Listener Feedback Q&A #13

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-068.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-068-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 68 for November 30, 2006: Q&A #[13].

Security Now! is brought to you by Astaro, makers of the Astaro Security Gateway, on the web at www.astaro.com. And by Dell. For this week's specials, visit TWiT.tv/dell.

Time for Security Now! with Steve Gibson, where we talk about the latest security news and explain security issues. We talk about all kinds of things, everything from VPN to virtualization to encryption and security. And this time it's our Mod 4 episode. Hi, Steve.

Steve Gibson: Hi, Leo, good to see you. Or good to hear your voice.

Leo: I'm seeing you. I'm imagining you. I can see you in your Security Now! lair, surrounded by security devices.

Steve: Next month we will be together again in Toronto, and so we'll...

Leo: I think your last time in Toronto. You're not coming up in January, are you?

Steve: Exactly, it'll be my last time.

Leo: Yeah. Then we're moving to Vancouver. And we've got to find Steve a direct flight or he'll never come up.

Steve: Yeah, you know, I did look at even the direct LAX to Vancouver. And there's enough flying time that there isn't time for me to get up there – oh, and there aren't any super early flights. And so there's just no way for me to get up in time. But I've got some friends in Vancouver, so I can come up the night before and hang out.

Leo: There you go. Come visit your friends. There you go. We're your friends. Come visit your friends.

Steve: Well, I was listening to – I've really been listening to Windows Weekly with you and Paul because clearly Vista is very much on my mind; it's very much on the minds of our listeners.

Leo: Well, and today is the launch date for business. In fact, I'm doing a bunch of interviews. I'll be doing "All Things Considered" this afternoon. Everybody wants to cover it right now. Even though in fact it's not available for months for consumers.

Steve: Well, I know. It's sort of interesting and frustrating. I got email from someone who has, through his corporate volume purchase agreement, they just got access to it. And it's funny because he was explaining that his experience exactly echoed mine. He initially set up Vista on a machine, and it was just, like, doggy slow. I mean, it was really lagging. And his graphics had a 2.0 rating. It turns out that he updated to the newest NVIDIA beta drivers. And his graphics performance jumped to 3.6. He got whatever the heck that – what is that translucency called? Is it...

Leo: Aero Glass.

Steve: Aero Glass, yeah. Which to my mind, you know, I'll be really interested to see what people think after they've had an experience with it.

Leo: It's pretty.

Steve: It just bugs me. It's like my contact lenses are fogged up or something.

Leo: You're a Windows 2000 guy.

Steve: Yes, I am. But I'll tell you, Vista really looks nice. And he also turned off the glass and then saw, as he said, and this has been my experience, it performs every bit as well as XP does. I don't see any...

Leo: And there are some enhancements that will, I think, make it faster, things like this hybrid drive support.

Steve: Well, it's funny you mention that because that's where I was going with this. I wanted to mention that the hybrid drives are not for performance. They are for battery life.

Leo: Well, but there is some performance benefit. For instance, it wakes up faster, things like that.

Steve: Yes. So there are some benefits from caching. But the reason it's on on small drives first and on laptop drives first, as you and Paul were mentioning in last Friday's Windows Weekly, is the whole idea is that laptop drives, or really any drives, really spend most of their time spinning and doing nothing. But that's burning power. So what the hybrid drive, the whole concept of the hybrid drive is that the spindle will be stopped most of the time; and that as the OS writes things to the drive, it fills up this non-volatile buffer with stuff eventually to be written. And only as that buffer gets near full does the drive then autonomously say, whoops, looks like I've got to get this buffer cleared. So it spins the drive up. It then flushes this non-volatile cache from inside itself, which is where this non-volatile cache is, out to the actual magnetic media. It then erases the cache and shuts the spindle down again. So it completely changes the model of how a laptop drive works. So that instead of mostly spinning, it's rarely spinning. And the whole motivation for this was battery life on portable drives.

Leo: But wouldn't there be some performance benefit because of this cache at all?

Steve: Well, you mentioned at some point, I don't know if it was last Friday or before, that EEPROM technology does not write very fast.

Leo: And that was my concern, is how are we speeding things up by using flash memory?

Steve: Exactly. And in fact we're really not that much. I mean, this USB boost, that's the technology which can be used to enhance the booting of XP. And the reason that the hybrid drive can't be used that way is its buffer is all about not requiring the spindle to be spinning all the time. So it isn't actually saving quick boot stuff. That's what the USB boosts do.

Leo: So you might want to use both.

Steve: I think you probably will. And in fact I'll be experimenting with it here shortly so I can begin – and doing some benchmarks to tell people, you know, like how this thing performs either way.

Leo: Now, you can't get a hybrid drive yet, so it's moot. But can you do the USB boot, the ReadyBoost, right now?

Steve: Yes, all that technology is available.

Leo: Okay. I just ordered a laptop from Dell for Vista, an M1210, that has a good video card, lots of RAM, core 2 processors. And so I'll try that. I guess we have to find some thumb drives that are ReadyBoost ready, I guess. I don't know if they're selling such things yet.

Steve: I haven't looked at it, so I don't know one way or the other. My guess would be that that's not necessary because...

Leo: Any relatively fast USB would be fine.

Steve: I would think so. And then the idea would be that XP would write to it a bunch of stuff that it knows it's going to need when it's booting, that enough of XP would boot to get the USB drivers going, at which point it would suck in everything from the USB that it would normally get from the disk. But again, Leo, yes, everyone's annoyed that Windows takes so long to start up. But I want people to understand that, from my perspective, Vista boots as quickly as XP. It runs as quickly as XP. I want to really make sure I back out fully from that first perception I had, which was entirely due to the fact that I was using debugging version and the prerelease version. And it was clearly not happy with my hardware. Whereas the final one really got happy.

The other thing is that I wanted to mention I also installed it on, I think I said last week, on a tablet. And one of the things, I mean, they've done some things for the tablet which is so cool. But when it first installed – and this relates to something that Paul mentioned last Windows Weekly – when it first installed it didn't have the Wacom drivers for the tablet technology. It went out to Windows Update, got the drivers, and suddenly enabled all of the tablet stuff. And the cursor, when you're using the tablet, since you're using a stylus you really don't need a cursor under the stylus. So it's always sort of annoyed me a little bit that I'm dragging this standard Windows cursor around when I already know where I'm pointing. They changed that so that the stylus with the tablet is this little tiny sort of twinkly star, I mean, very subtle. And when you tap the surface of the tablet, this little ripple comes out. It's just gorgeous.

Leo: That's Aero Glass, Steve. It is, because it's alpha transparencies. That's got to be Aero Glass.

Steve: Except that it doesn't give me the glass UI because I've got a relatively low performance graphics. So this is just an animated cursor. I think it just uses cursor animation.

Leo: Oh, it's not rippling the desktop.

Steve: No, no. But, I mean, it's just a little ping sort of, like, echoes outwards from where I tap the stylus. I mean, it's just little touches like that that make me think, well, I'll be running Vista on my tablet here as soon as it's fully supported.

Leo: Are you talking tablet PC, or are you talking a Wacom tablet on a regular PC?

Steve: In this case it's one of the HP Compaq – it's the same tablet that Jen is using because I really was admiring it when I was up at Call For Help, so I got one. HP doesn't sell them anymore, but you can get them on eBay. It was the TC1100.

Leo: Yeah, I had a 4200, actually.

Steve: I really like it.

Leo: So they haven't yet shipped the Windows Vista Tablet Edition.

Steve: Correct.

Leo: Or you just don't have a copy of it.

Steve: Well, it's interesting. I don't have that specifically, but there is all of the apparent tablet support is in Vista.

Leo: It's in there. Okay.

Steve: I'm guessing that there is no more a tablet edition of Vista; but in fact that Vista will just have that, and it'll work when it's on a tablet.

Leo: Very good. I'm looking forward to it, I have to say. The jury's still out, as you said, on security and reliability and so forth. And it's not really a target yet. Hackers haven't had a chance to really bang on it. But it can't be any worse. Well, maybe it could. I take that back after hearing the Vista Virgin Stack episode. It could be worse.

Steve: I think there will be some problems. We're still suffering stack overflow sorts of buffer overrun stuff. I don't see anything that's going to fix that until we get the data execution prevention, the DEP technology, ubiquitously present in users' machines.

Leo: But any new machine will support that. And I would hope Vista has that turned on by default.

Steve: Well, that was exactly what I was going to say. All of the lessons we have learned have told us that technology that's not on by default might as well not be there at all. It has to be on, like Service Pack 2's firewall. That's what finally shut down all of the worms. And so Vista will be better than XP was in the beginning because it'll have the firewall turned on by default. But the question will be, does it have DEP turned on by default? You could certainly imagine that the 64-bit version may, although that's user-mode protection, and that's very different than kernel-mode protection. Actually we've got a couple questions in today's Q&A that are going to be addressing that. So we'll be covering that.

