them my banking information. And there was also something about my address being unverified.

**Leo:** Right.

**Steve:** I mean, there was a lot of pressure on me to give them my banking information. And they have it now, and it defaults to using it. So every single time I make...

**Leo:** That's a pain, yeah.

**Steve:** ...a payment, I have to go in, and I manually override that and say, no, take the money from my card. Because again, you know, if I'm buying something that may not work out, I want that control. It's worth mentioning, too, there was a little news blurb – I haven't seen lots of attention or press on this. Google has now started a payment service in competition with PayPal.

**Leo:** Yeah. And in fact I'm looking into using it. There's some real advantages. If you use Google AdWords, the costs are much lower. In fact, it's probably free. So if you're a merchant that uses Google AdWords to buy your Google ads, there's a lot of incentive to do it. Also if you have Google ads, and you accept Google payments, you'll see a little button in your ad all of a sudden which makes the ad stand out a little bit more. So I think for anybody who buys Google ads there's a lot of incentive to do this.

**Steve:** I think it's called CheckOut, isn't it, Google CheckOut?

**Leo:** Yeah, yeah. And eBay doesn't like it. eBay, who owns PayPal, has refused to accept Google CheckOut. So there's a little battle going on. You know, I think that we will become more and more comfortable with these kinds of online transfer systems. And I think that they'll eventually take over. They make a lot of sense. Maybe it'll be a Microsoft Wallet, maybe it'll be some other company. But right now PayPal is in the forefront. And because they're owned by eBay, and there's a lot of scrutiny of them, I think that they're trustworthy.

**Steve:** Also it's worth mentioning, there was a credit card transaction system called SET that never got off the ground. Apparently it's limping around somewhere. But it offered exactly this kind of third-party insulation, where you and your credit card company would together pay a third party. The third party would never have access to your information directly. And so, you know, again, it's a variation on this idea. And I agree, in general, in the long term, I wouldn't be at all surprised if people who are receiving money don't end up with your personal information, they just get the money from a third party that is isolating them. And again, it

allows you to control how many individuals you must trust with this, you know, financial information.

**Leo:** I think ultimately it's more secure.

**Steve:** Anyway, to finally...

**Leo:** Yeah.

**Steve:** ...to finally answer Fred's question about if he uses Password Safe to copy and paste his 10-character banking password, yes, copy and pasting would probably bypass keyloggers. But there are many technologies that it would not bypass that are, like, reading fields or watching the data leaving your computer and so forth. So, I mean, if your computer – and there are several questions this week that sort of involve this issue of, you know, how safe am I if my computer's got bad stuff in it? Well, the answer is you're not. I mean, so it is so important to keep the bad stuff out of your machine. And we will see various takes on that in the Q&A that we answer this week. But there were a lot of questions that sort of were involved with, okay, if something's bad in my machine, how bad is it? It's like, well, it's really bad.

**Leo:** Kind of the wrong question to ask. If you think you have a keystroke logger, don't worry about whether cutting and pasting's going to work. You've got bigger problems.

**Steve:** Anyway, so Fred is increasingly uncomfortable with online banking. And so I wanted to address that, you know, the issue is that he's not wrong to be increasingly aware of the problems. And so, you know, if you're using online banking, certainly the security of your machine is more of a concern. I remember that, you know, years ago people were saying, ah, I don't worry about computer security. I don't use my computer for anything other than, you know, surfing the 'Net and, you know, doing some Google research and things. Well, and my feeling was at the time, well, yes, that's today. But, you know, it's very clear that services are going to be offered over the Internet, and the security of your machine is going to become increasingly important as we're depending upon them more and not just using them as an Internet surfing toy, but as, you know, more of our life, which of course is what online banking is doing.

**Leo:** Indeed. Matt Jordahl of Laveen, Arizona, says: You've recommended the Kerio Firewall – actually the Kerio Personal Firewall is the one you recommend, the free one from Sunbelt Software – several times on the podcast. He says: I just downloaded it to try it out, and I thought it was great until I realized it was limiting my speed on network file transfers. I have a Linux box with several large drives that run Samba, and I'm connected to it over a gigabit switch. Now, normally I can transfer files either way at around 40 megabytes a second. And with Kerio installed I was limited to 20 to 30 megabytes a second receiving, 10 megabytes sending. Have you ever heard of this before? I have a pretty new motherboard, socket AM2, using the nForce 5 SLI chipset – in other words, he's got a very, very fast computer. It's got a built-inboard gigabit NIC. Why is it slowing down?

**Steve:** Well, there's an interesting question that I liked because it raises the specter of do our personal firewalls impede and slow down our transfers in any significant fashion. It's something I remember being aware of because the early personal firewalls, if you were transferring data over the 'Net, you could see your CPU being a lot busier with the firewall involved than it would be if the firewall weren't involved. Now, Matt's case is really extreme. I mean, he's got gigabit

Ethernet. He's talking about a local transfer among machines. And he's, you know, 40 megabytes per second that he's normally getting is 360 megabits per second. So, I mean, he's really got normally super high-speed transfer.

