Backtracking Spoofed Spam eMail

Description: Leo's 'TWiT.tv' and Steve's 'GRC.com' domains are used by spambots which spoof their domains as the source of bogus eMail. This week they discuss the details of eMail "Received:" headers and explain how the examination of those headers can penetrate any spoofing to reveal the true originating IP of any spoofed spam eMail.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-079.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-079-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 79 for February 15, 2007: Spambots.

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

It's time to talk about security with everybody's favorite security maven. I'm going to call you a "maven" from now on, Steve Gibson.

Steve Gibson: I'm not sure what a maven is.

Leo: It's good. I know it's a good thing. It's like a, hmm, I don't know how to describe it.

Steve: An advocate, maybe?

Leo: No, no, like a big shot. A maven is regarded by cohorts as a trusted expert in a particular field and who seeks to pass his or her knowledge on to others.

Steve: Boy, you are fast with that dictionary, Leo.

Leo: That's Wikipedia for you, baby. It's right there. So you are – I would say that's pretty

much you, a trusted expert in security, and you seek to pass your information to others.

Steve: I'd go along with that.

Leo: You da maven, man. And today we're going to talk about what, Mr. Maven?

Steve: Well, a couple things. It's funny, I was listening to one of your recent TWiT broadcasts, I sort of listen to them in the background when I'm doing work where I can have one ear on that. And you guys were talking about – this was in the last couple weeks – the really continuing expansion of zombie botnet fleets. Do you remember that it was a couple weeks ago somebody had some stats about the percentage of PCs that are now believed to be infected with some sort of remote-control bot.

Leo: Yeah, it was Vint Cerf, the father of the Internet, who said that, of the 600 million PCs out there, he believed 150 million, one quarter, were infected. And you know, I've talked to a number of people since then who think that number is low. It's certainly a very big number.

Steve: Well, it's phenomenal. And of course that corresponds with two things because these bot fleets are being used both for spamming and for attacking with denial-of-service attacks. And what I wanted to talk about today was the issue of spoofing email. It was triggered by – well, before we get into that we probably ought to talk about errata. And then we'll start talking about it.

Leo: There's also a news story that we'll talk about because botnets were used to try to bring down the Internet just last week.

Steve: Right. Again, an attack on DNS servers.

Leo: The root servers, which you have always said is the weak link.

Steve: It's a vulnerability, for sure.

Leo: And at some point, maybe this isn't the episode, but I would love to talk to you – we did it on Call For Help, and I don't know if we've ever done it on Security Now!, about that vulnerability. And they are trying to do things to protect you. But I would love to talk about that at some point on the show.

Steve: For sure.

Leo: But meanwhile, let's cover errata from previous episodes.

Steve: Well, actually this is not from previous episodes. But the big news today, and actually it's big enough news that it's been picked up by a number of different news stories on the 'Net,

is Microsoft released a surprisingly large batch of vulnerability fixes on Tuesday.

Leo: For Vista or XP?

Steve: Well, for Windows. But they included Vista, which was one of the things that I picked up on immediately, and I was looking at it. What caught my eye – okay. First of all, Microsoft tied themselves for the largest batch of vulnerability updates they've ever released. There were six critical and six important. But together those 12 updates fixed about 20 problems because a couple of the updates fixed multiple problems. What caught my eye as I was just sort of scanning it, first of all, I looked at the length of the email that I received from Microsoft. And it's like, whoa, this is more than your normal couple one or two.

Leo: So you get an email that says, watch out, here it comes, this is what it's going to be.

Steve: Exactly.

Leo: I don't know if I get that email.

Steve: You're able to subscribe to it somewhere on Microsoft's site. And...

Leo: So it gives you a heads up.

Steve: It does, and it sort of keeps me in the loop. Now, the reason that I use that is I don't like the idea of this Windows Update running behind my back and automatically downloading and installing things for me.

Leo: I do. I let it do it. I just say, go ahead, you do anything you want.

Steve: Well, and I think for most users it's probably the right thing to do. At least maybe to download them and then advise you when they're ready to install. But for example, on my old creaky Windows 2K machine that I'm still using, I'm not moving from IE6 to IE7. And of course Microsoft is really pushing IE7 hard. So I don't – and then they've got that .NET stuff. Now they're at .NET 3.0. And that's 50 megs of blob that I just don't want on my machine. So, you know, and I think like many of our listeners probably, I'd like to have more control over what Microsoft is pushing onto my hard drive behind my back.

So the reason I bring up this security issue is, if there are other listeners among our growing base, I think you and I were just talking a minute ago, the numbers seem to be increasing of Security Now! listeners. I wanted to make sure that people would do an explicit verification, use Windows Update under your Start menu to grab this stuff. Because there were six zero-day vulnerabilities, which is to say, as we know from having talked about those before, those are exploits which appear and surprise everybody, exploits which are active and discovered before anyone knows there's even a vulnerability, let alone a patch for it. Now, they were not widespread exploits. Otherwise our listeners would already have heard about them from us. They were only being used in more selective, targeted attacks. But they were discovered on the 'Net being used. And it was from them that these vulnerabilities were reverse-engineered to find out what it was these things were exploiting. So there were, among these 20 problems, six of them were zero-day...

