SECURITY NOW!

Transcript of Episode #20

A SERIOUS new Windows vulnerability - and Listener Q&A #2

Description: On December 28th a serious new Windows vulnerability appeared and was immediately exploited by a growing number of malicious web sites to install malware. Many worse viruses and worms are expected soon. We start off discussing this, and our show notes provide a quick necessary workaround until Microsoft provides a patch. Then we spend the next 45 minutes answering and discussing interesting listener questions.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-020.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-020-lq.mp3

Leo Laporte: This is Security Now! with Steve Gibson, Episode 20, for December 29, 2005.

Steve Gibson: Last episode of this year.

Leo: The last episode of 2005. And we've done 20 of them.

Steve: Yeah.

Leo: And we'll do another 52, I predict, next year.

Steve: I have a feeling...

Leo: Something like that.

Steve: ...we're on track.

Leo: We've been very consistent, and there's always lots to talk about. But one of the things that you started a few episodes ago, and I think it's a really great, is every fourth episode we answer listeners' questions.

Steve: Well, yeah, you know, just the nature of what we're doing, suggesting things to people that they're able to go try, sometimes people come back, you know, like more than one person, with, like, whoops, what about this or what about that? So it really gives us a chance just to, you know, sort of know what people are thinking, too.

Leo: Well, before we do that, let's talk about a big hole in Windows. And this was my original concern with doing every fourth episode on, you know, a schedule, you know, answering questions, that we might miss a big security story. So, but we're not going to do that because we're going to tell you about it.

Steve: Yeah, this week we've got one. This all exploded yesterday, on December 28. There's a brand newly discovered Windows metafile exploit which is really bad. It's been called a "zero day exploit." It is able to install malware in people's computers just by visiting a website. In fact, the guys at F-Secure, while they were fetching a file in a DOS box, it infected their machine because they had Google's desktop search system going. And it turns out, when they fetched the file, Google's desktop system indexed it. And the process of indexing the file caused the exploit to run.

Leo: Wow.

Steve: So...

Leo: So let me define a couple of terms real quickly, or ask you to define a couple of terms real quickly before we go further. What is a metafile?

Steve: A Windows metafile has been a format that's existed from the very beginning of Windows. It's sort of a scripting language, sort of in the feeling of a PDF, the way a PDF is able to define rectangles and squares and things. So it's a different kind of image format than people are used to. You know, we're used to GIFs and JPGs and PNG files. We're also used to...

Leo: What's the extension for it? Is it...

Steve: Well, it's WMF.

Leo: Okay.

Steve: Except that just filtering WMF files doesn't solve the problem because it turns out that other image formats are able to cause Windows to execute the same code. Windows actually looks at the content of the file, not just the file extension. So you're able to masquerade WMF files in non-WMF image formats.

Leo: So you could distribute this as an attachment to an email with any arbitrary extension. But more likely it's going to happen on a website. As an image, or just something that's part of the page, or...

Steve: Well, when IE displays it - this will affect Firefox, Opera, IE. So it's...

Leo: Doesn't matter, wow.

Steve: Yup, it's browser agnostic. At the moment, this is so new that they were only seeing websites exploiting it. But we're expecting a virus to appear at any moment, an email-borne virus, because, when you open the email, your viewer shows an image. That can install malware on your machine.

Leo: And this kind of underscores the fact that even what appears to be a pure data file can contain a virus or malware, as you say.

Steve: Right.

Leo: Now, you said it's a "zero day exploit." What's that mean?

Steve: Well, it means that it has the ability to propagate very quickly, and that it does not require any direct user interaction. I mean, there are many vectors that will get this thing into people's machines. Now, this happened yesterday. We have no response from Microsoft yet, although they do have a page on their site acknowledging the exploit. The good news is there's something that anyone who knows about this, anyone listening to this podcast, can do immediately. It's possible to unregister Windows's handling of the vulnerable DLL. What will happen is that, in Windows Explorer, the image thumbnails that Windows Explorer would normally show, they will stop functioning. But you want them to stop functioning because...

Leo: That's part of the problem.

Steve: Exactly.

Leo: Those files, yeah.

Steve: Exactly. It's Windows's display of a file which is infected with this malicious exploit does cause the problem. So on the show notes for this episode, which we always have posted - Leo, you can copy the string if you want to or just point people to the show notes page on my site. But I have an extensive dialogue explaining what the problem is, with lots of links to other sources of information, but also with a step-by-step explanation for how anybody can just drop a string into the Run field under their Start Menu to run a little utility program that Windows provides to unregister the handling, basically, of this exploit. And they will be safe, then, from any exploitation of this until Microsoft catches up with a patch. Now, as we know, Microsoft generally patches on the second Tuesday of every month. This is so bad, I would expect something from them as quickly as they can react because, I mean, already we're seeing, I mean, the malware guys are jumping on this fast.

Leo: Oh, so there's already exploits.

Steve: Many exploits.

Leo: Many. Within 24 hours. Wow.

Steve: Yes, yes. And, I mean, sites are exploiting it. The exploitation code has gone, you know, like wildfire across the Internet. So listeners want to do this immediately and probably, you know, email their friends with, again, a link to our show notes page to get people to protect themselves.

