Listener Feedback #189

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-459.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-459-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about Patch Tuesday, the updates, and whether they'll affect XP. And we'll talk a little bit about, well, your questions and Steve's answers. We've got eight great ones, coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 459, recorded June 10th, 2014: Your questions, Steve's answers, #189.

It's time for Security Now!, the show that protects you and your loved ones online with the help of this guy right here, Mr. Steven Gibson of GRC.com, a security guru, the author of SpinRite - the world's best hard drive maintenance and recovery utility - the guy who discovered spyware, coined the term, and wrote the first antispyware tool, and I can go on and on and on. But let's just say hi to Steve and welcome to the show, a Q&A episode.

Steve Gibson: Yes. Yes, we've got 189. And I have to say that, as I was going through the mailbag, not surprisingly, almost everything was about Net Neutrality from week before last or TrueCrypt from last week. And there wasn't really much more to say about those things than we'd already said. I managed to find some other things. But my sense is...

Leo: It just shows you those are hot-button topics that everybody's interested in; right?

Steve: Yes. And they are admittedly sort of political. And there's, like, gray areas. And, boy, I mean, it's not like how many bytes are in a packet, which we all pretty much can

agree on. It's so loosey-goosey. And so, I mean, opinions were from one end of the spectrum to the other, but there really wasn't anything significant that I had, that I saw that seemed uncovered from...

Leo: Yeah, you know, I'm sure it was mostly, "Steve, you're so brilliant; Leo's such an idiot," and like that.

Steve: Or the other way around.

Leo: Or the other way around. And 'cause we disagreed a little, not disagreed, but we had different maybe opinions about the TrueCrypt, whether to keep using TrueCrypt. And Brett Glass and I disagreed somewhat vehemently, although Brett's been very gracious about it and has continued the conversation on Twitter.

Steve: Yeah.

Leo: With a lot of people. So that's good. Certainly good to hear other points of view.

Steve: Well, we have - this is our second Tuesday of the month. And it is a Patch Tuesday. And there are two interesting critical vulnerabilities, and we'll need to see whether they reach down and affect XP. It's not clear yet. We'll talk about those. I did spend some time looking at Google's browser-based PGP, so I have some comments about that. We have more trouble with OpenSSL that came to light of course the day after we did last week's podcast, when all these things happened like TrueCrypt and so forth. We have the first Internet registry to hit critical levels of remaining IPv4 addresses; some interesting iOS 8 privacy news; some typical nonsense from TheRegister.co.uk; a couple things about network congestion that I thought were interesting. And of course it's a Q&A, so we've got some feedback and interactions from our listeners.

Leo: A jam-packed show today.

Steve: A great podcast, as we generally have.

Leo: Yeah, we're always - lots to talk about when it comes to security, despite your concerns early on that we would...

Steve: Boy, we sure didn't run out...

Leo: ...run out of things to say.

Steve: And every so often I see people tweet, Steve, you know, you and Leo should do

this three times a week. How about a Monday, Wednesday, Friday?

Leo: There's plenty to do. But I - yeah.

Steve: Yeah.

Leo: Steve's got other things to do, too, you know. He's not just sitting there waiting to get on the air. All right. Here we go. It is Patch Tuesday.

Steve: So, yeah. And we'll have to see whether anything happens from this. This is potentially bad news for XP. At this point, none of this is in the wild. These are vulnerabilities which have been found in Windows. They affect all, across the board, Microsoft operating systems. There's a problem in the, well, there's a problem in both desktop and server OSes that is remote code execution in the Unicode scripts processor, which is - it's a DLL, usp10.dll.

However, in order to exploit it, you have to have WebDAV, the Web Distributed Authoring and Versioning system running; and you have to have your ports 139 or 445, which are the traditional Windows filesharing ports, exposed to the Internet. Well, routers block that. ISPs block that, typically. And the firewall, Windows Firewall blocks that. So this doesn't seem like there's any way to access this vulnerability from behind all of those wrappers. But we'll have to keep an eye on it.