The last thing I wanted to mention was something's just happened that I thought was very significant, and that is that SANS, the SANS Internet Storm Center, the ISC honeypot has found recently, out of 12 malware specimens, three of those were virtual machine aware.

Leo: Oh, boy. You mean Blue Pill-style attacks?

Steve: Well, no. What it was, this was malware that was deliberately detecting if it was being run in a virtual machine and being benign, if so.

Leo: Oh, interesting.

Steve: Yes. It would see that it was in a VM, and it would not do its bad stuff.

Leo: Why not?

Steve: Because they want to thwart researchers who are trying to use virtual machines as containers...

Leo: As do most of these honeypots. In fact, we learned this in Episode 1 of Security Now!.

Steve: Exactly. Exactly. So the idea was that, you know, the malware will not misbehave if it's in a virtual machine container so that researchers have a problem. And the other insidious thing is that, if end users were using, like if security-aware, high-end tech users were using a virtual machine to test software for its containing anything malicious, and they were putting it in a virtual machine as a test environment, deciding that, oh, look, this is not doing anything bad, it hasn't hooked the kernel, it hasn't modified any system files, blah blah blah, then they take it out of that container to install it on their native machine platform.

Leo: And then, boom. Gotcha.

Steve: Exactly. Only then would the malware launch itself and hook itself into their system.

Leo: These guys are amazing, these guys writing this stuff. I may not want to give them a lot of credit, but they really do work hard to keep up with what's going on.

Steve: Well, and this is the nature of any sort of a cat-and-mouse scenario where we don't really have robust protection. All we've got is the good guys battling the bad guys in a more or less level playing ground. And when the playing ground or playing field is level, and the bad guys are able to do as much, they have the same resources as the good guys, I mean, that's always been the virus/antivirus, spyware/antispyware battle is this back-and-forth thing. And so it's the reason, for example, that I salute Microsoft for being so determined not to allow the kernel to be modified under the 64-bit version of XP and Vista, and because this will eventually, even though as we saw last week that the Windows kernel protection doesn't really solve the problem, it's a stepping stone to them introducing a hypervisor that will lock down the kernel eventually and finally really prevent rootkits from ever being able to get a hold. So, neat stuff.

Leo: Good. I'm wondering if I should – I guess I will try to install the 64-bit version of Vista first on this new laptop, just to see if it runs.

Steve: I'm definitely going to be doing experiments, too. I downloaded the 64-bit version. The ISO for the 32-bit version, I think it was 2.687 gigs. The 64 is 3.6 or something. So it's an extra gig of stuff. So it's a bigger ISO. Still fits on one DVD. It's a substantial download. But it's definitely something that needs to be taken a look at. And I want to see what drivers are available because right now Vista had all the drivers built in that I needed for the two machines that I've run it on, which are not new machines. And so why wouldn't the 64-bit version, since

all these drivers are coming from Microsoft, why wouldn't the 64-bit version also know about all the same hardware?

Leo: You have to have 64-bit drivers, as you pointed out. And I guess they haven't gotten around to writing all of those drivers yet.

Steve: Well, it's the third parties who have wacky hardware that isn't...

Leo: I guess if Microsoft has the support, if it's a Microsoft driver, why wouldn't they?

Steve: That's my point.

Leo: That's your point? Yeah.

Steve: Yeah. And Vista, as Paul said on Friday and as my experience has been, Vista already has known about a much wider range of hardware right out of the box than XP or certainly any earlier operating systems have.

Leo: You've got to get something for five years of work. I'm looking forward to trying it on this, I mean, it's a brand new Dell laptop. It should have all of the, I mean, shouldn't be anything weirdo in it. I mean, certainly Microsoft knows about every bit and piece of that laptop. So you'd think 64 would install right out of the box.

Steve: I have a feeling.

Leo: We'll see. I'll let you know. Shall we get to our questions? We've got a lot of them. This is, as I said, Episode 68 means it's divisible by four, and it's Question-and-Answer # [13]. We've got a dozen excellent questions from you, our esteemed listeners.

Oh, I want to say before we do that – we'll get to the questions in a second. But I do want to mention Dell, since I already did, give them a little plug and remind you that the laptop I just bought, the big desktop that we just bought for Call For Help, and I put one affordable system in there, too, are all available at the Leo's Picks Page, TWiT.tv/dell. If you're buying a new computer, looking for a computer for Vista, may I recommend Dell. We've had such good results with them over a decade now, using them on TV, where it really is a challenging environment, and we're constantly banging on these things, installing new stuff, really working it harder than you ever would at home. And they just stand up and work and work and work and work. TWiT.tv/dell. Any Dell you buy through that page, any Dell at all will benefit the TWiT.tv network. So we thank you and Dell for their support.

Now let us – do you want to quickly mention SpinRite before we go to these questions, or should we save that for later?

Steve: We'll talk about it next week, actually. I've got another couple cool stories about SpinRite for next week.

Leo: All right. Because I do want to make sure everybody knows about SpinRite. And you have some neat new software coming up on GRC.com, too. We'll talk about that when it's ready. I know you're working on it. I've seen a picture.

Steve: I mentioned it last week. I wasn't sure I was going to do it. Now I'm in the middle of it. And so next week we will be announcing the first piece of security software to come directly from the content in Security Now!. It's been my talking about this stuff in the last couple weeks that I thought, you know, I've got an idea for something I want to write. So it's a new piece of GRC freeware we will announce, and it'll be available next week.

Leo: Good. Just in time for Christmas. GRC.com. All right. Question 1 from J. Sisco of Seattle, Washington: You were talking about the BitLocker technology in the forthcoming Windows Vista, its full disk encryption. I work at a company that's considering rolling out the technology to all of their laptop and desktop machines. Good idea on the laptops, I think. You see people like the Veterans Administration losing laptops with private data on there. He says: Since you're intimately aware of how a hard drive functions for data, what are the reliability implications of using full disk encryption? For example, if a non-encrypted hard drive were suddenly to get a bad sector, well, only the files sitting in that particular sector would be unreadable. Does full disk encryption fail if there are errors on reading a sector? In other words, am I going to lose a whole lot? Does the entire volume fail over one bad sector? And how would data recovery tools deal with the randomized data? Very good question.

Steve: Really good question. First of all, the good news is the encryption is only at the sector or cluster level, depending upon what type of encryption you're using. So the granularity of loss is no larger for encrypted drives than it is for non-encrypted drives. As for data recovery software, I can't speak to any other applications, but I know that SpinRite will work just fine. It won't care that it's being asked to recover the data and the sectors on an encrypted volume because SpinRite deals with the drive at the hardware level, underneath the encryption, underneath the file system. So it will recover the data on a sector, even one that it can't understand. It still does data recovery.

Leo: And I would imagine that even a recovery system that works at the file system level, if what you're saying is true, that the encryption is sector by sector, so the FAT 32 or whatever the file allocation table, that's going to be the same, right, that's still going to work.

Steve: Well, except that the software would have to be viewing through the encryption. So you'd have to have something running first...

Leo: So it couldn't recover an encrypted file. It couldn't say, oh, well, here's the FAT, and here's where the things are, let's recover those sectors.

Steve: Again, it's impossible to be accurate. You'd have to know exactly how the software was working, and what the nature of the damage was, to know whether it was going to be able to help you.

Leo: Okay. But a qualified yes, it's okay.

Steve: Yes. And certainly in the case of SpinRite, I know how SpinRite works, and it will have no problem at all. But definitely there is no greater loss potential. It's not like you lose, as this guy asks, the entire encrypted volume if you do lose a sector, any more than if it were non-encrypted.

Leo: On something like TrueCrypt, where it's making a big blob of encrypted data, that's risky, yes?

Steve: Well, except that internally it's going to – individual file sectors are being encrypted standalone. I guess the point is no sector depends upon any other sector. And so there's no...

Leo: I see what you're saying.

Steve: ...blanket problem that is being increased by just encrypting individual sectors because each one stands alone.

Leo: You may not know where a file begins and ends, but each sector stands alone. So if you lose a sector, you've only lost whatever that sector belongs to, not everything else.

Steve: Right.

Leo: No interdependency. Okay, that makes sense. Joe Rodricks in Massachusetts asks: I'm over my head here, but I'm curious. Why can't a kernel be more like a man? Oh, no, no. Why can't a kernel – Colonel Pickering. Why can't a kernel – this is inspired by the Vista kernel lockdown talk we had earlier. Why can't a kernel have an MD5 or other hash on a non-writeable medium, say a burned CD of a kernel hash that could be installed so at least the kernel can be aware it's been modified? Oh, that's interesting. So you've got an MD5 checksum on it. You've put it somewhere that can't be modified. Then you could validate whether that kernel's been changed or not? He says: I'm sure something along these lines could be helpful. Is he right?