Leo: Yeah.

Steve: The word really needs to get out about this.

Leo: And I won't put the string on the TWiT website. I'll just put a link to your - because you describe it, you tell how to do - I mean, it's much more elaborate than just giving people a string. So we'll point people to grc.com/sn/notes-020.htm. I know that's a lot to remember, but ThisWeekinTech.com will have it on the entry for this page. So, wow, that's scary. That's really scary.

Steve: Yeah. Well, and also, as I always do, the top of the GRC.com homepage will have a link to our Security Now! page. And this week I will have a special link directly to the show notes.

Leo: Good, all right.

Steve: So just going to GRC.com's homepage will also take people. Because, I mean, this is a big, bad problem.

Leo: This strikes me as the kind of thing that an antivirus probably wouldn't be much help in. I mean, heuristics might catch this kind of thing, but likely not.

Steve: I would imagine you could design something specifically for it. So it might very well be that the antivirus guys will jump on this even more quickly than Microsoft can.

Leo: But you're not currently protected against it, in all likelihood; right? And it happened so fast that...

Steve: Yes. Absolutely not. It's funny, too, because it's taking advantage of a deliberate feature which was built into the metafiles. Metafiles, as I indicated, is sort of like a scripting format. On the web there's something called "scalable vector graphics."

Leo: Right.

Steve: SVG. And the metafiles are similar to that in that you define rectangles and lines. And so Windows sort of draws them. In fact, what a metafile actually is, is, like, GDI, the so-called...

Leo: It's a small program, really.

Steve: It's a little program that's basically calling the Windows graphics primitives to draw rectangles and text and so forth.

Leo: And that's why it can be taken advantage of. It's really not a pure data file, it's a program file.

Steve: It is a little scripting format.

Leo: Yeah, yeah.

Steve: And what's interesting is that there's a hook in there saying, if the metafile fails, execute the following procedure. Well, that's what this thing does.

Leo: Oh, interesting.

Steve: Is it's actually executing - it's taking advantage of the ability to execute a procedure in the event of a metafile failure. So the malware registers itself as this metafile procedure and then fails the metafile, causing the procedure to execute. I mean, it's a perfect example of just, you know, sort of a nice thing that represents an inherent security vulnerability which has finally been leveraged. And I've seen some reports saying that all versions of Windows, even like back in Millennium Edition and probably before, are vulnerable to this because it's built into the protocol.

Leo: It's interesting, too, because it's not a buffer overflow exploit. You don't have to trick anything. You're actually behaving as Microsoft intended you to.

Steve: Right.

Leo: Unfortunately, Microsoft - and this is the big flaw in security - assumed that people were kind and benevolent, and nobody would ever take advantage of this to do bad things.

Steve: Well, like putting scripting in email, which is, you know, still my biggest source of, you know, head-shaking.

Leo: Well, and TrueType fonts are also scripts. I mean, a lot of graphics file formats are not really bits. They're program code.

Steve: And it's a very powerful technique, but it is prone to exploitation if you're not very, very careful.

Leo: Got to sandbox it.

Steve: In terms of errata, before we start our Q&A, one person made the comment that I've been promoting the idea of cutting and pasting these really long, hairy passwords using Windows Clipboard as the transport. And he said, you know, Steve, it's worthwhile mentioning that you don't want to leave the password on the clipboard.

Leo: On the clipboard.

Steve: It's like, oh, good point. I'm definitely going to mention that.

Leo: You know, it's funny because we leave data trails, little crumbs behind us all the time, you know, our cache, our swap files, page files, you know, there's a lot of data left behind. And even the slack space on your hard drive contains data.

Steve: And our own browser caches. I mean, browser caching is done so that, as you move around a site, you're not having to always reload the same images that are, like, common among multiple pages on a website. But the browser caches everything it can, and it really does leave a trail that can be taken advantage of.

And then one last thing before we get into our questions. OpenSSH, the forthcoming version, it'll be 4.3, of OpenSSH is offering a VPN-style tunnel in addition to the traditional port forwarding that we've talked about many times with SSH. And so it's going to be extending the protocol to actually do this sort of tunneling. However, it'll have the same problems that TCP tunnels always have. And in fact, in an interview that Damien Miller, the developer of OpenSSH, gives, he says, and I'm quoting him, "Like any VPN system that uses a reliable transport like TCP, an OpenSSH tunnel can alter packet delivery dynamics, e.g., a dropped transport packet will stall all tunneling traffic." So it probably isn't good for things like VOIP over a lossy network. Use IPSec for that. But it's still used for most other things. So several people have written to me mentioning the fact that they've heard OpenSSH is going to be offering a VPN tunnel, you know, what does that mean? Well, it means it's another good thing, but it's not as good as using UDP as the tunneling transport. So OpenSSH is never quite going to give us what we want, whereas OpenVPN does.

Leo: All right. Let's move on to the questions. We have quite a few for you, Steve, starting with this one. And again, these are composites, questions you receive over and over again from a variety of different people. So credit to everybody who asked this question: I have a question about secure web pages. Some bank and credit cards ask for a username and password from what appears to be an unsecured homepage. So Washington Mutual does it; MBNA does it. When using such sites, I normally click the log-on button without entering any info, so I get forwarded then to a secure page informing me of my error and allowing me to try again. I have this question, too. He says, am I being too paranoid, or is this method of logging in actually secure?