So first of all, in Matt's case, yes, the firewall being another layer of filtering is having to take a look at every connection, and on some level every packet. You know, it's moving it through another layer of software which is slowing things down. Now, the question would be whether he could change his permissions on the firewall or get it less involved in his LAN traffic than it is in his traffic going out onto the internet. But for most people who are using a firewall to insulate them from online stuff, they don't have a gigabit connection. They're not transferring, and can't, 360 megabits of data that, oh, darn, is going to be throttled down to only 160. I mean, you know, they've got maybe a megabit if they're lucky. So my point is that, in general, firewalls are going to take a tiny little ding in your per-packet performance. But it's, I mean, it's nothing to worry about because what's happening is the Internet's overall delay across the 'Net is completely masking any little packet delay with the packets coming in.

In a worst-case scenario, which is what Matt describes, where there's no effective per-packet delay in the network because he's got a gigabit network and a fast server just pumping files into his system, there you're going to see a firewall effect. But normally, absolutely none. So for most users who aren't in exactly Matt's situation, there would be no problem.

**Leo:** True of a router, too?

**Steve:** Yes. Yes.

**Leo:** Anything that's going to process the traffic is going to slow it down a little bit. I measured the iPhantom, which is kind of an interesting device. It goes out over the Internet and uses the Phantom Technologies servers to do the processing. And that's about 10 to 20 percent. Given that it's traveling out over to the Internet to their servers, where it's being sanitized, and then out to the real world, that seems like a minor difference.

**Steve:** Yeah, well, and there it's not quite applicable because...

**Leo:** They're doing a lot more.

**Steve:** You're actually running through, yes, through their own servers and doing encryption and decryption and so forth.

**Leo:** They're doing antivirus and – yeah.

**Steve:** Right.

**Leo:** But even then, 10 to 20 percent's not bad. What would you say, it's 5 percent or less for a normal firewall or router?

**Steve:** Actually, it's zero. And I guess I didn't describe it too well because I didn't want to go on about this. But the delay in, I mean, the Internet is all about...

**Leo:** It's waiting for the Internet anyway.

**Steve:** Exactly.

**Leo:** Is what you're saying.

**Steve:** The Internet is...

**Leo:** So while it's waiting it has time to do what it does.

**Steve:** Well, the Internet's protocols are designed brilliantly to hide the delay. For example, computers are able to send packets ahead of time, and the Internet only acknowledges them. Acknowledgments can come late, and everything still goes. So the guys who designed this packeting system understood that they had to have protocols that could intelligently anticipate the size of the buffers available at each end and send things in advance of technically receiving permission or acknowledgments. So zero is the overhead in a typical situation because your firewall is negligible compared to the trans-Internet delays, which are completely masked by the brilliant protocols we have.

**Leo:** Now, you did one other thing, and I'm just going to bring this up because I don't want you to get email saying your math is wrong.

**Steve:** Oh, I...

**Leo:** It's not wrong. I know why you did it, I just wanted to clarify it. You said 40 megabytes a second is 360 megabits per second.

**Steve:** Oh, 320.

**Leo:** Well, but I was wondering because isn't there – so it's normally 8 bits per megabyte.

**Steve:** Right.

**Leo:** But isn't it – or whatever it is, you normally multiply by 8.

**Steve:** 8 bits per byte.

**Leo:** 8 bits per byte, I should say. But sometimes there's an overhead byte, a ninth bit. And I thought maybe you were including that in your calculation.

**Steve:** Nope, I just multiplied wrong.

**Leo:** It just – okay. So I did save you some email.

**Steve:** Yes.

**Leo:** Brian in Toronto asks – I thought you were being really clever. And there you go. I give you a lot of credit, Steve. Brian in Toronto asks, again, about Kerio's personal firewall. He says it has a network intrusion protection, I'm sorry, prevention system – I like to call it NIPS – and a host intrusion system, HIPS. Can you explain what intrusion prevention systems are and what type of attacks they prevent?

**Steve:** Yeah, real quickly, without getting into the specific details because I want to focus on the question, he talks about a network intrusion prevention and a host intrusion prevention. I thought this was a good question because it highlights that personal firewalls are beginning to evolve and take more responsibility for host-side problems, i.e., rootkit-style attacks. There are many ways in which software running in our computer is able to involve itself with other software in our computer. A classic instance is, for example, a programmer's debugger which is able to reach into another program and stop it and single-step it and allow it to be analyzed. Well, the fact that one program, called the debugger, is able to insert itself into another process, you know, that's nice, except that it also means that malware is able to do the same sort of thing. So there isn't much isolation among processes, not nearly as much as most users would like.