Leo: Six.

Steve: Yes.

Leo: That's unbelievable.

Steve: Well, and, okay, the first thing that caught my eye, as I was saying, when I was scanning through this email I saw "Vulnerability in HTML Help ActiveX control could allow remote code execution." Well, it's like, okay, that's a bad one. Because once again that's your typical browser-based exploit that we've talked about a lot where doing something that brings up a page allows a remote site to have its way with you and your computer behind your back. Essentially, you know, the kind of problems we're going to be seeing are mostly this type now because Microsoft has finally got a firewall that's running by default. So I don't expect to see wide-scale worm problems the way we have been before. That is, as we said recently, in fact, mostly that the problems that we're seeing now are people visiting unsafe sites with scripting active that allows their machine to basically run a script that the website provides, which you know always makes me very nervous, and they get their machines taken over. So...

Leo: It's not that the worms aren't out there. They still are, but you're protected as long as you're running that Windows firewall.

Steve: Correct.

Leo: So there's got to be some other way.

Steve: Well, and even if you didn't have the firewall, any currently patched system will not fall victim to the old Code Red or MSBlast worm or any of the worms that are still out there poking around and we'll probably never get rid of. On the other hand, it's just so fundamentally wrong to have your ports open and exposing services that you don't really need to have exposed, which are what worms are looking for. So you really want the firewall up. But then the other vulnerability out of this list of 20 that caught my eye, get this, was in Microsoft's own Malware Protection Engine. There was a remote code execution vulnerability there in the way it parses PDF files. So, and this is Windows Defender for Vista. So...

Leo: Wow. Defender has a bug in it that will allow a bad guy to remotely execute code on your system?

Steve: Yes. It says...

Leo: Via PDF.

Steve: Via a PDF. And the problem is, Vista's Windows Defender scans incoming stuff without user intervention. It's trying to protect you. In the process, it exposes a buffer overrun vulnerability, exactly like we've been talking about a lot here recently. This was found and reported by the IBM X-Force guys, the ISS guys. And they say, "By sending a specially crafted PDF file, an attacker can trigger a heap overflow, resulting in remote code execution. This file

may be sent over common protocols such as SMTP" – you know, email – "HTTP, FTP, et cetera. In many cases, this vulnerability may be triggered without user interaction. The vulnerability exists because an arbitrary integer from a PDF file is used in a memory allocation calculation without proper bounds checking. As a result, an attacker may provide a large integer value, creating an integer overflow in the calculation. This causes a heap overflow with arbitrary file data. There are several structures on the heap that an attacker may abuse to obtain remote code execution."

So, I mean, here is exactly the kind of stuff we've been talking about where a PDF file contains a very large integer, some math is done, it wraps around to a smaller value, then a small amount of memory is allocated when a large amount of data is provided which overflows the buffer that was allocated and, wham, you've got a compromised system. So, and I think the little – a couple articles that I've seen have commented that, well, you know, the one thing you want to have really be secure is the software which is securing you. So here's a vulnerability...

Leo: Yeah, no kidding.

Steve: ...in Windows Defender, and it affects Vista as well. So again, these are really difficult things to fix. Now, what I haven't been able to find, nor have I been able to test this, is whether having hardware DEP running would protect against this. I'm going to keep my eyes out to see if I can find any demos for this. Oftentimes exploits will surface a few weeks after the patch, just looking for people who haven't made their machines current and kept them updated. So I may be able to find some sample code to exercise this exploit, which would then allow me to play with hardware DEP under Vista and XP and see whether having it enabled would have already protected people. Because of course it would be very cool if this was never a problem for people who had hardware DEP running. We don't know one way or another about that yet. And anyway, so I absolutely want to let people know that if they don't have Windows constantly checking for updates, go and do one now because there's a big package of goodies waiting for you at Microsoft that you really do want to have installed in your machines. And this is virtually all versions of Windows. I mean, Windows 2K, XP, even Vista, all have issues that are being fixed by last Tuesday's – second Tuesday of the month. It was also the 13th of February. So I thought, well, that's the right day for...

Leo: It was a little Valentine's Day gift early.

Steve: A little pre-Valentine's Day, yeah, exactly.

Leo: Got the love. You know, I'm looking at my system, and it has yet to tell me I have an update, so that's interesting. Even in and of itself.

Steve: Yeah, it is taking a while for these things to get pushed out by Microsoft. I notice that it sometimes is a day or two after the second Tuesday that Windows finally gets around – I think the problem is so many people are using this, and of course Microsoft has all this enabled now for many years in XP and certainly all in Vista, that there's got to be a tremendous load on Microsoft's servers, especially when you get patches of this size. So they're probably metering out the rate at which they notify people of these updates in order to get some control over it.

Leo: But with six zero-day exploits it seems a little foolish to wait too long. And this is why I, you know, have turned on the automatic updates and just say go for it, you know, because I don't want to take the chance. Is there something people should do before they

receive the downloads, if they haven't received them yet? Should they stop surfing?