Steve: That's a great question. And as you said, a bunch of people have asked. Most secure sites give people a sort of a warm and fuzzy feeling about entering data into a secure form by moving them onto a secure page before they ask them to fill anything in. You know, I wrote my own ecommerce system at GRC. And I do exactly that. I put people onto a secure page before I ask them to provide any information. The advantage is that you're able to right-click on the page, bring up the properties, check the actual security certificate that the page has to make sure that you're on the page that you think you are, and that you're not victim to some sort of redirection scam or phishing exploitation, as has happened in the past.

Leo: But now I'm on my Bank of America site, I'm at the log-in page, it isn't a secure page, and I just checked, there's no certificate for it. It says "website identity not verified." So that is a good question. Am I safe to enter my data here?

Steve: Yeah, it's a great question. The answer is probably. But this is the problem. There's no real way to guarantee it. Now, if you were to look at the code, if you were to view the source of the page, the actual HTML, chances are that the link that you would see, that is, the URL associated with the Submit button, it would be HTTPS:// and then whatever page is going to be accepting the data. So the point is that it's not the security of the page you're filling the form data in on that matters. It's the security of the page you go to that carries, that accepts that data. So it's possible and likely that, for example, Bank of America or Washington Mutual or whomever has set up the site, it's almost certainly the case that you're going to a secure page, and that's what's receiving your confidential data.

But, I mean, the reason this is a great question is that, unless you somehow verify that, you don't know for sure. And so I really think it's a fault of the website designer that they don't move you onto a secure page where the form is being filled out, even though technically it's not that page, it's the page that you're going to submit the data to, which is the next page you go to, which needs to be secure. And similarly, if they put you on a secure page, then it's possible that they could use an unsecure button to accept the data. I mean, it's really...

Leo: Oh, I see, because it's really the page you're going to, not the page you're on, that matters.

Steve: Exactly.

Leo: Yeah.

Steve: The idea is - this is all sort of a kluge which evolved after HTML and web pages existed. You know, the original concept of the web was just static delivery of pages that would - you know, basically people reading pages. There was no original concept of data going the other direction, that a web user would be submitting data back to the server. So that was grafted on later as an extension to the original concept, and it's got these problems. It's that, you know, it's the link you go to that needs to be secured because that's where your data is moving. It's not the page you're on. So it's a really great question. There isn't, without looking at the code of the page that you're submitting, there isn't a way to tell whether you're about to be secure or you're about to be nonsecure.

Leo: In fact, I've been, while you've been talking, browsing the Bank of America page. And it is - the post does go to an HTTPS secure site. So...

Steve: Right.

Leo: ...in a way this is guaranteed secure because I'm looking at it. On the other hand, if I were on a secure site, and I didn't look at the source, I might be going to an insecure site to do the submit. That's good to know, and it's actually kind of parallel to a question I've asked you, which is, I have secure web-based email, but the log-in doesn't seem to be secure. And you've reassured me now that it's still secure.

Steve: Right. Oftentimes, if you float your mouse over the button, sometimes your browser will show you the link down in the status bar at the bottom of your browser.

Leo: Oh, the button might tell you. Ah.

Steve: The button might. Normally it's more the case with normal links instead of Post links, where the data is not a Get request, it's a Post request, which is what most forms now use because that's a much more secure way of submitting data, because otherwise third-party advertiser things, you know, we've talked before about URL information leakage, which is also a problem.

Leo: Yeah. And Bank of America, which is - I have mixed feelings about this. They do put a padlock next to this - I'm hovering over the Submit button. It doesn't show me any security.

Steve: Right.

Leo: But they put a graphic of a padlock in this box, which is - I'm a little of mixed feelings because that is deceptive, in a way, because it isn't the same as the padlock you see in your browser.

Steve: Right.

Leo: Anybody can put a padlock graphic. But when you click it, it does explain exactly what you've just said. And it says, to provide the fastest access, we don't make the page secure, but when you do press Submit it is secure.

Steve: That's really interesting. Because, I mean, it no longer is the case that this represents any sort of overhead. I mean, lots of sites do it. I do it.

Leo: It's not that - it's not slow, is it.

Steve: It's dumb.

Leo: I mean, there is some handshaking; but it's not, I mean, come on.

Steve: So the idea that they've gone to all the trouble of putting up an explanation, rather than just securing that page, that's just quacky.

Leo: Question two, and actually it's a really interesting question: The Hamachi server - and we all know how much you love Hamachi, and I do, too, and how you recommended it - was down a couple of days ago. It's back online now. That's kind of scary. You want to talk about that? I mean, we are - when you use Hamachi, you're relying on it, aren't you, on their server?

Steve: Yes, totally. And so, you know, the guy wonders if it's useful for, I mean, he wants to make the point that we're relying on Hamachi to basically build our little network for us. And we're completely dependent upon it. Now, it's certainly the case with the VOIP systems, you know, you and I are using Google Talk. We relied on Google Talk being up...

Leo: Right.