And I'm annoyed that, in fact, that the security is as bad internally as it is. You know, the people who say, well, but what can you do, once malware is in your computer, anything could happen. Which is true, and there are other technologies we'll be talking about in the future. We're going to do some episodes on virtual machine technology, Leo, to talk about, you know, what virtual machines are and how they work. So it is possible to create very good isolation. But by default, ROSes don't typically provide that kind of protection. So what's happening is that personal firewalls are acquiring features, not just about the incoming network traffic, but about the behavior of the programs in the computer themselves.

And so that's what this whole difference is between network intrusion prevention, which is just packet stuff on network communications, and a host intrusion prevention. It deals with behavior and catching malicious conduct of programs, like rootkits, trying to mess around with your computer internally.

**Leo:** Is this related to IDS, which – I've never heard of the HIPS and NIPS, but I have heard of IDS, Intrusion Detection Systems. And many firewalls have those.

**Steve:** Yeah. IDS is a technology basically that looks at the traffic going by. So it's not a firewall from the standpoint of simply blocking or allowing packets. But an IDS actually interprets the traffic and tries to find bad things going on. For example, you might have an IDS on front of a server that is looking at, like, bad conduct in the URLs which remote users are submitting. There have been many problems in, for example, in the past with Microsoft's IIS server, where you could basically take it over by using a malformed URL. So an IDS could be in front of that, sort of pruning and purifying the traffic on the way in, and also detecting anything that looked suspicious.

**Leo:** Got it. Doug Dorbuck, writing from his Hotmail account, has a question about rootkits. He writes: If you have a secure firewall installed in your system, such as ZoneAlarm or,

again, the Kerio firewall – that's three in a row now – will these be able to detect the traffic generated by a rootkit in the event your system gets infected by one? You know, he's talking about the fact that software firewalls watch outbound connections. He says: Fortunately I haven't been infected by a rootkit, but I always closely monitor any communications of outbound programs or services. I imagine that, unless the rootkit installs its own TCP/IP stack, some of the better software firewalls will detect the traffic. True?

**Steve:** So here he's asking, if I've got a personal software firewall whose job it is to control the outbound traffic on my system so that I know which applications I'm giving permission, is there a way around it. And the answer is yes. The firewall vendors have certainly done everything they can to prevent being circumvented. And so – and over time they have gotten much better. The very first version 1.0 software firewalls would have been easy to circumvent if there was malware that was smart enough to do so. Back then the malware wasn't that smart. So it's been an arms race. The malware is getting better. The firewalls are also getting better to keep pace with the malware. And, I mean, this is what the firewall companies are doing is working to keep their products as secure as possible. The problem is, as we've said, once something is in our computer, all bets are off.

**Leo:** It could turn the firewall off.

**Steve:** There are...

**Leo:** They frequently do.

**Steve:** There is malware, yes, which has turned off well-known firewalls. And there has been malware which knows about many firewalls, so it's able to deal with whatever you happen to be using. Then the firewall vendors countered by making their firewalls much harder to turn off. But again, the rule is, if something's in your machine, you just can't be sure what's going on. So the outbound monitoring is more useful for non-malicious programs that you want to control than – you know, technically – than for something that is absolutely determined to get data out of your machine. The problem is, it's just all software, and it's all pretty much running in the same environment, using the same services. So, you know, it's an arms race.

**Leo:** They don't call it malware for nothing.

**Steve:** Right.

**Leo:** Lydell Anderson of East Hartford, Connecticut, has listened to all of our podcasts. He says he's read the GRC pages on DoS, DDoS, and DRDoS attacks. DRDoS? What's that?

**Steve:** That's Distributed Reflection Denial of Service.

**Leo:** I thought it was the old Digital Research operating system. Shows what I know. He still has a question, though: How can I protect myself from a DDoS that's aimed at port 80 on my web server? I can't, obviously, I can't block traffic to port 80. That would defeat the

purpose of the website. So what do I do to protect?

**Steve:** Well, this of course follows on last week's talk about Internet Weaponry. And I liked the question because, you know, it sounds like here's a guy who's got a web server. And maybe he's had problems with denial of service attacks before; maybe not. But he wonders, what can he himself do? The answer, which really follows from what we were talking about last week, is unfortunately not much. I mean, virtually nothing.

**Leo:** This is comparable to a DoS attack on your router.

**Steve:** Well, yes. The idea would be that he needs to offer port 80 to the world so that anybody anywhere is able to send traffic to port 80. The problem is, that means that botnets everywhere in the world can send traffic to port 80 and flood him.