The second problem is an image-parsing vulnerability in GDI+. And of course these are a problem because once upon a time Microsoft moved GDI from user space down into the kernel when they wanted to increase performance. And so now this is kernel-resident code where there's a known problem. So if this is going to get exploited, the path will be that these updates come out, and people will reverse-engineer an exploit from the difference in the code between this month and last month and see if they could turn that into something that attacks XP.

Computerworld had an article that I originally had in my notes, but I thought, eh, it doesn't quite make the cut, which was saying, you know, where's the XPocalypse? And, you know, it might be upon us. This could get leveraged into that. So we'll have to keep our eyes out and see if that in fact happens, in which case I'll be the first to say, okay, XP is no longer safe to use because someone actually did find a way of exploiting those systems.

Leo: But you've been pretty clear that you feel like it's safe to use XP, at least for now; right?

Steve: Well, if we take the position...

Leo: You know what you're doing. You know what you're doing.

Steve: Yeah, if we take the position that everything has bugs, I mean, we're about to talk about a serious man-in-the-middle attack on OpenSSL that's been there for 15

years.

Leo: Oh, man.

Steve: So it's not, I mean, and all the evidence demonstrates that everything has problems. So it's the known problems which are exploitable that are a concern, rather than just, like, I mean, if we were worried about using any software, we ought to just completely disconnect from the 'Net and say, okay, well, I refuse to use insecure software. All software has some potential insecurity, so I can't use any software.

Leo: Kind of the point the TrueCrypt guys were making. Right?

Steve: Exactly. And my position is, since we live in the real world, we're going to - and there's a choice between being proactive and reactive. And the notion of open source, it enables proactivity, but no one really seems to take advantage of it. You could be proactive by reading the source and find the problems and fix it. But in the real world, again, we end up being reactive. We discover a problem, and everyone scrambles around to patch it. So that's the world we live in is unfortunately a reactive world. And so, you know, that's where I'm coming from.

Where, for example, with TrueCrypt, as far as we know, there are no problems. When we learn differently, then it becomes insecure because what was always there we then find out about. And the point is that that's when probably the bad guys find out about it. They probably didn't know, either. And the good thing is we're generally pretty good about discovering exploitation. So when something gets exploited, we're pretty good about recognizing, oh, something just went bump. How did they do that? And then the word goes out. You know, that's the whole intrusion detection system idea. So it's not an ideal world. But it is the one that we've ended up with.

Leo: I don't want people to think that they're - yes, of course we're reactive. But I don't want people to think that we're not paying attention. I mean, all good security, all good software is security audited; right? I mean, companies like Microsoft and Apple go through their, you know, they have audit processes that go through their code. They look for flaws proactively.

Steve: Nobody wants to have flaws. Yet I'll never forget Steve Ballmer dancing around the stage before the release of XP, saying it's the most secure operating system we've ever produced. It was the least secure operating system they had ever produced. And at the time I said you cannot declare that. No one can declare that something is the most secure blah, whatever. It's history that judges that, after the fact, for exactly this reason. And that is that something really simple, you know, you make a cube, and you say this is the most perfect cube ever made. Well, it's simple enough that you can probably stand by that assertion. But an operating system is so complicated and has so many moving pieces that there are going to be problems that no one has found. And here we are later, still finding them in an operating system which is so old that its support has been discontinued after a decade. We're still finding problems with it.

So, yeah, all we can do - and that's why every month that we talk about this, we run through these patches, and I say to people, update your system because, now that we

know about these problems, now, I mean, even if it's been latent, even if the problem's been there for 10 years, if no one knew about it, then, okay, we wish it weren't there, but it's not hurting us. It's when it becomes public that it has the potential to hurt us because bad guys find out at the same time we find out. And so, in fact, this ridiculous, oh, no, I'll stay on script here, otherwise I'll get myself tangled up. But this ridiculous story about The Register, I just - it cracked me up. But anyway, so again, second Tuesday of the month, everybody should update.