Steve: No.

Leo: And that's because...

Steve: But it's a great question, which is why it's here. It's because the kernel is the thing that would be doing the testing of itself.

Leo: Right. You'd have to use an external operating system to check the kernel.

Steve: Well, exactly.

Leo: You could do that, I guess.

Steve: Well, and we talked a long time ago about the idea of using an external, in the case of rootkit detection, using something outside to look in and check for modifications that a rootkit had made to the OS.

Leo: That's how RootkitRevealer works. It's its own operating system, essentially.

Steve: Exactly. So the problem here is that if the kernel could be – and the reason I wanted to bring this up is that it makes it more clear, I think, that nothing can protect itself at the same level of capability. Malware and antimalware are fighting on a level playing field. Rootkits and kernels are fighting in the kernel space on a level playing field. Anything the kernel can do, a rootkit running at the same privilege level can undo. So that the first thing a rootkit might do is neuter the kernel's check of itself. And there's no way the kernel could know that it had been neutered because it could be prevented from knowing that by the software which has been designed to prevent it. I mean, it's literally, it's just not possible.

The only way we're going to get this problem resolved is eventually when we have a hypervisor running at a level above the kernel, or below, depending upon how you want to look at it, you know, more protected than the kernel, which is then able to enforce protection on the kernel. This is very much the way the kernel is able to enforce protection on user applications. And so it's sort of a hierarchy. And at the moment, the current Windows technology doesn't have anything running at a level that is more protected than the kernel. That's what we need. And only then will we actually have anything that can be verifiably and usefully bulletproof.

And this is exactly why, as we said last week, kernel patch protection, while a good thing, doesn't actually provide rootkit protection because rootkits already exist that are able to defeat it. The idea is, this prevents, as we said, good guys from modifying the kernel, which will finally allow Microsoft the flexibility of locking down the kernel at some point in the future. They just haven't gotten there yet.

Leo: From Sweden, Simon Lingh is worried about the future of third-party firewalls and A/V. He writes: It seems like Microsoft wanted to prevent all third-party kernel modification; which means, as I understand it, there won't be any, for example, personal firewalls in the future. Or did I miss something?

Steve: Great question. It's certainly the case, as we were just saying, that Microsoft is determined on the 64-bit platforms to draw a line in the sand and say, no more kernel modification. What they have done to compensate – and I believe this fully, that they really don't intend to prevent Symantec and McAfee and the other personal firewall and A/V vendors from being able to develop products. They just need to develop products that don't modify the kernel. To make that more possible – and first of all, it has been possible, so they say. I mean, there's a lot of dispute about that. I'm on the side of believing that it has not been possible to do the sorts of things McAfee, Symantec, ZoneAlarm and so forth are doing without reaching down and modifying the kernel because Microsoft has non-published APIs and just hasn't provided the hooks to allow that. Microsoft was – the operating system was hostile to things mucking around in the kernel. So there's a huge amount of reverse engineering that was needed to be done in order to make this possible.

What Microsoft has done with the next generation of Windows, that is, in Vista, they have published APIs that will allow third-party vendors to do what they want to do without needing to modify the kernel. There are some advanced API technologies, for example, one that allows software drivers in the kernel to filter and monitor all of the network traffic coming and going through the machine without needing to modify the network stack. That's never been possible before. It is now possible in Vista because Microsoft has added that functionality. So there are now ways in Vista for third parties to successfully do what they want to do without needing the

kernel to be modified.

Leo: And I guess it also comes down to the question of should you let Microsoft and only Microsoft control what goes on at the kernel level. I mean, I think about Linux, for instance, where the kernel is open source. It's published. There's no kernel protection. You can do anything you want. I mean, you can't modify a running systems kernel necessarily, but you could write your own kernel with all sorts of hacks in it.

Steve: Sure.

Leo: Which is more secure? The closed, only Microsoft has access in here, clean room environment; or the open, everybody understands what's going on environment? Right now it's the open environment.

Steve: Certainly we need to remember that rootkits were invented on UNIX.

Leo: That's true.

Steve: That's where they came from.

Leo: Good point. It's actually fairly easy to put a rootkit on a Linux system, come to think of it.

Steve: Well, exactly. And the documentation, the open source-ness of it makes it that much easier. So at the same time we have the issue of doesn't having the open source mean that lots of people can look at the code and scrutinize it and work on ways to tighten it up more than Microsoft has with their closed source approach?

Leo: The difference is we're not at the application level, or even the kind of the OS user interface level. We're at the kernel level.

Steve: Right. And there are two different issues that we're sort of commingling here. There's the issue of open source/closed source and hypervisor kernel protection versus the PatchGuard approach, which is just trying to detect modifications to the kernel, as opposed to absolutely preventing modifications to the kernel. So it gets complex. I believe that when Microsoft for their 64-bit systems has added a hypervisor, we're going to see a new era of security because applications will not be able to muck with the kernel at all.

Now, it is the case, though, that mischief could be performed by using the new APIs, that is, if now you don't need to modify the kernel, yet you can hook the stack, and anybody who wants to can hook the stack, then nothing prevents malware from creating a driver which they arrange to sign somehow. And driver signing is another thing that's been enforced in the 64-bit kernel. So malware gets a signed driver that hooks the stack. Now it can talk out through the network and potentially bypass the firewall.

Leo: I guess it's what you've been saying all along, and I guess I've learned this lesson, is

there is no such thing as perfect security. It's a process. And there's always something.

Steve: Yes. And doing the best job we can of keeping this stuff out of our machines in the first place is absolutely what you want to shoot for. You don't want to let stuff get into your system. Once it has, you just can't trust it.

Leo: A subject we don't talk a lot about, spam, is on Paul's mind. He's fed up with spam. He writes from the U.K.: Spam, spam, spam, spam, spam. Security Now! is preaching to the converted on firewalls, bots, spyware, et cetera. I doubt any of Security Now's listeners' PCs are sources of spam. I hope not. The fact is that there are millions of PCs out there that are vulnerable to being sources of spam, and there always will be because they don't listen to Security Now!. Leo said recently that his email filters catch a million spam emails a month. True. As effective as spam filtering might be, the long-term solution to this problem has got to be applied at source, otherwise the pipes will stay clogged. The solution has got to involve security. Hmm, I'm not sure I'd agree with him on this. I would like to hear your views on how it could be done. Should SMTP, the mail transport protocol, be replaced over time? Should we adopt the proprietary email authentication as proposed by Microsoft? There are others, by the way, that are not proprietary. Should LSPs limit users to 1,000 messages a month? Should all email be certificated? Or has the horse already bolted through the open stable door? He asks a lot of questions.

Steve: Well, yeah. And I liked it because he sort of painted a picture of the way things are today.

Leo: Bill Gates thought, and he announced, that spam would be conquered. He announced this two years ago, and it's long past the deadline. But he thought authentication was a solution.

Steve: Well, and I actually am a fan of that. In that famous TWiT that I participated in a long time ago which generated Dvorak's "I get no spam" comment, where he was complaining that my email gateway was rejecting his email because it used HTML. And what I had found was that by blocking HTML email, 99 percent of the spam GRC was receiving was blocked. And I was bouncing a message back saying, I'm sorry, we don't accept HTML email, please send text-only email, we'd love to hear from you. And Dvorak style, it made for some great comedy on TWiT.

Leo: I'm wondering if that still works, by the way, your HTML filter. Because a lot of my spam now is just embedded GIFs. It's not HTML, it's just a picture.

Steve: Yeah. Well, what I did, as you know, Leo, is I finally gave up my original email address. I was just steve@grc.com for, well, forever, for 20 years. And as happens, and you know this because your email address is well known publicly, and what I have seen, because I've looked at remote servers connecting to my server, I've watched the traffic, I see them just guessing names, people who never had accounts at GRC ever.

Leo: Oh, yeah, most of my spam is that.

Steve: Doing dictionary attacks, just going through a list of first names and second names,

everything that they can come up with, trying to get something, to find a legitimate account. What I did was I changed my email address, finally after all that time, basically after giving up, to something not in any dictionary that is no longer guessable. And bang, problem solved. So I then removed the HTML filter. Now people can send anything that they want to to us, as long as they have our addresses. And this was important also for sales and support because I wanted people...

Leo: You don't want to block customers' emails.

Steve: Exactly.

Leo: Now, I have to say there is an issue because they have to know what your new address is. You change it regularly. I don't want to say too much because I don't want to give away your algorithm.