Steve: Well, it's always the case that, as you know, my standard advice is be very wary of scripting. So many of these are, again, scripting-based vulnerabilities that are caused by someone browsing to somewhere unsafe. One of the other ones was really interesting, it's 016, it's technically the 16th one of 2007, and that's the cumulative security update for IE, both IE6 and IE7. And there were – that's one of the updates that contained multiple vulnerability fixes. It turns out that, if a script tries to instantiate, which is the term used in all this ActiveX or COM stuff, if a browser script tries to instantiate a COM object which was not intended to be instantiated by IE, and there is just a bazillion of those, it turns out that a vulnerability was found in the way that happens that allows remote code execution. So we've got vulnerabilities in ActiveX controls that are designed for IE to be able to run them. And now we've got even a bigger class of problems because it's been found that it's possible to exploit the instantiation process for COM objects in the system which were not intended to be instantiated by IE. I mean, it's a mess.

Leo: Let me just make this clear. Which of these are for Vista, and which of these are not for Vista? I mean, Defender obviously is in Vista. IE7 is in Vista. So I presume...

Steve: Yes, but IE7 in Vista is not vulnerable to this particular problem. And IE7 is also not because there are – IE7's enhanced security over IE6 does prevent, by default, it prevents this particular exploitation. The problem is, if IE6 had been configured to allow some of these things to run – because you might have some corporate system where for some corporate website or internal Intranet you were using some COM objects in your normal daily business. If you then update IE6 to IE7, in order for the update process not to break things, it will carry those permissions forward, and you'll still be vulnerable.

So anyway, unfortunately it's not a simple answer. But it is the case that IE7 under Vista is not vulnerable to this particular 016 set of problems, although Vista is vulnerable, as we said, the Windows Defender component of Vista is vulnerable. And in fact there's, like, eight different modules that Microsoft is now using for its antiviral stuff. It's like Windows OneCare, that's vulnerable. And there are several other places where they've got email scanning engines and things that are more IT oriented than end-user and consumer. All of those, all eight of those things use the common Windows Defender core, and they all have this PDF file vulnerability.

Leo: Wow. And so Defender, and then some of the other patches, are they XP specific? Are there any Vista-specific patches, just for Vista?

Steve: No. There was nothing that was only Vista except Windows Defender. Of course you are able to put Defender, I guess, on XP. So that sort of falls back there, too.

Leo: It's both, yeah. But that's actually quite interesting. I mean, in a way that's encouraging. There were no zero-day exploits that were particular to Vista.

Steve: Right. And at the moment what we're seeing is we're seeing the common code that has always been in Windows that unfortunately Vista has carried forward is still causing a problem. It is also troublesome that, okay, it's troublesome that Windows Defender has this problem because it was clearly recently written. So you can't argue that there was lots of legacy code in Windows Defender.

Leo: Well, no, because they based it on an antispyware program from Giant that is fairly old.

Steve: Ah, that's a good point.

Leo: So they could have inherited bad code.

Steve: Exactly. And probably this...

Leo: I'm going to let them off the hook there.

Steve: The PDF file parser – well, actually I was going to let them off the hook by just again saying that this kind of stuff is so hard to find. I mean, a programmer could stare at that, I mean, and as a programmer I have stared at my own code, a programmer could stare at that integer math and look at it and just not see, I mean, just not see how it could be misused. It's so difficult to find these things.

Leo: Yeah, but when you're accepting input from a user, those are the places you check.

Steve: Yes, exactly. Well, especially, see, as we know, a PDF file is basically a form of encapsulated PostScript. And PostScript...

Leo: Right, so it's programming.

Steve: It's a scripting language, exactly. So when you open a PDF, you're running a scripting code. And that's exactly why a static file like a PDF would have an integer that something would be using to drive an allocation of memory that creates this problem.

Leo: Very interesting.

Steve: So anyway, the takeaway message is make sure, everybody who's listening to this, that you know that last Tuesday was a major patch event, the second Tuesday of the month, for Windows. And you want to make sure your machines are current because, although there's no widespread exploitation of these things, they generally do get more popular after the news gets out, after a patch occurs. We see more exploitation then in many cases than beforehand. And I wouldn't be at all surprised to see that happen here.

Leo: But again, I'm going to look at the bright – I'm going to be the optimist here and say the fact that Vista now has been out for, what...

Steve: Couple weeks.

Leo: ...couple weeks, almost three, well, two weeks, and has not had any major exploits that are specific, that are particular to it is encouraging. I mean, we're on this countdown clock, frankly. The minute Vista shipped, I've been waiting to hear of that first Vista exploit.

Steve: Yeah, they will happen. There's no doubt.

Leo: I'm saying half full; you're saying half-empty. That's just the way it is.

Steve: Well, for you to be saying half full about Windows security, Leo...

Leo: That's pretty surprising, isn't it.

Steve: That's a real change.

Leo: I think it might just be that I want it to be true, you know?

Steve: I got a really neat note that I just wanted to share with our listeners about – this is a different kind, sort of, about SpinRite. We have an avid Security Now! listener named Russell Gordon, who has been listening to the testimonials that I've been sharing with our listeners and got to the point where he was looking forward to his hard drive crashing.