**Leo:** And it's legitimate traffic. You can't distinguish it from illegitimate traffic. That's the problem.

**Steve:** Right, right. Now, if he were hosted by a company which was offering protective services, like we were saying, you know, for example, some gambling sites now are using that technology in order to be more resistant towards – against really large attacks. If he were there, then he'd be protected by the services offered by his host. But most servers in the 'Net, because that kind of protection is very expensive, are not. And so denial of service attacks are just something you tolerate. It's just, I mean, there isn't a simple solution for this guy, unfortunately.

**Leo:** How do they protect? What do they do to protect you?

**Steve:** As we talked about last week, they just have really big pipes that are able to absorb a phenomenal amount of traffic.

**Leo:** That's it. That's the only thing you can do.

**Steve:** That's really the only thing you can do.

**Leo:** We had talked some years ago about maybe changing the way you respond to the SYN requests by delaying or waiting for a second SYN request. Any of techniques like that work?

**Steve:** Well, there are many different types of appliances now being sold. But fundamentally, because of the way packets convene, as we were talking about the analogy of a magnifying glass focusing the sunlight down to a single point on the palm of your hand, similarly, unless you have a really huge pipe in the first place, that pipe is going to get flooded. And, you know, it sounds like this guy Lydell, you know, just has a web server, and he'd like to be, you know, have protection from denial of service attacks. And he's just, you know, like a regular guy. Problem is, there just isn't any solution. Yeah.

**Leo:** Typically when you buy a hosting solution, especially if you buy a dedicated server, so you spend some money, you'll get a 10-megabit pipe. Now, that's easily flooded by a distributed denial of service attack. 10 megabits is nothing.

**Steve:** Right.

**Leo:** A larger pipe that you might pay a little extra for is 100 megabits. That takes a little bit bigger of an attack. But last week you talked about attacks that were considerably bigger than that, in the gigabit range, so...

**Steve:** 10 megabits is 40 bots, each sending a quarter megabit out of a cable modem. 100 megabits is 400 bots. Those, by today's standards, those are small networks.

**Leo:** So you've got thousands of bots in your IRCBot network. You can choke pretty much any pipe unless the guy's got, you know, gigabit pipes.

**Steve:** Yup.

**Leo:** And that's very expensive. Trust me. I know. Let's see. We'll go to Daniel Hummer of Modesto, California, who asks: Is there a file size limit, or maybe a limit to the number of sites you can block using the hosts file – we talked about that a couple episodes ago – before your system starts really bogging down?

**Steve:** Yeah, we got a bunch of great mail after that. Some people found large files on the Internet, that we talked about, that are being maintained. And they said it's amazing, when they added that large file to their hosts file, all kinds of stuff just stopped, you know, ads and nonsense, because their computer was now looking up in the hosts file before going out onto the Internet to see whether the hosts file provided an IP which, in this case, is a, like, you know, 0.0.0.0 or 127.0.0.1, something other than the real IP address that just prevents your computer from looking any further. So his question, and I really liked Daniel's question, is because some of these files are really big, I mean, they're very comprehensive, they go on and on and on, and in fact it's an education just to read through. You see all these weird sites and things. The answer is no, the hosts file can be very big because Windows is very quick in looking through it. But more importantly, it's the speed of that versus the speed of going out across your connection to your ISP's DNS server and doing a lookup. It would have – your computer would have to be incredibly slow in checking the hosts file for that not to be much faster than making an external request. So there really is no problem with hosts files getting really large in terms of their own performance. And they end up really speeding things up because they offload that traffic from your Internet connection.

**Leo:** Yeah, makes sense. I mean, even if you had hundreds and hundreds, maybe even thousands of lines, it'll just still be a small text file, relatively. And computers are very fast at getting through those.

**Steve:** Especially now compared to, for example, modem, or even a – well, an old analog modem or a cable modem. And it just – it's a second of delay versus just milliseconds to check the file.

**Leo:** Plus think about what you're saving in terms of images and so forth you'd be downloading. I think it's probably always a net gain in speed.

**Steve:** Right.

**Leo:** David Cockrell of Bossier City, Louisiana, wants to monitor his system's traffic, but he worries: Will freeware like TCPView and TDIMon always show network traffic related to these tools that the bad guys use? I was real curious about how those tools would show RDP, for instance, across a Hamachi connection. That's the Remote Desktop Protocol Microsoft uses for Windows Remote Desktop.

**Steve:** Right.

**Leo:** TDIMon did show the traffic across the Hamachi connection. But I wonder if there were a way for traffic to be covered up by the tricks that are used in the bot world. You'd better explain this question because it's a little complicated.

**Steve:** Well, this again relates to, could something in my computer alter the proper behavior of other software in the computer? And the answer, unfortunately, is yes. So we're going to be talking about network monitoring. Actually I think next week we're going to talk about these various freeware and built-in ways to see what your computer is doing at any instant in time. So then the question is, is there a way for that to be fooled? And unfortunately, yes. If something is in your computer, it can do anything it wants. Now, tools like TCPView and TDIMon by the Sysinternals guys are terrific tools, and they're great. But there again, you can only really know that they're going to be doing the right thing for monitoring programs that are sort of playing by the rules. There are ways that they can be circumvented if somebody wanted to get traffic out of your machine and its software, maliciously trying to do so, it probably can. Or at least you cannot know that it can't.