Steve: Right. And we do have the addresses for sales and support are on our web page. And the current address is in every receipt that they receive after they purchase SpinRite, blah blah blah. So we basically closed the loop so that nobody is inconvenienced and dodged all of this flood of spam.

Leo: But nobody can guess your email address.

Steve: Correct. And so the idea is...

Leo: And I can't do this because – I can't do it.

Steve: You're too dependent upon that. And it is painful to change one's email address after you've had it for a long time. I really understand that.

Leo: Everything we've talked about so far really is filtering. Even that is a form of filtering.

Steve: Correct. So the first thing I would tell people is, if you finally give up – and I talk to people more and more who have given up trying to fight this. They're like me, who were really annoyed that this was happening, and just filtering it. And basically, Leo, that's exactly what you're doing is using antispam filters to stem the flood and try to find the good email from the bad. And in fact you and I had a problem communicating because yours was false positiving on email from me, so you had to explicitly allow me.

Leo: You were getting through the primary server-side filtering. It was the secondary, actually it was the secondary filtering that was catching you, Spam Assassin. I use three-step filtering. But he makes a good point, which is that doesn't solve the root cause.

Steve: Correct. So the one thing I wanted to say was that if anyone finally does give up and change email addresses, do something that is not going to be in a dictionary. Because if you just use your same, like your first name or whatever you use that might be in a dictionary, could be guessed, you'll find yourself immediately getting more spam.

Leo: If you're joe@earthlink.net, you're screwed.

Steve: Basically, yes.

Leo: Give up. That is not – joe537499 would be fine.

Steve: Right. It's very cool, Joe, that you got that name at EarthLink. But unfortunately it's pretty much useless to you now.

Leo: In fact, I think I got some email from – it was joe@aol.com, saying, yeah, this is my email address, but I can't use it. So please don't try to send me mail at this address. The reason I got it was because a spam message was bounced back to him from, I guess, my address.

Steve: Well, now, I wanted to address the second part of this, and that is this issue of authentication. The technology, which I really like, and I'm using it, and it's becoming more and more pervasive, is this notion of authenticating the source of email. It's this SPF technology, which is Microsoft tried to sort of adopt it and then, what was that old expression, "extend and enhance" or something, which basically means...

Leo: Embrace and extend was the original idea, but then extend and devour is what we ended up calling it.

Steve: Exactly. And they were unable to because it turns out that this was bigger than they were.

Leo: Unembraceable and unextendable.

Steve: Well, they wanted to patent it and license it, and people said no, we're not going to adopt some Microsoft thing that is not in the open domain. The good news is that this Sender Policy Framework, SPF, is a really cool and simple addition to an existing system because all it requires is that some special DNS records be added to the DNS server of anyone who wants to generate email.

So, for example, GRC has a text record which has been added to our DNS server that says the addresses of the servers, either by domain name or IP address or IP range, that are valid originators of GRC mail. So any incoming mail to a third party, like say EarthLink, that is using SPF, Yahoo! is, Gmail is. So some email comes in, apparently from me at GRC.com, for example, to EarthLink server. It performs a DNS query asking for the text records of GRC.com. Our DNS, over which we have absolute control, our DNS sends back these text records, among which is one that says GRC email will only be originated from these IP addresses. And what happens is EarthLink knows the IP address from – it has a connection to. And as we know, IP addresses of valid TCP connections cannot be spoofed because you have to have traffic succeed in that three-way handshake in order to get a TCP connection. And SMTP, which is the transport for email, uses TCP as its transport. So EarthLink absolutely knows the IP address of the server that's trying to send it mail ostensibly from GRC. It does a DNS query to see whether that IP address we are saying, that is, GRC is saying this is a valid source of GRC email. That allows it to authenticate my GRC server as being the only source of GRC mail. And if it works, it allows

it. If not, it drops it.

And it really is a great solution. It prevents these spambots running in botnets from originating email from other sources. Now, unfortunately this hasn't been adopted widely yet. But it is picking up and spreading because it is a simple thing to do.

Leo: You said that Hotmail does it, EarthLink. Does AOL do it? I don't think they do. As I remember, yeah, they had their own form. That's the problem.

Steve: There is a domain keys technology, which is another approach. In fact, I think that's the one that Yahoo! adopted was domain keys. And I don't know for sure whether Yahoo! is also doing SPF. There's nothing to prevent servers from doing both. And it just checks to see which ones you're using. I think Gmail is doing both. Because I know that I've looked at my headers, and I've seen that GRC server has been authenticated by Gmail. So anyway, there are things moving forward that just take time to adopt which ultimately are going to help, although it's not clear that anything is a complete solution. And boy, doing something like dropping SMTP and switching to a different protocol, we're talking about obsoleting the entire email structure of the Internet in order to do that.

Leo: I have to think basically spam is like cockroaches. We've tried to fight them for a long time, roughly 40,000 years, with no success. And I think it's going to take us that long, if not longer.

Steve: Yeah. I've got to say, changing your address and protecting it...

Leo: It works for a year, though; right? You change it regularly. Because it doesn't work forever.

Steve: I do change it regularly.

Leo: Because after a while you start getting spam.

Steve: Yup. But I just talked to Mark Thompson, my buddy at AnalogX this morning, who said that he's starting to get some on his current address, and he's looking forward to making a change next month, essentially, and switching over to a new address next year.

Leo: He does it about every year.

Steve: He does the same thing I do.

Leo: Yeah. So that's interesting. Have you started seeing spam? Because we're getting towards the end of the year.

Steve: No, I haven't. Although I'm very careful, as you know, because I've said, Leo, you know, please be careful...

Leo: Oh, I don't give out your address to anybody.

Steve: ...about letting this thing escape.

Leo: Yeah. No, I don't give your address to anybody. I always ask permission. I have to say that, to be honest with you, I think there are larger problems with email. Even with all the filtering I do, I still get too many emails to respond to. It's rapidly getting to the level where I get too many emails to read without spending hours a day on email. It's just too easy to send an email. When you're a public figure, you get too much. And I'm not – Amber and I were talking about this. We're public enough to get a lot of mail, but not successful enough to hire somebody to do it. So we're stuck. We're in that middle thing. I'm sure President Bush has plenty of people answering his email.

Steve: Well, and speaking of that, I love the fact that these questions that we're discussing are from Security Now! listeners. And they go to the Security Now! page. Down at the bottom of the page that allows them to submit questions.

Leo: But it doesn't come in your email.

Steve: It doesn't come in my email. Well, no, actually it is a form submitted. It is emailed to me. But the problem is the same as you, Leo. I mean, I love that I get these. But I get so many of them that, for example, if I were to read them, let alone answer them, I could not be developing this new freeware that I'm excited about that I'll be announcing next week. So like you I've had to make that tradeoff. So I really appreciate people writing. I just have to apologize if we're not able – if I'm not able to respond, and we're not even able to answer your question because there's just so much.

Leo: And I say the same thing. And it's sad because I get a lot – I love my email. I get a lot of great stuff. And I do so far still read it all. But I just can't respond to it all. And I think the time will rapidly come where I can't even read it all. It's so easy to send email. Anybody can do it. And it's just a push of a button. And if you get a thousand emails a day, you're just not going to read it all. And so I think ultimately the problem may not be spam. The problem may be email. It's a larger problem here. But let's not get philosophical.

Steve: Well, basically we're talking about developing larger communities of communicating users. And I've got to say that the newsgroup solution that GRC has is fantastic because we have a Security Now! newsgroup at GRC, a forum there. And even that I can't keep up with. But I dip in from time to time and see what's going on. And here we've got our Security Now! listeners talking to each other, answering each other's questions and discussing things. And it's fantastic to be able to create that kind of an online community.

Leo: Absolutely. Greg Rudd in Seattle gets it. He writes: If a computer becomes infected with malware, is running an antispyware program installed on the infected machine a reliable method to delete the malware? How can any program running on an infected machine be trusted? That's kind of the bottom line, isn't it.

Steve: And that's why I said that he "gets it." It is the golden rule that, if your machine is ever compromised, you can never trust it again because we don't know what has happened.

Leo: Increasingly I've been telling people that on Call For Help. In the past we would talk about how to disinfect, how to get rid of viruses. But these things have gotten so nasty, so adept at burrowing into your system, that...

Steve: Oh, and tenacious. You just can't take them out. You can't delete the files because then other things break because they've inserted themselves as a filter in between some other things. And you take that chunk out in the middle, and you break the communication flow.

Leo: You're killing the patient.

Steve: So, yes, if all you can do is run a malware remover, then that's what you have to do. The better solution is to think about the structure of your machine. The best thing you can do is to set up a C partition which is big enough for the OS and your applications. And that only needs to be maybe, you know, maybe 15 gigs, depending upon what you're doing. Then have a D partition or, you know, D, E, F, G. It's funny, I run across people sometimes from the old DOS days who've got the whole alphabet covered. "On my Q drive I've got this."