Leo: He wanted to use SpinRite.

Steve: He did. I just love this note. So he says: "Dear Steve, I'm an avid listener of your Security Now! podcasts and..."

Leo: That's so funny.

Steve: Oh, just listen, it's so great, it says: "...and have always heard you and Leo reading the SpinRite emails that you get. When I would hear those stories, I would think to myself, one day I'll be needing SpinRite. As a matter of fact, after hearing the stories and how much success people were having with SpinRite, I was almost eagerly anticipating getting to use it. Well, yesterday was my day. I'm a controls engineer who programs industrial computers called PLCs, Programmable Logic Controllers. Those are RISC-based processors that are extremely reliable and hardened for the industrial environments they reside in. Yesterday I was at one of my customers' sites, which is a brewery in Texas. I'd been working for several days on upgrading one of their programs to change the way the beer was being filtered." That sounds like a good idea.

He says: "Yesterday afternoon I started getting these Windows messages about how memory was corrupt, and I noticed that my hard drive light was staying on a lot more than usual. At one point I had to power the laptop down and turn it back on due to unresponsiveness. When it started to boot up I got a message saying something to the effect of ntoskrnl" – the NTOS kernel – ".exe is missing or corrupt. I rebooted again, and it never even got far enough to tell me what was wrong. Dead laptop." He says: "I knew it was time for me to purchase and

download SpinRite."

Leo: Lucky boy, he gets to use SpinRite.

Steve: He said: "What was interesting to me was that, after hearing all the stories, I was not even the least bit worried. For some reason I was in the mindset of, 'Bring it on, baby...'"

Leo: We do not recommend this, by the way.

Steve: Yeah, don't go dropping your laptops just so you have a reason to run SpinRite. Anyway, he says: "...because I knew I could probably resolve the problem with my drive. It was being recognized by the BIOS, so I knew that SpinRite would have a good chance of helping me out. I purchased and downloaded SpinRite and started it running in Level 2. It found quite a few problems with my disk, and all but one of them was recoverable. I let it run overnight and came in the next morning and it was complete. I pulled the floppy out of the laptop and rebooted, and there was my Windows again, running perfectly, just like it should be. Thanks to SpinRite, Texas will still have a good supply of Shiner Beer."

Leo: And filtered, no less.

Steve: Well-filtered Shiner Beer.

Leo: Wow, that's really great. Another happy story.

Steve: I just got a kick out of that. It was like, bring it on, baby.

Leo: SpinRite, of course, is Steve's disk recovery and maintenance program, available at GRC.com. Before we get into the show – are you ready?

Steve: Yeah.

Leo: You're ready. Are you prepared? I would like to mention our sponsor, Astaro, because Version 7 of – this is good news – of the Astaro Security Gateway is here now. And of course Astaro makes this great security software, and they've been a sponsor of the show for some time. And we are really happy to have them as a sponsor. But let me just talk a little bit about Version 7 because there are a number of new improvements, including transparent email encryption and decryption. You know I'm a big proponent of that.

Steve: Well, and Leo, let me just interrupt you to say I think that is so cool because it means that the machine that is your gateway machine does all the encryption and decryption.

Leo: For you.

Steve: So inside your network you don't have to have any of that stuff, and you don't need to see any of that stuff.

Leo: You don't have to be aware of it. I think this is great. You can use S/MIME or OpenPGP standards. Inbound email is automatically decrypted, too. So it's just really a great solution, completely automatic. They also now have secure remote access on their VPN via SSL, which as you know makes VPN so much easier. Let's see, what else? In fact, I think it's the only appliance on the market with VPN solutions that are SSL, IPSec, L2TP and PPTP with SSL VP. I mean, they really cover the waterfront there. Can be clustered for scalability so as you grow, so will your Astaro Gateway. In fact, you can even get up to ten clustered together without load balancers, it does it automatically. And they are going to continue to offer the home use package, but the v7 package will now be free of charge – get this – and will include the base license, all subscriptions, and Astaro up-to-date, but it'll be limited to ten IP addresses or ten users or 1,000 concurrent connections, plenty for a home user. So you don't even have to purchase the home user subscription anymore, the one that we were talking about for 79 euros. You get that for free. These guys are the good guys. Aren't they great? You can download Astaro v7 for free on the website, or call them for a free demo unit, Astaro.com. I am just really pleased that they've made that free now, including the subscriptions. That just makes it so much of a great product for everybody.

Now, let's get to our topic of the day, and I think it couldn't be more apropos, the use of bot networks. And you're going to talk about one particular use of bot networks.

Steve: Yeah. Okay. The problem is, and we've talked about this before, that bot networks are one of the new ways of disseminating spam. So when these little bots that are running on, by some reports, what is it, 150 million machines, one quarter of all PCs, when they're not being used to go blow some website off the 'Net through denial-of-service attack, they are being used, and it's been confirmed, for sending spam. They're like sources of spam. A couple times we've received some complaints from people saying, hey, you – GRC, Steve Gibson – sent me this spam. And it's like, I guarantee you I didn't send you any spam. And I'll scroll down, and they'll have attached the spam. And it's, you know, some horrible blob of keyword stuff meant to get through the spam filters. And then typically it'll have some binary attachment which is, you know, some evil thing. So it's not only spam, it's probably malware that is being sent. And something in the headers has led them to believe that I sent it.