Now, there is an alternative, though, and that is, if you were to run these tools in a separate machine, that is, if your – for example, if you have a router, either a consumer router that shows you what's going on or a Linux system or a BSD UNIX machine, which is on your network perimeter, it has the ability not to be victim to any software running in your host machines. So if someone really wanted to know for sure what was going on, the idea would be to get that monitoring software out of the environment that you're trying to monitor, put it on a separate machine which you are not using as a workstation, it's not going to be victim to random Windows infestations and malware and things. And by virtue of being isolated by a network connection, software is not going to have the same access, malicious software is not going to have the same access to your network monitoring tools that it would if those tools were running in the same system. So there is a way to do this by separating those functions.

**Leo:** And that sounds like it's always a good idea for security, but of course it requires extra hardware and so forth.

**Steve:** Right.

**Leo:** Shawn Doyle writes from somewhere on Planet Earth – did he say that, or did you?

**Steve:** No, he actually put that down as his address.

**Leo:** Somewhere on Planet Earth. My goal is to run a game server that my friend and I can play together on. My problem is the game server requires information from port 8080, a standard Internet port. But as soon as I turn on port forwarding to make that port available, your ShieldsUP service at GRC.com – highly recommended, ShieldsUP, GRC.com – shows a big red flag indicating I'm open to attack from the whole Internet. Is there a way for me to do this safely?

**Steve:** Well, this was a great question because it highlights what we've talked about with port forwarding. Basically it sounds like Shawn's behind a router, which is good because it gives him protection from any unsolicited inbound traffic. The problem is, his friend is unable to connect to his server, which is running behind his router. So he needs to do port forwarding, that is, open up a port on his router so that unsolicited incoming traffic to port 8080 is able to get through into his server.

The problem is, then, if he checks his security with, you know, my ShieldsUP service at GRC, I'm checking incoming traffic on port 8080 because it is a commonly used port for web servers. The normal port is port 80; but non-high-permission servers, servers that don't have root permissions, which are not allowed to open ports below 1024, often will run on 8080, just because it's a higher port that they have access to. So the cool thing is that Shawn doesn't want to offer that server to everyone in the world. If he needed to offer it to everyone in the world, like that prior question asked, where he was just wanting to run a server on port 80 and let everybody in, the problem is everybody could be malicious. In Shawn's case, he only wants one friend to be able to connect. Well, if he knows that friend's IP address, he can only allow traffic in from that one IP.

Now, if his router allows him to set up basically a firewall rule – many of the newer firewall routers will allow you to say the remote IP needs to be this on port 8080, and send the traffic on. So in his router he might be able to permit only that, you know, his one friend to have incoming traffic on port 8080. But even if the router doesn't do that, his server certainly can. So because any traffic coming through the router will be forwarded to his server machine, a firewall running there could be instructed, drop everything, don't bounce anything, don't respond, be stealthful on port 8080 except for traffic coming from this one remote IP. So there should be a way he could configure himself. And of course he can use ShieldsUP to make sure that port 8080 is stealthed for everyone but his one friend.

**Leo:** I suspect he, like I have, has a cheap old router that doesn't do the firewall rules. So that would be a good reason to upgrade to a newer router.

**Steve:** Either that or he ought to be able to do the same thing on his server.

**Leo:** Right, right, right. Vericha – Verha – I'm usually good with names. This one's eluded me. It's probably...

**Steve:** Verachert?

**Leo:** It's in Flemish, which is part of the problem. Verachert Armand in Antwerp, Belgium, was worried – I hope it's in Flemish, I hope I'm not making something up – was worried by a website he visited. He says: A firewall test revealed that my private address – whether

that's whatever, 192.168 or 10.0.0 – can be seen from the Internet. Can it be reached from the Internet?

**Steve:** Yeah. This is so annoying.

**Leo:** It happens all the time, though.

**Steve:** It does. And that's why I wanted to bring it up, is that there are websites that terrify people by saying "this is your private address." And they'll show 192.168.0.1 or .2 or whatever. Well, okay, there are two things going on. First of all, what's happening is he's browsing with scripting enabled, which I frown on, but it does allow the 'Net to work...

**Leo:** You kind of have to, yeah. I mean, even our own site doesn't work without a little JavaScript running. So you can turn JavaScript off, but you're going to get a very different experience on the web.

**Steve:** Well, yeah. You know that I've talked about Internet zones and how you can get the best of both worlds by locking down, if you're using Internet Explorer, by locking down your script processing for unknown sites and only selectively enabling it for sites that you trust.