Leo: Removable drives I always make be W, X, Y, Z, like CD-ROMs and so forth, because then if I install something it won't screw up anything I've installed on the CD-ROM. It's going to stay at the end of the alphabet.

Steve: Well, the advantage to me of keeping my system partitions small is that it can be imaged relatively practically. That is, I'm able to do an image onto a couple DVDs or onto an external spinning Firewire or USB drive. And so my point is that, if malware infects you, it's infecting not your entire drive, but it's infecting your system partition. It's not like you're going to have...

Leo: It wouldn't spread to a data partition? Sometimes they do, viruses will install themselves into...

Steve: Well, and we've seen instances, for example, where JPGs are carrying malware. But you have to ask yourself, what are people doing with 250GB drives doing with the drive? That can't be 250GB of programs. Certainly most of it is music and movies and media files that are massive. So my point is that, if you were to deliberately keep your system partitions small, only as big as it needs to be, I mean, I've got a 17GB partition on my main workstation that is still only half full. And this is everything I use is in half of a 17GB partition. I don't even know where I came up with 17, but it made sense once.

Leo: Magic number.

Steve: And so my point is that I can make images. It's practical to make an image of my C drive because I kept it small. It's not practical to make an image of a 250GB C drive. Okay, you could do it by using an external drive, but it's going to take a long time, and you're imaging all your media, which is much more transient and probably doesn't need to be – well, basically it doesn't need to be saved in the event of a malware incursion.

Leo: Do you still use Drive Snapshot to do your images?

Steve: Yup.

Leo: I'm going to get that.

Steve: I love Drive Snapshot. It runs in Windows. You don't need to leave Windows. You're able just to run it, for example, against the C drive. It makes a snapshot. And what's cool is they're mountable snapshots. So you're able to mount that as a drive letter if you want to poke around and get something...

Leo: Just get a file or something.

Steve: Exactly, from a prior snapshot. And so my point is that, if you then got infected, if your data is separate from your operating system, you can restore your operating system from a recent snapshot and not lose the data that you've been working on since. So increasingly I think it really makes sense not to have everything as just one big blob in a C drive.

Leo: I think as I'm getting ready to create this Vista machine it's an opportunity to kind of sit down and think about these strategies. And I think that that's exactly what I'm going to do. I'm going to put it on the network, back up the partition – create a good Windows 64, Vista 64 partition, back it up as an image. And that way, if I should get spyware – I don't anticipate it, but if I get spyware or a virus infection, I could quickly restore to the way it was. Backing up is now becoming the best strategy of all. There is no interim, in-between thing to do.

Steve: It really is true. As I said, if all you can do is run antimalware, that's all you can do. But if you've got an image, it's much better to restore an image. And you just know, you have the peace of mind that you've stepped back in time to a point when your system was clean.

Leo: Best time to do it, when you get a new system. And then I guess, you know, this is what we always did on the TV show, we'd create a couple of images. First an image of just the installed OS. Then an image of the installed OS and applications. And then what is key is that you know that that's good because you haven't even gone online yet. And when you want to build a new image, perhaps with new applications or new drivers or, more likely, all the Windows updates, you actually go back to the known good image, restore that, add the applications at that time...

Steve: And then move forward again.

Leo: And then move forward from that point so that you're always moving from a known good spot. Otherwise...

Steve: That's exactly my approach, too, Leo.

Leo: Otherwise you could be restoring a compromised system. That's the point nowadays, you can't tell.

Steve: Well, and it's funny, too, because here we've only been talking about malware. But as we know, something about Windows just kind of goes a little wonky over time.

Leo: Yeah, it's a good thing to do anyway. I reinstall all the time.

Steve: Exactly. So backing up to a known good point and then installing a few more apps and then moving forward, or installing a few more apps and then making an updated snapshot, it does allow you to sort of reset your Windows system to like an almost new condition.

Leo: This is where a gigabit Ethernet and a big terabyte mass backup system is really helpful because anything that I put on my network I could snapshot out to that terabyte storage. And I have a number of good images. But they're not even on the system. But as long as I can get on the network, I can do it.

Eric in San Jose is in the market for a powerful new machine: I'm interested to know your opinion and maybe have some explanation on computer hardware. Oh, good, this is fun. I was looking for the most advanced systems on Dell and found only the servers have the option for dual quad core processors. I think he's looking at Xeons. It seems to me any enthusiast would want that. Is it better to buy a server to get the latest in performance? They definitely cost more. Is the difference between computers designed for home versus business versus servers in terms of performance and quality? I know the business-level tech support is better, but do these systems offer higher levels of performance and quality control as well? I can answer one thing. A lot of this comes down to Intel and the fact that they won't let you sell a Xeon-based computer as a computer, as a desktop. It's sold as a workstation or a server. That's just a branding thing. They tell Dell you can't call it a desktop computer.

Steve: Interesting.

Leo: So they're always at least a workstation.

Steve: Yeah, I've been looking at that, too, as a matter of fact, because as I think I mentioned, I'm still running Windows 2000 on an aging hardware platform.

Leo: Time for a Dell, dude.

Steve: I've got a pair of 866 MHz Pentium 2 or 3, I think Pentium 3.

Leo: Oh, dear. What are you using that for?

Steve: Well, the problem is it's so hard to move. I mean, I've got this fantastic mature environment.

Leo: Mature? Belongs in an old-age home.

Steve: Actually it does. And as a matter of fact it's a mixed blessing. I look at all the crud that I've installed on this thing over the years, and it's amazing it still runs.

Leo: I think it's time for a new system, Steve.

Steve: And so my point is that obviously I don't change hardware often because, I mean, this is proof of it because here I am running on an old dual P3 866. And so my point is, since I don't change hardware...

Leo: I have one of those for a doorstop, actually.

Steve: Right. So my point is, since I don't update my hardware often, when I do go through the pain of switching to a new platform, I want to reach far ahead, beyond what I need now, knowing that I'll probably be stuck with it for the next five or ten years.

Leo: That's kind of what I did when I bought the Mac Pro. What we're working on right now is a quad processor. It's dual Xeons.

Steve: Dual dual core.

Leo: Dual dual cores.

Steve: Right.

Leo: And I'll tell you, I feel like I'm reaching into the future. I can't get this thing more than 10 or 20 percent usage.

Steve: Well, that is where I'm going to go. I'm going to go to a dual dual core. I think dual quad core may be a little overkill.

Leo: I think that's one reason, to answer this guy specifically, that's an awful lot. That's what I have, but that's just because I'm silly. I think dual core is plenty.

Steve: What I have found, doing some research, and I was just looking at all the Intel motherboards available, is that you want to look at what kind of audio, what kind of video, how many of what type of PCI slots you've got, and then what the maximum amount of memory is and what type of processors and how many you're able to put on the motherboard. So those are sort of the things that, completely aside from branding and what you're going to call it, whether it's called a workstation or a server, blah blah blah, Intel does make high-end workstation motherboards that are dual processor, dual core, so you can go to the high end of the workstation or the low end of the server and sort of get basically what you want.

I like to run a lot of displays. I've got three in front of me. I will definitely go to a system that allows me to run three displays because I'm hooked on the idea of being able to have that much screen real estate. I can't imagine going back to a single display now for my main system. So I need to be able to have enough PCI slots to put graphics cards in, in order to run LCD panels. So that's a consideration for me. I don't care about audio for my workstation. It's not a media machine. And having integrated RAID is, I think, now that RAID is on motherboards to a greater and greater degree, just having it built onto the motherboard so you can run three drives in RAID 5 and just know that, no matter what happens, you're going to be okay, for any next-generation system I think at least running a mirror and maybe even a RAID 5 makes a lot of sense, too.

Leo: On the computer or as an external system?

Steve: Lots of motherboards build it in. I know that Intel has...

Leo: Let me ask you about this, though, because this actually is an important question to me. My impression of most of the motherboard RAID controllers is that they were software. They're BIOS-based RAID controllers, not hardware RAID controllers.

Steve: Yes, that's the case.

Leo: But RAID 5 would require hardware. Is that correct?

Steve: I know that Intel's motherboard says they support RAID 5.

Leo: See, I have some questions about how good the RAID 5 implementation would be.

Steve: Well, you're right. Although I have seen, when you boot FreeBSD, for example, or maybe it's Linux, I think it's Linux in fact, where it runs through a bunch of tests to determine what the optimal RAID strategy would be using different approaches for generating RAID checksums and RAID images. And the performance is really getting pretty far up there. And you need something for one of those extra cores to be doing.