Leo: Often it's the from address is xyz33@grc.com. I'm getting those to TWiT.tv all the time. And I hear from people all the time who are saying, they're using my address.

Steve: Yes, exactly. And so what the spam generators have is a whole list of reputable websites, like mine, like yours, like many others, you know, CNET and Amazon. And so they have a huge list of reputable websites, and they'll randomly pick those websites and spoof them as the source for the spam they're sending out in order to get through people's filters, in order to basically try to steal the credibility which is built up by good domains and ride their junk through on that credibility.

Leo: I wonder if that works, though. I think antispam solutions aren't fooled by the return address.

Steve: Well, no. It takes – actually it takes more than that. But in the samples that I have been sent from spam that has been sent in GRC's name, the headers do provide some reason to believe that it came from us. And that's what I wanted to talk about. I wanted to talk about how the headers can – well, there's one mechanism in email which is it's very cool, and it turns

out it's very useful for tracking down the true source of spam that is received. And so this is of interest to listeners because if they get some stuff and want to understand a little bit more about the path it took from wherever it was originated to their machine, that information is in the headers.

The idea is that email was originally designed with what's called a "store-and-forward model," where you would have so-called post offices, and those would receive email, not only for their own recipients, but actually technically for anyone. That is, in the old days, in the good old days of the original design of the Internet, before it became so subject to abuse, you'd have an SMTP server, Simple Mail Transfer Protocol, SMTP, which would be out on the 'Net, listening on port 25 for anybody who wanted to connect to it. And it would accept a connection coming in on port 25 and say hi there, identify who it was, and accept email.

Now, if the email was bound for its own domain, its users, it would store it locally, and then the users would use a different protocol, originally something called POP, the Post Office Protocol. That's what clients use to hook up to their POP server. And in the case of POP it's on port 110. They would hook up to it in order to retrieve their mail. IMAP is the other now arguably as popular, if not more, protocol that essentially does the same thing, but it's more recent and has additional features.

In the case, though, that this SMTP server received email not bound for one of its users, it would still accept it and say, oh, this isn't for me, but I'll send it on toward its direction. So it would look up, using DNS, it would look up what's called the "MX records" in DNS – MX stands for Mail Exchange – and find the IP address of the SMTP server for that mail's intended recipient, and it would connect to that server and forward the mail. So it was sort of a go-between.

Well, this quickly became exploited by spammers. And a server operating like this is known now as an "open relay." Because it is a relay point, it will accept email intended for people other than its own users, that is, other than its own domain, and happily forward it on. So spammers quickly discovered that they could simply dump their spam on other people's SMTP servers and get them forwarded on their behalf. Well, one of the consequences of this sort of store-and-forward model was realized by the original designers of the email system. And this is, again, this is part of, you know, the original brilliant conception of all of this. They realized that it would be possible for email routing loops to exist in this store-and-forward model, that is, they deliberately designed SMTP servers to be friendly back when the Internet was all just good guys. And so you might not be able, for example, to send email directly to its destination. You'd have to stick it on an intermediate server. And then it would try and try and try, for days in some cases, to ultimately move the email to its final destination.

The problem is, if DNS was configured in erroneous ways, it might be that an SMTP server would send email through a destination that would accept it; and then, in getting ready to forward it on, it might send it back to the server it just got it from. Or it could send it to a third server that would send it back to the first server, or any kind of round-robin looping like that was possible because the protocol itself didn't disallow that from happening. So you could see how you could end up with a piece of email just being handed off among servers in a particular messed up configuration of DNS.

So the designers, recognizing this problem, said okay, we need a way of tagging email such that, if an email comes back to a server that it has previously sent, it will know that, wait a minute, I've already seen this mail. I've already received this mail before. We've got a problem here. And it will not then continue to send it on. Instead, it'll send an error back to the email's apparent sender, saying this mail cannot be delivered as addressed.

Leo: And that's to prevent just a loop, right, an endless loop.

Steve: Exactly. And so the way this works is, any incoming email, any email received by an SMTP server has a new header appended to the front. And this is important for people to visualize, that there will be a bunch of so-called email headers. It'll say "To:," "From:," "Subject:," "Date:." Those are just lines with a word, followed by a colon, followed by data. And those are the email headers. What the receiving server does is it adds its own header. It appends it or pre-appends it to the front, to the top of the email as the first header now on the email, and that starts out saying "Received by." So that received-by header is a way of establishing that it has received this piece of mail. It identifies itself.

It also identifies, and this is really cool and critical, the IP address of the connection which it accepted. That is, it accepted a connection coming in, as I said before, on its port 25. Well, that means it has a remote IP which we know cannot be spoofed. As we said early on in Security Now! episodes, TCP, because it requires this three-way handshake, that is, in establishing the connection packets have to go back and forth three times. That validates the IP of each endpoint in the two endpoints of a TCP connection. So the IP that it logs in this received-by header cannot be spoofed. That is the IP of something. It doesn't know what, necessarily, but it's definitely the IP of the sender of that piece of email. So this received-by header gets added to the top of the email.