**Leo:** The thing I would point out, though, that unlike ActionScript, which really is dangerous, JavaScript is sandboxed. And it's pretty difficult. In fact, I don't know of many exploits, except where there are browser holes that take advantage of JavaScript. So...

**Steve:** And that's the problem, of course, is that scripting is complex, and we have a long...

**Leo:** There could be holes.

**Steve:** ...history of browser holes.

**Leo:** Yeah, yeah, yeah.

**Steve:** So it's just, you know, from a security standpoint, if you can, better not to have it enabled. Anyway, the point is that this site, and there are many of them that are scaring people, is using scripting in the browser to determine the machine's IP address and then presenting it in a Window.

**Leo:** As if they knew it.

**Steve:** As if they knew, exactly. But they may have been sent it by the script, and then sent it back to his browser, or it could just be displaying it in the script locally, saying this is your IP. Well, okay. In the first – so the annoying thing is that he could have full Internet security with a perfectly operating router blocking all traffic, he is not exposed in any way, yet this web page is scaring him because it's looking at his IP inside his computer and showing it to him on his

screen. Which means nothing from a standpoint of security except, as I said, that he's got scripting enabled, which he'd be better off without, if he wants to get into that whole battle.

Okay. The second thing is, it's a private IP, 192.168.something.something, you know, 10.anything, 172.16 through whatever, I mean, and we've talked about this several times historically. Those IP addresses are nonroutable IPs. They are specifically set aside for people to have in their own local networks so that they're able to create networks that don't conflict with IPs on the public Internet. So even if this company had his IP address, or anyone had his IP address, you can't use it at all. It will go nowhere. Packets that are stuck on the Internet aimed at those IPs are just dropped immediately by routers. So he has nothing to worry about. And he went to a website that was unfortunately working to scare him.

**Leo:** Yeah. And I guess that would be the thing I would say is that I guess you could turn off JavaScript, but better to know that it's harmless. There's nothing they can do with that number, even if they knew it.

**Steve:** Right.

**Leo:** And they probably don't know it because most of the time what this is is client-side, and they don't send the number back. They just display it. Brian Voeller in Ashland, Oregon, has been busy. He writes: I would like to offer you a suggestion on combating phishing. That's phishing with a "ph," those online scams or email scams where they try to get you to give up your private information.

**Steve:** Yup.

**Leo:** In hopes that you'll popularize this technique. A few days ago I got a phishing email claiming to be from PayPal. I decided to respond. Oh, boy.

**Steve:** I know.

**Leo:** Already I don't like this.

**Steve:** No, just read it.

**Leo:** I decided to respond with a fake login and password. I used "I_AM_A_PHISHER" as the email address and a series of random letters for the password. I submitted it several times very quickly in hopes that the phishing site was feeding them directly into PayPal instead of saving them for later, in hopes of triggering the PayPal intrusion detection system. If PayPal looked at the logs, this would be 100 percent proof that it was a phisher trying to get in. I would like to know if you think this would be effective, especially if many tech-savvy users would do the same thing en masse. Of course we'd have to submit logs that appeared legit but could be recognized only by the target site as fake. A bit like injecting radioactive dye in the body to trace a chemical's path. I think the best way would be for the users to make up a fake login name or trigger password and enter that, etcetera, etcetera, etcetera. It's funny, we get a lot of emails like this from people who think they have conquered spam or phishing or viruses.

**Steve:** Well, yes.

**Leo:** This is a common thing. Everybody always wants some magic bullet.

**Steve:** Well, and the reason I wanted to address this is there are people who – I mean, and I understand the emotional side of this – who are really frustrated and annoyed by all the nonsense going on the 'Net. And they want to take some action themselves to deal with the problem. It's like people who historically looked at their logs of Internet traffic and took the time to track down every random IP address that was sending junk at their IP and figure out who to send an email complaint letter to.

**Leo:** Or the same thing with spam, tracking it back to the, quote, "originating server" and complaining to the source.

**Steve:** Right. And phishing is the same sort of thing. So, I mean, I don't want to blunt Brian's enthusiasm for this. But, you know, my experience has been that you just, you know, you want as little involvement with this junk as possible. If you get phishing email, just delete it. You know, how much time do you have?

**Leo:** Well, there's a legitimate security concern here, too. By clicking on that link in the email, he's going to a website, and he doesn't know what malware might also exist on that website. It's not merely a form.

**Steve:** That's a very good point. In fact, by having any involvement with email that you suspect is not legitimate, I mean, he says he knew it was phishing email. The only proper thing to do is delete it.

**Leo:** Hit the delete key.

**Steve:** It really is the right thing.

**Leo:** And just get in the habit of not clicking links in email, period, no matter who they come from. Don't click links in email. That's a very dangerous thing to do.