Steve: ...in order to initiate our connection. But it's not the case that this system would work without Hamachi being up. Now, one of the things that Alex mentioned, as he's moved his versions forward, is that the clients now do not depend upon Hamachi's server to be up to maintain an existing set of tunnels; but it does have to be there in order to establish a new network. And in fact he's been doing some versioning on his server in the last couple days. And apparently there was an actual outage, also. And so people need to recognize that, yeah, it's really cool that this thing is free, but there is a dependence on the server being up. And that's of course the same with HotSpotVPN or iPig or any of the stuff that we've talked about. It will not be true once users are using their own OpenVPN scenario, which is where we're headed, of course. But, you know, nothing beats Hamachi for ease of use. It comes with this downside.

Leo: I set up my home network, someone writes, with WPA, as you recommended, and a preshared key using Steve's password generator. WPA worked fine on all my wireless GNICs - those network cards that support 802.11g - but not with the old 802.11b card in my laptop. Question. WPA doesn't work with a wireless "b" network card unless the new driver for that NIC supports it; correct?

Steve: Almost. Certainly having - if you can update your drivers, there is a chance that "b" will work. There's nothing about the...

Leo: "B" doesn't preclude it.

Steve: Exactly. There's nothing about the Bness versus the Gness of the Wi-Fi that does or doesn't mean WPA. So if you can get later model drivers for the card, and they support WPA, you're in great shape. If not, McAfee does offer a free WPA client which I've actually personally seen it work on a laptop whose built-in Wi-Fi driver did not support WPA. I have a link to it on our show notes page. It's a long link, so it's not something I can even say over the air. But they're trying to sell you, on that page this link goes to, much fancier sort of radius-style dynamic real-time authentication, which you don't need. So read the page carefully. There's a place when you're installing it where you say "disable the dynamic authentication." You definitely want to do that. Then you're just using the preshared key, which is what I'm promoting as really all the security you need. So it may be that you can use this free client to upgrade your WEP-only system to run WPA. And I know that it has worked for people.

Leo: So it's not necessarily a firmware update to the card you have to make. It's merely a driver or maybe not even a driver, a client that you need.

Steve: Yeah. They call it a Windows client software.

Leo: It's not even a driver. It's just instead of the - well, see, I mean, the problem is on some versions of Windows there is no WPA support. So that would fix that.

Steve: Exactly.

Leo: Yeah. But the driver would also have to support it, or not?

Steve: No, it does not need to support it.

Leo: Okay. It's just merely that you need the client.

Steve: Right.

Leo: Got it. What's the security of Windows Remote Desktop? Oh, we talked about this before, which is so easy to use, one man writes, and which you and Leo often talk about. One thing that I've always done is connected via RDP to my home PC from the office. I work as a network administrator, so no worry about the IT department catching on. But your show made me wonder about the security of Remote Desktop. Is it at all encrypted? Could it be sniffed? Would it be better to conduct the RDP session on a port other than its default, 3389? Perhaps that all should be a show on its own. And just to share what I do in terms of security, it's pretty basic. I have my router configured to allow traffic in on that port, which could be viewed as somewhat of a security threat. However, I also do some filtering with my ZoneAlarm firewall to only allow my office IP address through. So even though the port's open, only one IP address is allowed. It seems to work. When I connect from a friend's PC and log in, even as administrator, I'm almost instantly dropped by my firewall before I can even attempt to shut down the firewall or anything.

Steve: Well, it's a long question.

Leo: Yeah.

Steve: But it's a really good one.

Leo: Okay.

Steve: So we have a guy who's running Remote Desktop at home. He's opened a port through his router, so - and he's using 3389, which is the default. And then he's using the firewall on his machine, he's basically set up a packet filter so that only that port will be accepted coming from his office IP. Well, okay. A couple things. First of all, always changing Remote Desktop's default port is a super idea. And the good news is there are pages on the web that provide you with instructions for how to do that. So the default port, 3389, can be changed. And that's the first thing I would do is change it.

Leo: Now, that's security by obscurity. It's not really security. And I suppose there are sniffers that would just try the protocol - of course it would take you a lot longer, but across a bunch of ports; right?

Steve: Right. I mean, certainly you don't want to rely on security by obscurity. But it's better than not having any security.

Leo: Or being wide open and saying, here I am, come get me at 3389.

Steve: Yeah. So, for example, if, as just happened yesterday, we were talking about this Windows metafile exploit - if a horrible new exploit was found in the Windows Remote Desktop, then what would immediately happen is that worms would start scanning for port 3389, trying to get in on that port. They would just all assume the default. So being on 33890, for example, or anything else, is, you know, really the first thing you want to do.

Now, secondly, doing what this guy has done, that is, establishing a port filter so that he's only accepting incoming connections from his office IP, that is really good security. It's, for example, it's what I do in order to have Telnet available on my remote equipment at Level 3. Unfortunately, the equipment I'm using doesn't let me use a super-strong password. So what I've done is, I'm filtering the traffic so that only my home network range is even able to see those ports open where they are. I mean, nobody else is able to see them. And you can't spoof those IPs because they're TCP connections, which are inherently immune to any kind of spoofing.

Leo: Oh, I didn't know that.

Steve: Oh, yeah.

Leo: So even raw sockets don't let you spoof TCP.