Leo: True, you have enough processor now that you could – software RAID isn't such a bad thing. So RAID 0 is striping, which speeds up the writes and reads. Although, again, in our tests on BIOS-based software motherboard RAIDs, we haven't seen much improvement. There's RAID 1, which is mirroring, and that's for redundancy. I have some issues with that, frankly, too, because in my experience you shouldn't count on that as a backup by any means. And then RAID 5, which is using three drives; right? Don't you have to have three drives for RAID 5?

Steve: A minimum of three, yes. You're able to use more.

Leo: So you're doing both striping and mirroring. Or is it only mirroring?

Steve: It's actually neither. It's a very sophisticated approach where basically...

Leo: It's a redundancy approach, though.

Steve: It's a redundancy approach, but it's much better. Striping is zero redundancy. Mirroring is 100 percent redundancy. RAID 5 you get an extra drive of sort of checksum information. And so if any one of the three dies, the system is able to rebuilt the dead drive from the data in the other two.

Leo: You just swap a new drive in. And that's how my NAS storage works. It's three drives, RAID 5.

Steve: Now, in my case, for a personal workstation, I'm not at all concerned about drive performance because drive performance is already so far beyond what a workstation needs that it's just not a concern for me.

Leo: I'm not sure I'd agree with you, though, Steve. I think especially, well for instance on this quad core, I think you're I/O bound. You have so much processor performance that you are I/O bound now.

Steve: My point, though, is that for workstation, that is, for me sitting here doing email...

Leo: You could go back to that P3, and you'd be fine.

Steve: That's why I'm so happy right now with my P3.

Leo: But not everybody – but if you're doing anything demanding, editing video or ripping CDs or anything like that, hard drive performance does become important. I put a Raptor in almost everything I do, a 10,000 RPM drive.

Steve: Then I certainly agree. And I agree, then, that not only do you want to get off the motherboard so that you're not doing software-based RAID, but you probably want some caching on a separate RAID adapter so that the OS is able to dump data onto the RAID and move on, and then the RAID uses its on-card cache and writes stuff out to the RAID asynchronously. You really want that additional overlapping.

Leo: I believe I have that in my in-front terabyte NAS.

Steve: So did we answer Eric's question?

Leo: I don't remember what his question was. Let me go back and look. I have it written down here. He's just basically asking all about hardware. He wanted to know quad, dual quad core, which I think we're both saying is not necessary for most applications. You'd get more benefit from, say, a faster video card or a faster hard drive than you would from

having four cores.

Steve: I can't imagine anything other than a compute-bound, high-end server, or somebody who's doing tons of media compression.

Leo: Or 3D rendering or something like that.

Steve: Yeah, exactly, 3D rendering or media compression. I mean, I pin my processors, when I'm compressing media, and the hard drive just kind of flickers, flickers, flickers, flickers because all of the time is spent by the processor doing MPEG-2, MPEG-4 compression.

Leo: So you should come over here and use my Mac Pro because this thing is fast.

Steve: Or get one of my own.

Leo: Yeah, this thing is fast. On that kind of stuff it's as fast as I've ever seen.

Steve: That's very cool because it's a pain. My solution is having multiple machines just running in the background. It's funny, too, because I keep wanting to write – you know I have my little Window gizmo called Wizmo that is, like, all kinds of little features that I've added. I want to add to it the ability to monitor a system for how busy it is and to start ringing a bell, playing a media file as soon as its level of busyness changes. Because I'll start a compression on a computer and forget that it's running, and come back hours later when it's been done for a long time and wish that I had something that notified me when it was through.

Leo: I guess to answer his question most specifically, he's asking should I get a server or a computer. There are really three choices. There's computer, workstation, or server. Servers are really tuned for being servers. You don't need a server. If you want quad core, you can get a workstation. But I think Steve and I agree, even though I bought one, I don't think most people need a quad core workstation.

Steve: I completely agree. And Leo, talk about expense going exponential.

Leo: Yeah, it goes up pretty quick.

Steve: Oh, those things are just so expensive.

Leo: But, see, I was buying for the future. I shouldn't have to buy another computer for quite some time now.

Steve: That's the plan.

Leo: Which will be sad because that's part of the joy of life. So that's why I got this laptop; right? Matt Carroll of Milwaukee, Wisconsin wants a bit of clarification: In Episode 65 you discussed configuring your regular mail client, Outlook for instance, to access Gmail via POP and SMTP with an SSL connection. Does that mean, if my PC is connected to an unencrypted open WiFi connection, and I send or receive email using an email client configured to use SSL, that messages cannot be intercepted by someone sniffing the connection? In other words, is it safe to transmit email using Outlook or a similar client via an open WiFi connection if I'm using SSL?

Steve: Yes.

Leo: It is. I thought you were going to say no.

Steve: Absolutely.

Leo: Because sometimes it opens up the connection in SSL, but then reverts back to unencrypted; doesn't it?

Steve: Not POP and SMTP. It's exactly like using a secure website where the SSL connection is established first. It cannot be sniffed. It cannot be spoofed. And you then have an encrypted tunnel through which all of your email traffic goes.

Leo: We should point out that on web-based servers, however, this is not the case necessarily.

Steve: You need to be careful. With Gmail, as long as you go to it with https://mail.google.com, as long as your initial contact with the server is secure, after you log in you stay secure. Whereas if you go http://, that is, you first go to it nonsecure, then only your log-in is secure, and Gmail reverts to insecure afterwards.

Leo: Yeah, and I think people get fooled by that.

Steve: But in his case, yes, if you're using SSL with POP and SMTP, you can absolutely not be sniffed. Now, if you're using open WiFi unencrypted blah blah blah, of course there are other concerns that you would need to be aware of. But certainly at least your email would be safe.

Leo: Yes. Yes. In fact, that's what I do. I still use VPN, but I do do that at least; so if I forget to turn on the VPN or whatever, my email is safe. I mean, the first thing you want to protect is your email password. But then it's nice to protect your email, too.

Eric Sarratt of Asheville, North Carolina, catching up on back podcasts. He had a question about drive capacity: I know Steve mentions getting the smallest drives he can find. Is there an ideal size? I mean, you can get 500GB drives now. Is it better to get a 160 or a 100GB drive? He said: I can actually get drives down to 1 or 2GB. Obviously that's too small. By the way, I usually get the old PATA drives as anything faster is usually a waste of money for my computing needs. But if you recommend SCSI – ugh – or SATA, I would be

open to this, given a reason. Is PATA IDE?

Steve: Yes, PATA is the original IDE style. I thought this was an interesting question.

Leo: Parallel ATA, I get it, okay.

Steve: Exactly, parallel versus serial. I thought it was an interesting question because I did talk about deliberately asking my computer supplier guy, when I'm buying drives just for random use here, what's the smallest drive I can still get. These are applications, though, for example GRC's server, I think I'm running a pair of 40GB drives in a mirror on one of my servers. This is not the main GRC server. There I've got a SCSI RAID 5. But there are many applications where I just don't need – I just know I'm not going to have lots of media files, and I don't need that much space. So given that I don't need that much space in that application, I'll get the smallest drive I can, just because I like them better. I trust them more. But most end-users, I think, are now obviously going for the 160s and the 200s and so forth. You'll get a kick out of this, Leo. I just increased the 80GB drive in my MacBook Pro to 160.

Leo: Why?

Steve: Because I'm doing more with it now.

Leo: You're putting music on it and stuff.

Steve: And media stuff. And it turned out that 80 gigs was just – it was cramping me. So there's an example where Mr. Get the Smallest Drive You Can has just gotten the biggest drive he could.

Leo: So what's the optimum? That's the biggest laptop drive. What's the optimum – there is no magic number, is there?

Steve: No.

Leo: There's no point beyond which – your problem is aerial density gets so high, so much data on a square inch, that you're concerned about reliability.

Steve: Yes, in fact, the drive I bought, I bought a 160GB Seagate SATA, which is what goes in the MacBook Pro. And it's Seagate's first vertical recording drive, which is storing the bits standing up instead of laying down, which is technology we've been hearing about for a decade now, and it's finally beginning to make it out of the lab into production. What I would say the other concerns, other than just cost and size, is power consumption and heat. And this is really the other reason I wanted to put this question in here is that bigger, faster, larger drives do tend to generate more heat and burn more power. Certainly the case with your 10K RPM drives, Leo. I'm sure those guys are running a lot hotter.

Leo: Oh, yeah.

Steve: So you want to make sure that your power supply has enough oomph to run all the drives and the super-speed processor and all that stuff. This stuff all ends up consuming a lot of power. A lot of that power is translated into heat. So your system also needs to be able to get rid of the heat. There have been some interesting cases where you'll see computer furniture that's got a really nice enclosure for your computer that's not ventilated. And it's like, what are these people thinking? You put any computer in there and close the door, and this thing's just going to cook inside.