**Steve:** Yeah.

**Leo:** Eric Stauffer in sunny Malibu has been thinking about Internet weaponry, our last show, and DoS attacks. He says: If a website is attacked, as you explained, they cannot track the origin of the traffic because the return IP addresses are often spoofed. However, since it appears that the ISPs are logging user traffic, couldn't that – it says "thanks President Bush," but I'll leave out the politics here – couldn't that information be used to help the community get rid of this malware? Ah, an interesting point. When a website is attacked, ISPs could be alerted to this fact. In a sense, local ISPs would have the list of originating computers. They could send messages to the probably unknown participants in the attack that they may have bots in their system. Instructions could be provided for removing these bots. He's got a hold of something very important, I think.

**Steve:** Well, yeah. I liked the question because it brings up an issue that we didn't discuss in the whole Internet weaponry side, and that is the role of the ISP. You know...

**Leo:** You can't spoof an IP address without, in effect, the collusion of your Internet service provider.

**Steve:** Well, yes. And the idea being that – and I've talked about this, I've written about it on my Denial of Service pages. The idea is that, if you have a bot running in your machine at home, your machine has an IP address which is known to your ISP. And your ISP essentially is offering a range of public IPs to all of its customers. So the ISP knows what are valid source IPs for the traffic leaving their network. That is to say, if there was a bot in your machine using random numbers for the IP of the originating computer sending the packet, that is, the source IP, the source of the packet was just random numbers, those packets egressing, that is, leaving your ISP's network are clearly invalid if the source IP could not be possible. That is, if it is not part of the network that the ISP controls, then it can't be a valid packet leaving the ISP. So there has been historically a real push for ISPs doing what's called "egress filtering," that is, filtering packets egressing, or leaving, their network. And by filtering we just mean dropping. Just drop any packets that your network can't have generated legitimately because the source IP is not one that you own, ISP. The problem is, this involves more work, more equipment, more trouble. It involves the ISP in a way that so far, as far as I know, very few if any ISPs are bothering with.

**Leo:** That's too bad because that single-handedly would eliminate IP address spoofing.

**Steve:** Well, and the counterargument is, well, yes, maybe, just for that one ISP. The thing that users could do, if this became prevalent, then what botnets would do is not randomly generate source IPs, but generate source IPs in the neighborhood of the computer that they're living on. And so that would essentially generate IPs as if from other local users in the ISP's network. So again, good as this idea is, there are ways around it. Which is not to say that it doesn't make sense to do it because it would solve some of the problem but not solve all of the problem.

**Leo:** Now, he's actually proposing that ISPs then use their secret, super-secret powers to prosecute people who are doing this, or at least turn them in.

**Steve:** Or to advise them what's going on. Now, it is the case that an ISP would know because on their network you're part of a LAN. And so...

**Leo:** They know everybody's name and address.

**Steve:** Well, yes. And they know your actual IP. They also probably have your MAC address of your adapter because that's the way, as we know, traffic within a LAN, over an Ethernet LAN, is actually moved by MAC address, which is not spoofable by a bot running in someone's computer. So it would be possible for an ISP to be more involved and, for example, detect that there are bot-infected machines within their network. I've heard of ISPs doing this, but it's not something that's done pervasively at this point.

**Leo:** You know, it's funny because when spam became a problem, Internet service

providers resisted the notion of blocking port 25 or...

**Steve:** Well, of getting involved.

**Leo:** Of getting involved at all.

**Steve:** They just didn't want to get involved.

**Leo:** But eventually we're compelled to. And I think maybe the same thing will happen here with...

**Steve:** Well, and we do see things happening. For example, my own cable modem provider, Cox in Southern California, it's blocking a whole bunch of different ports where there have been problems in the past. They are generally the various Microsoft security vulnerability – 135, 137-38-39, 445, things where there have been problems in the past. The ISP is saying, you know, basically for their own benefit, they don't want bots infecting their network and causing problems because it does create bandwidth traffic that they have to pay for. So at some level there is an expense to this.

**Leo:** Although an interesting point is raised in the Net Neutrality controversy here because if there were, for instance – a security researcher in his blog wrote this up, and I thought it was an interesting argument. I want to thank John [Puit ph] for passing it along to me. It was from – it's actually SANS, and I'll put a link in the show notes to this. SANS is the Internet Storm Center. And it's John Bambenek who says, well, if Net Neutrality were enforced, that would inhibit the ability of Internet service providers to start blocking certain ports or somehow protecting certain ports. I don't think so. I think that a law that prevents ISPs from discriminating against some traffic wouldn't go so far as to prevent them from doing it for security reasons. But it might scare Internet service providers to do that kind of thing.

**Steve:** All the ISP would have to do would be to add to their service agreement a waiver from that.

**Leo:** Right.

**Steve:** Saying that anyone who wants service from the ISP agrees that they are not enforcing Net Neutrality. For the benefit of their subscribers they're going to block the following ports, blah blah blah.