Well, if the email is bound for its own domain, for example, if GRC.com received this piece of email, it would say, okay, it would just store it. And when one of the GRC people, one of the clients connected, it would say, oh, got some more mail for you, and I would receive it. If I look at the headers in my email, I will see only one received-by line at the very top of the list of headers, which was pre-appended there by my server that received it at GRC.com. But say now that in the old happy days of open email relays this email was not intended for GRC, my server would then – the GRC server or the receiving server would then forward it to its destination, at least in that direction, towards its intended recipient. The server that received it would add its own received-by header, and now that one is on top of or in front of the first one. If that server forwarded it to a third server, that third server would again add its own received-by header.

So the point is, every hop that email makes, there is a received-by header always added to the front of the email. And what that does is, it creates basically a trace. It allows anyone looking at the headers to know what path that email took as it got towards them. And in fact, even though out on the public Internet you no longer have open email – you very rarely have open email relays because they're quickly found by spammers. The people who are looking at spam quickly identify open email systems and will notify you that you've got that problem. So it's generally not the case that public servers are open relays.

However, inside of networks, like inside of Yahoo! or inside of Google or inside of Cox, you know, in any large organization you may have multiple email servers. And so it's very often the case that email will be forwarded internally through multiple hops on SMTP servers. Which is why users – oh, and also Hotmail, for example. Which is why users who look at email headers will often see multiple received-by header lines. You would normally not see that, for example, if some remote user had sent email directly to its target server. You would normally see typically two received-by lines. For example, a user in the Cox network would – their client would send mail to the Cox SMTP server, which would add its received-by header line. The Cox server would then send it, for example, to GRC. My server adds its received-by header line, where the mail then sits until I receive it. So I would see, in the normal case, two received-by headers because that mail was received by two SMTP servers, the originator's SMTP server and the recipient's SMTP server. Sometimes you'll see many more.

Okay. So the way spoofing is done is the first time you see it, it's pretty clever. Some bot somewhere wants someone to believe that GRC has initiated, is the actual sender and initiator of a piece of spam. So the mail they send out has a fake received-by header stuck on the front of it. That is, most email that is sent out doesn't have a received-by header because that's only added to the front of mail by receiving SMTP servers. But remember that these headers are, like, they're in reverse order. They're stacked so that the first one is the most recent recipient, and then they go back in time. So there's nothing to prevent an email from being spoofed and

generated with a fake received-by header at the beginning of the mail, so that when that bot then sends it to some SMTP server, that SMTP server will stamp its received-by header in front.

But that means now that, for example, the spoofed GRC received-by header is underneath it, and it appears that someone at GRC sent the mail to our server, which created that stamp, and then our server sent it on. So basically what that does is it creates a spoofed piece of email. At first glance it hides the fact that some end user somewhere sent the mail. Except it turns out that that really falls down under closer scrutiny. Remember that I said that the received-by header contains the IP address that the receiving SMTP server connected to to receive the mail. Well, GRC's server, SMTP server, is at a known IP address. And, you know, it's in the GRC.com network. But since a spam-generating zombie connected to some other SMTP server trying to fake it out, the received-by header, which was added by that server, won't show a connection from the GRC.com IP. It'll show a connection from the IP of the infected computer.

And in fact, when I've responded to these spoofed emails that people have occasionally sent, saying hey, GRC.com sent me some spam, I take a look at the email they sent, and I write back a response explaining to them how, yes, it does look like GRC originated this because we're the lowest-down header in the received-by stack. Except that if you look at the second one up, you will find the IP address, not of GRC, because our server would have connected to the second SMTP server. Instead, you see the IP address of the zombie machine. And so on...

Leo: So you can actually tell who these zombie machines are.

Steve: Yes. And in fact, that's the cool thing...

Leo: That's not spoofed.

Steve: And here's the point, Leo, it cannot be. Because the receiving SMTP server always adds the IP of the other SMTP server that it's accepting the email from.

Leo: And it's not possible for you to have a rogue SMTP server that fakes those originating addresses?

Steve: Remember that it's actually using the TCP connection IP. And there's no way to spoof that. And so what's been fun is when I've responded to these people, and I've shown them how to read – you know, I figure I ought to try to educate them since they're sending out erroneous complaints about spam. I've educated them. I explain and show them this is the IP that actually connected to the second SMTP server, not GRC's IP. And then I'll use my very favorite site that I've shared with you, it's DNSstuff.com, it's www.dnsstuff.com. If people will go there and scroll down a little bit, there's just a beautiful array of very easy-to-use tools, it's web-based, that allows you, for example, to perform what's called a "Reverse DNS" lookup. I'll take the IP address out of the second received-by header, drop it into the DNSstuff, and often get and know a lot about the infected zombie computer. I'll know whose network it's in. And not only do I have the IP address, of course, but I actually get more information from the Reverse DNS. Is it DSL? Is it a cable modem? Probably it's a semi-static IP. And so I'll return that and say to the guy, this is the source of this spoofed email. Not me, but some random zombie in Norway or in China or in Russia or wherever, one of these 150 million...