Leo: Well, your Pentium 3 might run fine, but nowadays you're right. You know, it's funny because this new laptop I got for Vista, the choice was a bigger, 5400 RPM drive, or a smaller 7200 RPM drive. And even though I know there's lower battery life, more heat, and now I'll have less capacity, it's only an 80GB drive, I still ordered the 7200. Because this machine won't be a main primary machine for me. So speed is what I'm looking for on this.

Steve: I completely agree.

Leo: And 80GB, that's enough unless you start putting media files on there, which I'm not planning to. It is Windows, after all. Let's see here. Rodney writes from Trinidad in the West Indies. Nice to have you listening in the West Indies, Rodney. You're probably having better weather than we are. Regarding network firewalls, can you guys bypass a network firewall to use a downloading program? He wants to use a peer-to-peer program like Soulseek, Limewire. He says: Is that possible, to bypass a network firewall DSL connection? I'm thinking this guy's got an ISP that's blocking peer-to-peer networks.

Steve: Well, to me, saying "network firewall," I was sort of thinking he was inside a corporation.

Leo: Well, but I think, as I remember, in the Caribbean a lot of the ISPs act this way, as well.

Steve: Okay. In that case, this is where some sort of VPN is your friend. The only way to really bypass a network firewall is to create a true VPN tunnel, that is, a virtual private network. And I would suggest that Rodney go back and listen to our VPN episodes where we cover this topic extensively, and basically use a supplier – this is not somewhere that Hamachi would work because Hamachi is sort of a peer-to-peer VPN that's used to create a tunnel between machines. What Rodney wants is a tunnel out to a VPN supplier, somebody who's selling a VPN service. So the peer-to-peer network would actually be connecting to that service and then tunneling the peer-to-peer traffic through the tunnel, through the corporate firewall or the ISP firewall, into his machine. So I can't even remember now the name of the VPN service that we liked and have used, Leo.

Leo: HotspotVPN.com.

Steve: Exactly, HotSpotVPN would be a perfect example of someone who could do that. It would allow you to use a peer-to-peer network. And he mentioned Soulseek and Limewire and so forth as a couple. But basically all of your ports would then be available, even if your local

network were filtering them and preventing that kind of behavior.

Leo: That's about it, huh? An ISP couldn't block a VPN, could it?

Steve: Well, sure. They could block – and this is typical of, for example, PPTP and L2TP and all those crazy acronyms when we were talking about VPNs. The good news is that modern VPNs are using SSL technology in order to use nontraditional VPN ports, and then an ISP is not going to be able to block you.

Leo: Okay, because it won't know what to look for.

Steve: Right.

Leo: Brian Voeller writes from Ashland, Oregon, USA, Earth – did he write that?

Steve: Yeah, he did.

Leo: We're glad to know what planet you're listening on, Brian.

Steve: Put down as his location.

Leo: And a big shout-out to all of our listeners in other parts of the galaxy. You have mentioned several times how 64-bit Vista will not allow unsigned kernel drivers, thus breaking some devices that require them. It's not that they require them, they just haven't written them yet. I'm wondering how much of a problem this will be with old hardware. I understand why antivirus software needs to be deep in the kernel; but what about things like printers, webcams, video, TV tuner cards that should have no need to modify the kernel in the first place? I find it hard to believe that something as delicate as kernel modification has become so commonplace that it's used instead of regular APIs. Unless, of course, this is poetic justice for Microsoft for hoarding all the good APIs. All the good APIs are taken. Would it be possible for Vista to allow old unsigned 32-bit kernel drivers to run in some kind of protected or isolated mode, like DOS and Windows 98?

Steve: Well.

Leo: Where do we begin?

Steve: The kernel is a different dog from applications. So it's certainly the case, for example, that DOS apps are able to be hosted in a so-called DOS box and not know that they're not just running on old 16-bit DOS. Similarly, 32-bit apps we know will be compatible with 64-bit – well, they are compatible with 64-bit XP, and they will be compatible with 64-bit Vista. The 64-bit OSes expose a completely compatible 32-bit API for those applications to use. The problem is that kernel drivers are actually a part of the kernel. There's no way that you could run a 32-bit driver which has not been recompiled for 64 bits. There's no way you could run that driver in a 64-bit kernel.

Leo: So they just have to write a new driver, that's all.

Steve: Yes. Now, you could imagine...

Leo: There's no protected mode? Because isn't that how we ran 16-bit drivers in 32-bit?

Steve: Right. Well, in fact, the way to think of it is that the driver is underneath that protected mode level. It's participating with the operating system. It's not being served by the operating system. Now, it happens that Windows I/O architecture in the kernel sort of has a client-server approach. So I could imagine that somebody could write an emulation layer that would allow 32-bit well-behaved drivers to be hosted in a 64-bit environment. For example, FreeBSD, the latest version of FreeBSD, runs Windows drivers. So you could literally use Windows drivers on FreeBSD. It's done that by emulating the whole device driver interface to a Windows driver. The problem is that a Windows driver that then also tried to make kernel modifications, well, you can imagine, there's no Windows kernel on FreeBSD. So that would just explode completely. And similarly, if a 32-bit driver that was being hosted by this hypothetical 64-bit/32-bit driver compatibility thing, if it tried to make any modifications to the kernel, that would explode, too. So you could imagine that a well-behaved 32-bit driver could theoretically be hosted by some special, really cool layer...

Leo: An envelope thing.

Steve: An envelope that sort of encapsulated it and pretended to be 32-bit Windows when it really wasn't. Still, if that was an ill-behaved driver, there's no way that would work.

Leo: And we don't want it. We don't want it. The point of having 64-bit, if you're going to have 64-bit, is running 64-bit drivers. Don't be silly. Don't try to run 32-bit drivers.

Steve: It seems to me that this really does make a nice sort of a cut point where older hardware that has to have older drivers is just going to be stuck with 32 bits. And it's going to only be the currently supported drivers. It's not like converting these things from 32 bits to 64 is a huge problem. You have a 64-bit C compiler.

Leo: You just recompile?

Steve: You recompile. You need to change some things. But it's not like, I mean, basically the whole I/O architecture has pretty much stayed the same in Vista. So it's not like you're having to write the driver from scratch. You just need to expand it to 64 bits and then run...

Leo: And the compiler should do most of that; right? Unless you've got some sort of dependency on register size.

Steve: The compiler does most of it, and then you run it through various driver certification suites that Microsoft provides...

Leo: That's the problem, because that's expensive.

Steve: ...to verify that it all works.

Leo: And that's why a lot of people won't do this, because they don't want to spend the money. Until there seems to be a big market for this, they're not going to do it.

Steve: We definitely have the chicken and egg. The good news is that my Vista had all the drivers for my hardware out of the box. And it may very well be that Microsoft, recognizing this chicken-and-egg problem, has got so many drivers built into Vista now, both 32 and 64-bit, that this is not going to be as big a problem as some of us fear. We really won't know until we start playing with it.

Leo: And I think to really underscore, you shouldn't be using 64-bit and trying to get 32-bit things to work on it. If you've got 32-bit drivers that you need, then stay with 32-bit Windows.

Steve: It's just time to say goodbye to those.

Leo: Or get rid of them, that's right. But to try to make them work under 64-bit is misguided, to say the least.

Steve: I think so.

Leo: Sheldon Smith in Minneapolis has been thinking about security. He writes: I've been listening to the Security Now! podcasts since the very first day. Just finished 65 – he's a couple behind. Listening to the discussions of security makes me wonder why Windows has so many holes. Back in the early '80s, before DOS, Digital Equipment Corp. created a 32-bit operating system called VMS. It was designed to be secure from the start. Since then it's been evolved first to use DEC's Alpha so that now HP's OpenVMS runs on the Itanium, which is Alpha-based. In 20-some years there have been no viruses – mostly because I don't think there are very many people running it.

Steve: That's where we're headed here.

Leo: And only a handful of worms or trojans. The stack grows up instead of down, just as you often say would be a good idea. Also there are no problems from buffer overflows. And the hardware has always supported hardware address protection, so even if data were to overflow, the hardware has "stone walls" to prevent the data from overwriting regions of code or having an executable stack. And customers have systems with uptime measured in years, and clusters with uptime measured in over a decade. And new systems are still being installed. Of course I like my Windows XP notebook, it's great for writing documents and playing games. This is a common kind of thing we hear from people who run BigIron. Why doesn't BigIron have the same problems computers have?

Steve: I wanted to mention this because it does touch on the Mac. And the reason that the Mac

doesn't have as many problems as Windows is partially that it doesn't represent the same size of target that Windows does. And so certainly the reason that DEC's VMS hasn't had virus problems is as you said, Leo. I mean, it's just not being used to the same level, and it's not being used by the same type of people. Remember that virus writers are often, not always, but often younger programmers. And they are writing viruses for the systems they have.