**Leo:** Right, yeah. Our last question. Let me open the envelope here. It's been kept on Funk & Wagnall's porch for the last two weeks. Jason Partlow in Maryland has a terrific question about security zones in Internet Explorer. Terrific because you like security zones. Suppose, for example, that I have the Internet Zone set to high security – that's how Steve does it – and my Trusted Site Zone set to medium security. If I go to a trusted site, but it has a frame or an advertisement on the page that's from another site, will Internet Explorer treat that frame or ad like it's also part of the Trusted Site Zone? Because it is

coming from a different server, isn't it.

**Steve:** Yeah.

**Leo:** If I had my webmail provider, for instance, set it as a trusted site, but I get an email that has a body with a nasty script in it, will I be safe? I presume he means he's on the webmail provider's site, like Yahoo! Mail or whatever.

**Steve:** Correct.

**Leo:** Good questions.

**Steve:** Yeah. And the answer is, Internet Explorer does the right thing. To review very briefly what's going on, Internet Explorer allows you by domain to control the behavior of our browser so that, for example, by default, anything in the so-called Internet Zone could have higher security than domain names that you have decided to trust, like your email provider, your bank, and so forth. So Internet Explorer essentially is dynamically varying its security based on where you go on the web, which is really – it's a great concept. I love the idea because it allows someone to really bolt their security down, and then selectively trust sites where they're going to lessen their security in order to get the higher level of functionality which many sites increasingly are requiring, like having scripting running.

Anyway, so the question was, if you had a page that was a hybrid, the page was coming from someone you had told Internet Explorer to trust, yet it had components which it was fetching from untrusted sites, does that wrapper of trust extend to anyone that that web provider or the web server or the page trusts? And the answer is no. Internet Explorer looks at every single request individually and will restrict, will properly restrict the security of any components that you haven't told it to trust.

**Leo:** There are third-party programs out there that extend Internet Explorer, make it easier to lock down Internet Explorer. In fact, there's one, and I'm trying to remember the name of it, that locks down your entire computer, basically says – you can run as administrator, but it basically says that you don't have permission to do anything unless you specifically open it up. And I just can't remember the name of the program, so I'll have to leave that for the show notes or something like that. But you just do it by hand, and it seems to work very well.

**Steve:** Right.

**Leo:** Yeah. Well, that's it. Twelve down. Thank you, Steve.

**Steve:** Absolutely.

**Leo:** You're an amazing fella. If there are people who listen, and I know that there are many security pros who listen who are concerned with security, they know about Astaro. I don't need to tell them Astaro is one of the best known security firms out there, because,

of course, of their Astaro Security Gateway software. Really fantastic stuff. Astaro is great for your medium or small business network. If you need protection from spam, from viruses, from hackers, of course you get complete VPN capabilities, intrusion protection, as we were talking about, but the real deal, content filtering, and of course an industrial-strength firewall in a simple, easy-to-use high-performance appliance, Astaro.com. It's open source, too. I really like that. Or call 877-4AS-TARO to schedule a free trial of an Astaro Security Gateway appliance in your business. And I also want to point people to the really cool Astaro Command Center v1, which is now free for users of the ASG, that really gives you a neat way of monitoring your entire network. I use a 120, and I'm very happy with it. So when you're talking security, the name Astaro absolutely should come to mind. And we thank them for their support of Security Now!.

**Steve:** Yeah.

**Leo:** A reminder that Steve's site, GRC.com, is the place to go for, as he said, if I screw up, 16KB versions of the show, or anybody who wants an easy, bandwidth-friendly version of Security Now!. Transcriptions, too, thanks to our great transcriptionist, Elaine, at GRC.com/securitynow. That's where you'll find ShieldsUP, of course, Steve's great security software. He does a lot of free, pro bono stuff. And the bread and butter for Steve Gibson, which is SpinRite, the fantastic file recovery and maintenance utility. Everybody, if you've got a hard drive, you should have SpinRite. I mean, I'm serious. It has saved my bacon many a time. And if you want to see some great testimonials, go to SpinRite.info. And Steve's always putting up new letters there.

Ah, let's see. Anything else we want to cover? I think we're good for the day. Next week...

**Steve:** Good for the week.

**Leo:** Next week we're going to – well, maybe not next week.  But the next couple of weeks we're going to talk about virtualization.

**Steve:** Yup.

**Leo:** And that will be fun.

**Steve:** Virtual machine technology. But I want to talk about monitoring communications traffic, how end-users who want to kind of get a sense for what's going on in their computer can quickly make that determination.

**Leo:** Ooh.

**Steve:** And that's what we'll do next week.

**Leo:** Next week. Episode 49. Thank you for being here, Steve Gibson. Thank you all for listening. And remember, take a look at the new site, TWiT.tv, and those donation buttons are still there in a beautiful lime green.