Steve: Well, raw sockets would allow you to make up a...

Leo: Random.

Steve: ...fake source IP.

Leo: Right.

Steve: But the TCP connection handshake, it sends a packet in each direction. Both...

Leo: Oh, yeah, you couldn't - it couldn't get back, could it.

Steve: Exactly.

Leo: There's no return address. I get it.

Steve: Exactly. So what this guy has done is absolutely secure. The only threat would be somebody on his local Ethernet would be able to use ARP poisoning. But, you know, that would have to be somebody in his own home network, which is...

Leo: To pose as him on the IP address.

Steve: Exactly. And there you'd have a man-in-the-middle vulnerability. But you are completely invulnerable out on the Internet. So what he's done is really good. And I will just say that, otherwise, RDP is not secure. This is one of those things Microsoft keeps trying to do, and they're up to Version 4 now, and they still haven't got it. We talked about briefly there's a tool called Cain & Abel. And the current version of Cain & Abel, which is freely downloadable, has complete RDP man-in-the-middle spoofing and decoding so that, if it knew that that's what was going on, and if it was able to insert itself with a man-in-the-middle attack, it could, and it does, give you a complete transcript of everything the user does and types over his Remote Desktop connection. So you cannot rely on RDP security all by itself. You have to tunnel it through VPN, as we've talked about, you know, running Remote Desktop over Hamachi or iPig. Or, in this case, this user is just using it with port filtering so that it only accepts a connection from his office IP range. And that's really good security.

Leo: SSH, OpenSSH uses RSA fingerprinting in a similar way, saying, you know, this is the fingerprint for this machine that's contacting me. And if it's a known machine, allow it; otherwise, don't.

Steve: Well, what you're talking about is extremely good authentication.

Leo: Right.

Steve: And that's what you need in order to avoid man-in-the-middle attacks, strong endpoint authentication, which we'll be talking about very soon in 2006.

Leo: Even better than the IP authentication. You can't, you mean, you can't fake that, it's a public - it's a, you know, RSA key.

Play-Doh in Poughkeepsie writes: Biometrics. What do you guys think of all the cool new fingerprint scanners? I'm setting you up on that one.

Steve: That's okay. There's lots of people who are asking about biometrics.

Leo: Yeah.

Steve: I'm scared of them.

Leo: Well, you saw the - maybe you didn't see the story, that 83 percent of all fingerprint scanners can be fooled by a Play-Doh thumb.

Steve: Oh, I see what you mean. Oh, oh...

Leo: You didn't see this, you didn't see this.

Steve: I thought you said Plato and not...

Leo: Play-Doh.

Steve: ...not Play-Doh. Yeah. Now, I mean, that's a perfect example. I see these things that are, like, these little scanners on the end of a USB cord. And, I mean, I just know that they're going to be a problem.

Leo: It's kind of hard, I mean, you have to get an imprint of the actual guy's thumbprint and then use the Play-Doh to make it reverse, so...

Steve: Well, actually there are - unless you wipe one of those scanners off deliberately after you've removed your...

Leo: There's a fingerprint.

Steve: Yes. It turns out...

Leo: Of course.

Steve: ...you can fill a bag of water and push it against the plate again, and it thinks you're back.

Leo: Oh, that's not good.

Steve: Now, the good news is, on my little Toshiba Libretto, and also on the IBM ThinkPad, there's a different kind of scanner. It's a fingerprint scanner, but you have to draw your finger across it. It's a capacitive scanner instead of a static image scanner. And so the beauty of that is that at no point is your entire fingerprint available on a plate. Instead you're scanning across, like, a little single line scanner. And so it's a dynamic print, which it acquires over time. That's much more secure. But overall, the whole idea that any sort of single static biometric is being used for full authentication, I think that's always a bad idea.

For example, when I go to access my equipment at Level 3, I have two things I have to present. I have a badge, which I wave by their badge scanner. Then I stick my hand in a biometric reader, which reads the various parameter from my palm. So the idea is that - that's called two-factor authentication, and it is far better than single-factor authentication. So what that prevents is it would prevent somebody from getting my badge without my knowledge and going behind my back to the facility and getting in. Because the chances that they could have my badge and have an identical hand to mine is really much smaller than them just being able to get the badge by itself. So in real terms, I would say, if you had a system that required a password and a thumbprint or fingerprint, that makes lots of sense because then it prevents somebody from getting your password or discovering it independently and being able to log on as you, because, again, the chance that they're going to have a fingerprint which matches yours or have, you know, a spare ball of Play-Doh in their pocket, is very remote.

Leo: But don't assume that just the plain old, as convenient as it is, that the plain old thumbprint scanners are really secure. They're not, really.

Steve: You know, the rule of thumb is, anytime something seems really convenient, you've got a problem.

Leo: Secure does not equal convenient.

Steve: Exactly.

Leo: Regarding the information you have in Episode 19, using the secured connection that Google's Gmail offers - that's HTTPS://mail.google.com for Gmail - if one wanted to use this as their only email account when using an open Wi-Fi hotspot, would this provide enough security so I don't need VPN for my email? And I have, on FastMail, which I love, my IMAP email provider also has a secure log-in. In fact, they're nice. They have a public terminal and a secure log-in, so it saves no information, and it's HTTPS. Is that safe over Wi-Fi?