Leo: They don't have a PDP in the basement.

Steve: Right, exactly. You cannot write a virus for an operating system you don't have because inherently you need to know it intimately. And so we're already beginning to see more growth in Mac problems as the Macintosh becomes...

Leo: I will dispute that. I haven't seen any growth in Mac problems.

Steve: Really.

Leo: What Mac problems are you talking about?

Steve: Well, I don't know. Seems to me...

Leo: There's press all the time about, quote, "Mac problems never in the wild."

Steve: Okay.

Leo: They're all hypothetical.

Steve: Okay. Well, that's good. Let's hope it...

Leo: And in fact, I agree with you that a lot of the reason why the Mac is protected is because what you just said. But there's also a fundamental difference in the way it's written that provides more protection, I think.

Steve: We'll see.

Leo: We'll see. I think there's incentive to attack the Mac right now, just to prove it could be done. I think people are trying it all the time.

Steve: The fact that Windows has these problems is really not Microsoft's fault. I don't mean, you know, I'm no apologist for Microsoft. I've been on them since day one for policy things. It is just so hard to catch all these sorts of problems. And I really have to believe that, if the Mac were scrutinized to the same degree that Windows is, there would be problems there. On the other hand, having said that, things like ActiveX, which is inherently a bad idea, I mean, the Mac doesn't have those architectural mistakes. There's nothing you could call ActiveX other

than architectural mistake.

Leo: It has a real security model.

Steve: The idea that IE can invoke non-IE ActiveX controls rumbling around loose in the system and invoke overflows in those, I mean, that's just insanity. And so we're telling people to set the kill bit to prevent IE from having access to this ActiveX...

Leo: That's crazy.

Steve: Those are structural faults in Microsoft that now I'm sure they wish they hadn't done, just like they probably wish they had never allowed people to modify the 32-bit kernel. We see that Microsoft will never break compatibility. They'll just wait long enough to find some point where they can start things fresh. ActiveX was a clever componentized idea that they went too far with. So, yes, this is what's making Windows very vulnerable today. I have to believe, though, that there are unknown problems in the Mac which people may discover ultimately, although I doubt they will nearly be as significant as Windows.

Leo: I think it's very clear that there are application-level issues, especially with some of the older UNIX applications that are just kind of lying there on the Mac OS. We know of those. The thing that I think protects us a little bit on the Mac is just – unless you're running as an admin. But even then there's a little protection. Nobody really runs as root on OS X. And it's a lot harder for somebody from the outside to get into your system and modify it, even given these exploits and these holes. Unless you're running BIND on OS X.

Steve: Yes, well, and there's another feature in Vista that we haven't yet covered because it's extensive, and that's called User Account Control, UAC.

Leo: Right. They're trying to duplicate what's in the Mac.

Steve: Yes, they are. And so, again, it's so extensive we're going to do an entire episode on UAC in Vista because it is extensive and it is, well, it's wonderful, but it's also a mixed blessing because you're being hounded by this thing all the time, even if you're running as an admin.

Leo: They turned off – it did for a while. You had to enter the password a bunch of times. Now you just get a – and I don't know if this was a good idea. They made it simpler for the user.

Steve: Yeah, it keeps popping up "are you sure" dialogs.

Leo: Yeah, but you can just click OK now.

Steve: Or, I should say, "is this you" dialogs.

Leo: A little warning. But you don't have to enter a password anymore. And I'll be interested. I think the most vulnerable novice users just click OK all the time. And that's the problem.

Jonathan Green, a listener from London, asks: I'm running mostly Macs and a PC behind a Netgear wireless router at home. The last question, by the way. This might be the longest Security Now! ever. I might have to divide this into two parts. I'm running mostly Macs and a PC behind Netgear wireless router at home and have set it up to be what I believe as fairly secure. However, when I take my laptop out on the road or connect it to a corporate router, how can I be sure that I'm secure? What steps can I take within my laptop, rather than the router, to be more safe? I noticed when doing a ShieldsUP test at home, everything is stealthed. When I try the same test connected wirelessly to a work router away from home, some ports are open. Is it possible to make any security provisions on the laptop itself?

Steve: Well, what's happening in this case is that Jonathan is confusing the local security of his laptop and the security of the router. When he's at home, as he says, behind his Netgear wireless router, and he uses Security Now, we are testing the public IP, which is that IP of his router. And as we know, NAT is very good security as long as you have Universal Plug & Play disabled. So he's seeing stealth at home, not because his laptop is secure. He could have no firewall, for example, running on his laptop, and he would still see stealth at home because we're testing, and any hacker is attempting to penetrate, through the public IP, that is, the IP of the router, and it just blocks it cold.

Now, if in his corporate environment he runs ShieldsUP and he sees open ports, again, if this is a router and not just some sort of a non-routing gateway, then, again, the router's IP is what's being tested. So it means that the machine on the border in his corporate network has some ports open. Now, it's very likely that it has ports open, for example, it probably needs port 25 open in order for the corporate network to be able to receive incoming SMTP email. They may have port 80 open if there's a corporate web server running on the network. So corporate routers, corporate NAT routers which are connecting the corporation to the 'Net, they're inherently going to be offering services to the public which your typical home user is not doing. So a home user will be stealth; a corporate router will probably not be. But in both cases his laptop behind that router is probably safe.

What you definitely want to do, though, again in both cases, is be running your local firewall on your laptop. And this really goes for all users. It's nice to have the router protecting you. But it's still a good thing to run the firewall on your local machine. I know that you and I do, Leo. On my Macs I've got the Mac firewall turned on all the time.

Leo: And of course then there's also the HotSpotVPN that we talked about, which would be another way to secure it. That wouldn't change how your ShieldsUP is responding, though.

Steve: Well, exactly. ShieldsUP would then be testing the HotSpot server, and they've got a bunch of ports open. I remember because I was doing that at one point. I was interested in seeing what ports the VPN server had open. And also running a VPN wouldn't necessarily protect you from other access of your laptop by local people who are accessing your laptop's IP. The tunnel would only be allowing you to get a tunneled connection to the Internet. So really, you know, the short answer to Jonathan's question is, make sure you've got your software firewall up and running on your laptop.

Leo: All right. And with that we conclude this marathon episode, holy camoli, of Security

Now!. We thank our sponsor, Astaro Corporation, makers of the great Astaro Security Gateway, for their support. I'm very pleased to say that they're going to stick with us through 2007. We just got word from them, and that's great. We're happy to have them as a sponsor from very early on, not day one, but very early on of Security Now!. And we're glad that they're continuing forward with us into the new year.

Steve: Well, you know, Leo, we're doing a good job for them, and we know that they're doing a good job for our listeners because we get email from people saying, hey, I tried Astaro because we know that they're sponsoring Security Now!, and people really like it.

Leo: Well, it's nice because you can try it for free. Now, if you're a home user, of course, you can just download the software and put it on any old beater PC. Maybe that Pentium 3 you've got lying around would make a great Astaro Security Gateway. And then I think it's 79 euros a year if you decide that you want to subscribe to all the additional updates – the antivirus and the antispam and all that stuff. But the free software is very, very capable. Now, if you're a business, and you want superior protection from spam, from viruses, from hackers, complete VPN capabilities, you get intrusion protection, content filtering, and an industrial-strength firewall, I mean, this appliance is amazing, you can try it absolutely free by contacting Astaro at www.astaro.com, or call 877-4AS-TARO. I know a number of people have tried this and then gone on to become customers because they're so impressed. You can get a free trial of the Astaro Security Gateway appliance right there in your business. That's 1-877-4AS-TARO. And we thank them for their support and their ongoing support of Security Now!. It's really nice to have a sponsor that believes in the product so much.

Steve Gibson's website is GRC.com. That's where you'll find the 16KB versions of this podcast for the bandwidth impaired and full transcripts, too, if you'd like to read along. This will be a long one, I'm sorry, Elaine. Oh, she's going to be typing today. And it's mostly my fault because I kept interrupting. Let's see, what else? Oh, and that's where you'll find SpinRite, which is Steve's great product for hard drive recovery and maintenance. There is nothing better. I'm just going to say that. There is nothing better. It's the granddaddy of these applications and still the best. Version 6 now?

Steve: Pays all my bills, yes, SpinRite 6.

Leo: SpinRite 6, GRC...

Steve: Supports my coffee habit.

Leo: Something's got to. He's up to quenti venti lattes. If you want to hear Steve talk about that a little bit with Amber and me, he was on net@nite, which is up on the TWiT Network, last Sunday. It was really fun. Really great to have you on. Thanks for joining us.

Steve: It was fun with you and Amber again.

Leo: Yeah, I miss Amber. She's doing great at CityTV, though.