Leo: I just did one, and it's in Mexico. Interesting. Very interesting.

Steve: Yeah, yeah. It's very cool.

Leo: But it's posing as TWiT.tv. But it's really a relay. Very often it's in the U.S. Many of these zombies are in the U.S.

Steve: Actually 20 percent of spam is currently originating from the U.S., 20 percent from Russia – oh no, 17 percent from Russia, and I think 20 percent from China is the most recent stats. It turns out that 67 percent of all email is now spam. So two thirds of email is now spam. And, you know, obviously no end in sight because this continues to be a problem. But at least if our listeners are interested in taking a look at some of their own spam, they'll be able to see now by looking at, especially if it's spoofed, if they see, you know, multiple received-by lines, you can normally, if you take a look at the IPs in the received-by headers, since they are not spoofable, those are the actual IPs that each SMTP server in line received, had a connection to. And often you will find that they are end-user IPs, and those are infected machines.

Leo: They're zombies.

Steve: They've been taken over.

Leo: Yeah, the one I'm looking at looks like it was in fact an end-user machine. Fascinating.

Steve: Now, I'll close by talking about one last – and this is pretty quick – one last antispam measure which is very cool. And this is what GRC uses, and many others are, too. And that's the SPF, the Sender Provider Framework, or Sender Policy Framework. It's got multiple names, and the name has mutated a few times. The idea of that is that SPF is a very cool antispam measure which is also very useful for anti-spoofing. GRC has added a text record in DNS. In DNS you have different types of record. An A record in DNS lists the address of the machine. I mentioned MX records earlier, which give the IPs of mail exchange servers for that domain. And there are many different types of records.

Well, one type is a text record where it's basically freeform. You can put anything in there that you want, and anyone who wants to look it up can essentially ask for GRC.com's text records. Well, if they do, they will find one which is formatted with the protocol of this SPF, this Sender Provider Framework. What it does is it declares the valid IPs of email originators for GRC. That is, in there I say that any email coming from GRC that's valid has to be in this range of IPs. And it's basically GRC's network.

What that means is, since it is not possible to spoof a TCP connection, if a receiving SMTP server had somebody on the line, essentially, that is, accepted an SMTP connection, they would know the IP of the connection originator. If that person then says, hi there, I'm GRC, I've got an email for you, well, that receiving SMTP server can do a lookup in GRC's DNS records for this SPF record and determine what GRC's actual range of valid email originating IPs are. And if the connection it has from somebody claiming to be an authentic GRC email sender isn't in that range, it absolutely knows it's fraudulent, that it is spoofed, and can simply drop, I mean, they could drop the connection, they could tarpit it, they could do whatever they want to. Certainly just, you know, discard the email.

Leo: And that won't have any false positives?

Steve: No, and that's the beauty of it, it cannot false positive. There are some problems with this in relaying because, if GRC were to need to relay email to a third-party server for some reason...

Leo: I'll give you an example. I talk to people who have businesses. They have home addresses that they want their business email forwarded to, and they want to respond from that home address using the business address as the reply-to.

Steve: Exactly. Exactly.

Leo: And that would fail.

Steve: Exactly. And in fact my tech support guy, Greg, who's in Phoenix, he uses Cox and has in some cases tried to send GRC email out through his Cox system. And oftentimes he'll get a bounce-back because this system is beginning to come online, this SPF technology. Of course, it's only as good as it is adopted. And it's trivial to add these SPF records to DNS. It's so simple to do. In fact, the site that is promoting this even, you know, it has a little online web form where you're able to fill it out, and it builds for you the exact line of text that you need to put into your DNS. So...

Leo: Microsoft has a little form like that, as well.

Steve: Yeah, that's very cool. Anyway, so it does absolutely prevent spoofing. But it does have the problem, exactly as you say, Leo, that if you are trying to route email through some other server and not going direct, then you've got a problem.

Leo: You could add an SPF record if you wanted him to do that.

Steve: That's exactly right.

Leo: To say, oh, he's an accepted sender, or a Comcast accepted sender.

Steve: Yes. And in fact, well, I wouldn't – I'd be afraid, for example...

Leo: You wouldn't want to open Comcast, obviously.

Steve: Yes. I wouldn't want to whitelist all of Comcast's network. But for example, you know, Greg's got a cable modem. Of course he's got a NAT router, which tends to anchor his IP firmly. Also the SPF format is flexible enough that you can use a domain name. For example, in general, stuff coming from us comes from an IP at client.grc.com. In fact, I think in every case it's client.grc.com. Any outgoing traffic comes from that IP. So you are able to use domain names in that record. So, for example, if we really wanted to do it this way, Greg could have his router set up using a dynamic DNS service, which is giving his IP an automatically updating DNS record. I could put that in GRC's SPF record, in which case his IP would always be a legitimate source for GRC email.