Steve: Absolutely. I mean, that's the short, one-word answer.

Leo: Yeah.

Steve: That's going to be using SSL. And we're going to be talking, again, early in 2006 about SSL and certificates and authentication and all of this stuff. You would want to make sure that the connection and the certificate had not been spoofed by right-clicking on the secure page and looking at their certificate. I really - I'm going to start talking about this a lot because it's the way phishing scams work, and it's the way SSL man-in-the-middle attacks work, is by presenting people with fake credentials, because that breaks the authentication, which is inherently the way these protocols are secure. So you really want to make sure that you're actually talking to the server you think you are. If, however, you are, and you're, for example, using secure Google Mail or whatever, then, you know, all of your email work would then be secure in an otherwise insecure area, you know, hotspot or in a hotel or so forth.

Leo: Very good to know. How secure are NAT routers, Steve? Do stateful routers offer more security? I guess he's talking about stateful packet inspection?

Steve: The SPI, right.

Leo: SPI. Do firewall routers, or routers billed as firewall, offer more security than just plain old NAT routers?

Steve: It's a really good question. You know, we've talked about NAT routers extensively as providing inherent security because they don't allow unsolicited traffic to come into your network. My sense is that's really enough. Now, if manufacturers are going to offer more features, for example, stateful packet inspection, where they're actually tracking the state of the protocol as it goes back and forth, I mean, it provides theoretically better protection. And I like the idea that maybe it offers much more configuration flexibility. For example, the guy, for example, who was configuring ZoneAlarm only to allow incoming traffic on a certain port. A firewall router will probably have that built in, so you wouldn't have to use ZoneAlarm on your computer, but instead you could say allow incoming traffic on this port only from this IP range. So more features of a security nature are probably a good thing. But I wouldn't ever suggest that someone go out and buy, like, buy another router just to get those more features.

Leo: Well, and also you could legitimately label any router, NAT router, a firewall. So just because it says "firewall" doesn't even mean you get more features.

Steve: And any NAT router is doing...

Leo: They're doing it.

Steve: ...stateful packet inspection.

Leo: Right, right.

Steve: That's what it means to have a packet come in and inspect it and check the state of the connection table to figure out which machine to send it to. That's...

Leo: Does any form of SPI do more than that, look at the actual contents of the packet or, you know, try to assess what's going on, anything like that?

Steve: They say they do, and it's theoretically possible; but how would we ever know? I mean, no, I mean, really, it's just - don't buy one. I mean, if you don't own a router yet, buy the most secure fancy-schmancy stateful packet inspecting firewall, you know, router that you can afford. But...

Leo: And you still may not be getting anything better than your Linksys WRT54G; right?

Steve: Well, certainly if it's got features that are exposed on the user interface...

Leo: Right, right.

Steve: ...of the router. If it's got extra fancy, you know, firewall features, timed availability so, like, instant messaging is only available during these hours, you know, those kinds of things are cool. But...

Leo: And useful, yeah.

Steve: But really, any NAT router is providing you with great security. Beyond that, well, it doesn't hurt.

Leo: There's a company which shall remain nameless, mostly because I just can't remember the name, that sells what is essentially a router with an on/off switch, and - for twice the price - and claims it to be a, you know, high-security firewall router.

Steve: That's got to be the dumb idea of the century.

Leo: Well, you could turn off the Internet. I mean, that is secure.

Steve: Or you could just unplug your computer.

Leo: Yeah. Well, that's a good point, yeah. You just shut down. That'd have the same effect, wouldn't it.

Another correspondent writes - and by the way, if I could ask people, when they fill out those forms, if you want to get a question on the show, it would be great if you'd put your name and city, just so we can give you credit.

Steve: Yeah, that'd be cool.

Leo: Yeah, just, you know, say - make up a name if you want.

Steve: You know, before we leave the last topic, I'll mention that one of the cool things that Windows 2000 and XP do is, if you right-click on Network Neighborhood and then do Properties, it'll show you your various adapters. You're able to right-click and disable an adapter. And in fact I use it on my laptop all the time because I've got, you know, adapters coming out my ears. I've got multiple Wi-Fi, I've got, you know, VPN, I've got several different interfaces. And I deliberately keep them all disabled, and then I selectively enable the one that I want to use. So that's another nice way of sort of taking your machine off of your LAN if you're not using it. I mean, it's certainly not super secure, but it's a little nicer than literally unplugging your computer from your router.

Leo: And it would protect you against incoming traffic, although if you had a bad boy on your system...

Steve: Right.

Leo: ...they could probably turn it back on.

Steve: No doubt.

Leo: If they were smart, you know, if they were checking for that, I guess.

Steve: Yup.

Leo: I'll make up a name. Doleful in Duluth writes: If I use my home network from somewhere else to access the Internet, what would the resulting speed be?

Steve: Ah, that's really a good question. And it brings up a good point. We will be promoting the idea of people running some sort of an OpenVPN server on their home network, either their own little, you know, plastic Linksys router that they've reflashed with OpenVPN, or running OpenVPN server on a Windows machine, which is then used, not only to give them access inside their network, but as their connection to the Internet. But we know that many connections to the Internet are not the same speed upstream as they are downstream. Instead of being symmetric speed, they are asymmetric connections. And what's normally the case is that the traffic downstream toward them is high speed, for example, for downloading, but their outbound traffic is low speed. And this allows their providers to basically sort of change the bandwidth allocation to optimize it because upstream traffic actually takes more channel bandwidth if you're going to have high-speed upstream traffic. But when you're out roaming around, and you're connecting to your server at home, then your traffic has to come down your connection to your home and then, if you're going out on the Internet, back out your connection to go on the Internet. Basically what this means is that your effective speed is the lower of your upstream and downstream speed. It's not the greater. So surfing the Internet using your home connection as your entry point to the Internet would be slower than just using it from your home and not needing to have traffic coming downstream as well.

Leo: That's a very good point.

Steve: Yeah.

Leo: Can a Sony rootkit or any rootkit be removed by reformatting the hard drive and doing a clean install of the operating system? Isn't that one way to get rid of it?

Steve: Yes, absolutely. It's...

Leo: Nothing's going to survive a format and install.

Steve: Nothing. Now, it's a little trickier if you had a multiboot environment, where you had multiple operating systems on the drive. You'd want to make sure that you got it off of the drive. There are viruses which are able to install themselves in the partition table and in the first track of the hard drive, that are able to install or able survive reformats and reinstalls of the operating system. But that's not a rootkit. A rootkit is a compromise of the OS. So if you wipe it clean and reinstall it, you know, that rootkit is gone.

Leo: Good. Sony, by the way, recently announced that they were going to settle a class-action lawsuit against them over the rootkit. And they're going to offer people money for the CDs that they bought, and they're going to offer trade-ins, and they're going to stop doing it and offer an uninstaller. So maybe they've seen the light.

Steve: Oh, if this hasn't shown them the light, they'll never see it.

Leo: Nothing will. For the VPN, I would love to set one up on my home server. Unfortunately, my DSL provider gives me a private IP address, not a public IP address. Is there any way, short of getting a public IP address, to connect to my home network from the outside world?

Steve: Wow, that's another great question. So...

Leo: I haven't heard of this. So what is going on here?

Steve: And it is becoming more and more common. He's actually got a provider, his ISP is giving him, like, a 10. address, or a 192.168. So he doesn't actually get from his ISP a publicly routable IP like most cable modem users have, for example, you know, it's 66 or 64 or something.

Leo: That seems like a good idea from a security point of view.

Steve: Well, it is because he's never having any unsolicited stuff coming in. Basically what it means is, his provider is running a big NAT router on his behalf. The bad news is, it totally precludes him from running any servers.

Leo: Ah, which is probably another reason the ISP does it.

Steve: Well, exactly. So the only solution for him is Hamachi.

Leo: Hamachi gets around that.

Steve: Yes. If Hamachi will work for him - and he'll have to give it a try. But if Hamachi will work, it will in the same way that Hamachi is able to go outbound through a user's own local NAT router, it ought to be able to go outbound through his ISP's NAT router.

Leo: It does NAT traversal with the ISP's router.

Steve: Exactly.

Leo: Interesting. So it tunnels through, basically, the router, and you've established a connection.

Steve: Yeah. So I would imagine that, if Hamachi works for him, it's the only thing I can think of that would work for him.

Leo: Another reason to love Hamachi. This is kind of related to the slowdown question before. Isn't using a VPN inherently slower than just an unprotected connection? I guess the encryption would slow it down, wouldn't it?

Steve: The encryption can. And in fact, in the case, for example, of Hamachi, Alex will be offering a non-encrypted variant for Hamachi, specially for gaming guys who are extremely concerned about the latency through the network. The fact is, on modern computers, encryption is a very low overhead overall.

Leo: They're fast.

Steve: They really are. I mean, encryption, I mean, you could even use RC4, which is virtually, I mean, it's so fast and still very good encryption that is just representing no overhead. So, well, for example, Leo, I'm going to be using RC4 in my own experimenting with VOIP stuff because I don't want encryption to slow us down...

Leo: Right.

Steve: ...and there's just no need for anything more. However, VPNs are slow, not because of encryption, but for two reasons. And I want to hammer this home. In fact, people are complaining about iPig and slowdowns with iPig, and in fact iPig's stalling sometimes. Stalling is what happens when you are tunneling TCP through a TCP tunnel. If you suffer packet loss, then the external tunneling protocol, TCP, will start asking for packets to be retransmitted. But the internal TCP protocol will also start asking for packets to be retransmitted, which ends up creating a complete bog down of the VPN protocol. So this is really a fault, not of the VPNness, but of the particular technology used to implement the VPN. So unfortunately, iPig, cool and free as it is, it can have problems, which you will see if you're on a connection which is too slow or is suffering from packet loss. You can have problems. The traditional VPN protocols, IPSec and L2TP and PPTP that we've talked about before, they avoid using TCP just for this reason. They have their own tunneling protocols. And what is becoming my favorite solution, OpenVPN, it's able to use UDP in order to avoid that problem. So it's really, I mean, encryption is some overhead, but it really should no longer be a real problem. It's just the implementation of the VPN that can cause problems.