Leo: So you only have to change it for the domain that's the source of the email. You don't have to change the MX record for his sending SMTP server. He doesn't have to have Comcast's cooperation, in other words.

Steve: Right. The idea would be that he would send email to Comcast. If they were using SPF, they would check to see whether he's a valid generator, a valid sender of email from GRC. So they would check GRC's DNS.

Leo: They'd go back and ask you, yeah.

Steve: Exactly. They come back and ask us through our DNS. And we'd say, yes, that IP can, is authorized to generate GRC email.

Leo: So if I want to send email from Leoville.com via my FastMail account, I wouldn't change my FastMail's records, I would change my Leoville.com's records to accept anything from that IP address.

Steve: Right.

Leo: That makes sense.

Steve: It's a very cool...

Leo: It also raises the issue that, frankly, another way to kind of prevent these bots, now you tell me if I'm wrong, would be for Internet service providers to validate that the return email address was in fact in their network.

Steve: Well, yes. Actually there are many things that ISPs could do. For example, a bot, in order to do this, a bot is in a network where it is able to connect to an external SMTP server. Okay, you could argue that – and in fact, ISPs are often now blocking port 25.

Leo: For that reason.

Steve: Exactly, for that reason. So you are unable to get a port 25 connection outbound. The only thing you can do is plant your email on your ISP's server. And that completely shuts down these bot networks. Although what the bot networks then do is they just dump all their spam on their own ISP. And often the ISP forwards it on.

Leo: So that's why, by the way, I think the U.S. is now only 20 percent. I think it was a much larger number for a long time because some of the big, high-speed Internet service providers refused to block port 25. Now they're doing so. And so now it's going offshore to Mexico, Romania, Russia, and these other countries, where ISPs I guess are not so enlightened.

Steve: Right. And in fact there are various sorts of workarounds for getting through, of course, port 25 blocking. If you just allowed an SMTP server to run on some other port number, then you'd just have to configure your client. But, you know, that would work.

Leo: Right. By the way, that's OpenSPF.org was where that wizard lives, if you wanted to do that. And Microsoft has one as well, although I don't know, they don't use – do they use SPF? That's part of the reason this has become such a spaghetti is that there are competing standards.

Steve: Yes. Microsoft had one that they did not want to leave as completely open. They wanted to get intellectual property protection on it and then say, well, but we'll license it for no cost. It's like, wait a minute, you know, the open source community said...

Leo: No.

Steve: No. And so Microsoft has substantially muddied the waters by sort of trying to do their own. The other very popular one is a more sophisticated approach known as domain keys. And was it Yahoo! that was the early domain keys adopter?

Leo: I think Yahoo! Mail was going to use that, yeah. Hotmail wanted USPF and domain keys...

Steve: Right, right. So we're seeing those. But anyway, for people who are curious about the spam they receive, and certainly for people like you or I or anybody else who's a high-profile, good reputable source, we end up being the target of these sorts of spoofing attacks. It turns out the good news is it's very possible to demonstrate that we were not the source; and there is no way, if you know how to read those headers, that you can be fooled.

Leo: Yeah, would you send me the boilerplate that you use? Because I would love to use that. I get those comments fairly frequently. Actually more commonly I get people asking me, you know, I'm a businessman, I have a domain, why am I getting all this bounced-back spam? Who is using my domain name? And I have to explain the same thing to them. And apparently there's nothing you can do about it.

Steve: No.

Leo: As long as people don't require domain authentication, it's going to continue.

Steve: Right, well, and it is interesting, too, that, I mean, we notice that we are getting the IPs of infected machines. On the other hand, there's 150 million of them, and they're all over the globe, so what are you going to do?

Leo: Right, right. Well, fascinating subject, and just one of the ways botnets are used. I do hope we'll at some point talk about how botnets are used to attack systems like, as they were recently, used to attack the domain root servers. But that's a topic for another show.

Steve: That'll be a good one.

Leo: I think we're all done here, Steve. I will see you next Thursday. Have a happy week and a happy weekend. And thanks to our favorite security maven.

Steve: We'll be doing Episode No. 80, which of course is going to be one of our big Q&A episodes.

Leo: Oh, boy.

Steve: So we'll have lots of good questions, and we'll address the answers.

Leo: Go to GRC.com, that's Steve's site, and you can ask questions, pose questions there. It's also where you can find show notes and 16KB versions for the bandwidth-impaired, and of course transcripts so that you can follow along, read along as Steve speaks, which is often useful, frankly.

Steve: Oh, also I should mention there have been some questions, people saying how do they send questions to me for Security Now!. So I want to remind people, just go to GRC.com/securitynow. Go to the very bottom of the page, and there's a web form where you can provide as little or as much information about your identity as you choose. I don't really care one way or the other, but we do like to know names and locations in order to read these things on the air because it makes it more...

Leo: We'll give you credit, yeah.

Steve: ...know where people are. And that's how you can submit questions to me.

Leo: Yes. Steve doesn't take personal mail. He bounces it right back to you. All right, Steve. GRC.com's the place for SpinRite also, Steve's great disk recovery and maintenance utility, and now SecurAble, his new program that tests how secure your system can be. A lot of interest in that.