Leo: Won't disabling - another viewer writes, Secure in Encino - won't disabling my access point's SSID keep my network secure? Now, I thought we'd flattened this one. This is turning off the broadcast ID.

Steve: Yeah, we have flattened it. And it brings up a good point. We're going to be needing to sort of move forward with these security topics and not continually address things that we have in the past. So I'm going to be presuming that we're building on a foundation of knowledge as we move forward and be referring to things that we've talked about. So I want to answer the user's question; but I want to say that, you know, we've covered this, and it's really going to be important for people to spend some time to catch up on our prior podcasts so that they're moving along with everyone, and we're not penalizing the people that are, you know, keeping themselves current.

Leo: Yeah. I mean, we have the same problem on the radio show and the TV show, and there are certain things that just come up again and again and are worth, I guess, addressing. But we have covered this one pretty thoroughly. Just the short answer is...

Steve: It prevents - blanking or disabling the broadcast SSID from a node does not provide you with security, but it does provide you with the prevention of people using your bandwidth by mistake. It's much like MAC address filtering. If you tune your access point only to accept data from known adapters or known MAC addresses, then it will prevent your neighbors from inadvertently using your node and your access point. But it will not prevent someone who really wants to use your connection from doing so. So both MAC address filtering and SSID hiding, they will prevent inadvertent access; but they should not be confused with real security.

Leo: And finally, could you explain port knocking - that's with a K-N, knocking - and how I can use it? From what I've heard, it seems like a great solution for allowing controlled access to a system or network. What is port knocking?

Steve: It's very cool. Unfortunately, it's still very unavailable. The idea is, say that you have a router that's got no exposed open ports, yet you want to connect into a router from outside. For example, say we were running OpenVPN server, and a user wanted to connect into his OpenVPN server, but he didn't want to leave his router's port open and forwarding traffic in all the time. Port knocking is just so cool. You send a sequence of packets at your router on closed ports. And, you know, closed ports don't do anything. But the router knows they hit there, and it drops them. But the idea would be, you configure your own secret sequence of packets to send.

Leo: You mean, knock, knock-knock-knock-knock, knock knock.

Steve: Yeah, exactly. Exactly. And so you configure your own secret sequence to send. And when the router sees it, then it knows it's you because you've given it the secret knock. And then it opens the port and allows your data through. So...

Leo: I like that. It's kind of like a call-back, a little bit.

Steve: Yeah. It's, well, it's very cool. As far as I know, no commercial routers offer it yet. It's becoming available on, like, you know, in Linux systems and in other Open Source systems. I hope it becomes a widespread solution. It means you need something fancy over on your computer to send out the special knocking sequence. But, you know, that's not a big deal. And it's just - it's a very cool solution because it allows you to be absolutely stealth until it knows it's you. Oh, and the source IP of the knocking packets. If the router was really clever, it wouldn't open the port to everyone, it would only then allow unsolicited packets from your IP from which the knocking packets came.

Leo: Oh, very important, yeah.

Steve: So it's very - it's a very, very cool technology.

Leo: I like it. Knocking.

Steve: As it happens, we'll be letting people know.

Leo: Oh, we are? We're going to do some stuff on port knocking, huh?

Steve: Oh, sure. I mean, as it becomes available.

Leo: I hadn't really heard about it. I like it. Well, I thank you for making yourself so available. Steve Gibson, it's always a pleasure. Security Now! appears every Thursday afternoon, thanks to the good graces, I have to say, of your company's SpinRite and GRC.com. SpinRite is the world's best file recovery and disk maintenance utility, still available and always getting better.

Steve: Well, yeah. In fact, it's at Version 6, of course. And you're right, it's SpinRite's ongoing sales and success that basically, you know, pays for my bandwidth, pays for my time, and lets me do this.

Leo: Yeah. Because this has pretty much become a full-time job for you, I know. I know at least it was with OpenVPN. And of course, thanks to our friends at AOL Radio who broadcast this show on their podcast channel and offer us the bandwidth so that the downloads are free to you. That's at AOLmusic.com.

What are we going to talk about next week, Steve?

Steve: Next week we're going to sort of start off the New Year with talking about how the Internet works. It'll be a review for people who already think they know it all, but it will answer a question I had in the very beginning and that I know lots of people have, you know, what actually are packets? What actually are ports?

Leo: Oh, great.

Steve: And how does all that work? We're going to sort of start laying a foundation for security understanding by talking about some stuff that we've referred to but have never really taken the time to define.

Leo: There's a couple of people I saw complain, well, sometimes it's a really advanced show, and sometimes it's a really basic show. But I like the mix because I learn something on all of the episodes. And eventually, you put it all together, you've got a pretty good education in how this stuff works and how security works.

Steve: Yeah, in fact some people have written that the Security Now! podcast is becoming an assignment. Their teachers are telling...

Leo: That's great.

Steve: ...to listen to the podcasts. So...

Leo: I like that.

Steve: ...it's very cool.

Leo: You're the teacher we always wish we'd had. Thank you, Steve Gibson.

Steve: My pleasure, Leo.

Leo: Happy New Year. All the best in 2006. Look forward to another great year of Security Now!.

Steve: Here it comes.

Leo: Take care.