Now, I haven't yet done it, but I do have a tablet running Windows XP SP3. I'm going to add that registry key and just see if it updates itself. I ought to have a system which is using that hack to pretend to be an embedded XP and just keep an eye on these things because these patches are probably available for XP Embedded, and anybody can turn their XP into one that looks like that.

Leo: You could be the canary in the mine that lets us know when it's working, when it stops working, if it does anything, if it breaks anything.

Steve: Yeah.

Leo: And you don't mind losing that tablet.

Steve: No, it's just...

Leo: Does it go online?

Steve: No, I have it, yeah, I have it plugged in. It's got my stamps and an electronic scale. It's my little postage station.

Leo: Good, that's a good use for it.

Steve: Yeah.

Leo: Apparently, Considerate in our chatroom, Considerate1 in our chatroom has tried it. And this Patch Tuesday there is an update via that, so...

Steve: And what was interesting was that Microsoft, I looked for it, had absolutely no mention of XP Embedded in their security release. Now, I wonder, they have to be publishing that somewhere. So maybe there's a different link that we're not looking at right now for, like, the people who are licensed XP Embedded XP users. They've got to know that they're getting patches and so forth. So there's probably a channel somewhere for them. But anyway, thank you for the feedback from the chatroom. I'm not surprised. I mean, this is what we believed. And, what was it, is it five years? I think it's through 2019, if I remember right, another five years' worth of...

Leo: This is not, and we should be clear, this is not - there are two things when you say the word "embedded." It could be the embedded version of Windows. This is not that. This is Windows XP for embedded systems, which is the same bits as Windows XP.

Steve: Correct.

Leo: And so you change the registry to be saying, no, no, I'm not plain XP. And in fact you could make a strong case it's a point-of-sales system you've got there. It's Windows XP for embedded systems.

Steve: Good point.

Leo: All it does it postage.

Steve: Yes.

Leo: You kind of do have an embedded system there. But that's different from Windows Embedded, which is a different operating system.

Steve: Okay, yes, correct. And Embedded was really cool. I looked at it a lot, actually, once upon a time, considering it as maybe the platform for SpinRite. But I'd still be paying licensing fees, and so it just didn't make any sense.

Leo: Yeah, you put SpinRite on FreeDOS, which is a free DOS.

Steve: Now, yeah.

Leo: And it's lightweight. It seems silly to have Windows when all you need is a command line.

Steve: Yeah, yeah. Well, I just liked the idea of a strong platform. And anyway, the point I was going to make was that it is - what they did was they broke it up into tiny pieces so you're able to - they componentized it. So you're able to make a much, much smaller footprint in terms of, like, ROM for Windows by just discarding all this nonsense that Windows brings along. I mean, all the stuff that an end-user would need, like Internet Explorer that you can't get rid of unless you use the embedded version, where you just turn off the checkbox, and it builds one without it. So anyway, that was what that was.

So there's been some controversy about Google's so-called "End-to-End" is the official name of this thing, which is their browser-based PGP for email. I'm impressed by everything I've read that they understood the danger and have, like, faced up to the fact that some people consider the phrase "secure JavaScript crypto" to be an oxymoron.

That is, there's a famous blog posting, "JavaScript crypto considered hazardous" or something, or malignant. And the problem is, I mean, and it's an understandable concern, and that is that you are inherently downloading code that the browser runs which is - which has to be secure in order for it to do its job. And it just makes old-timers uncomfortable when the browser is connecting to a server to download code that has to be secure. But with sufficient protection, and with a real understanding of the dangers, Google makes the point that they're going in with their eyes open, they're not wanting to hurt anybody, and they understand the risks.

For example, critics have said, well, you just don't have enough control in JavaScript. For example, you're using a virtual machine to interpret your JavaScript, and it's doing memory management. Well, memory can be sensitive. You can have stuff in there that you don't want to let loose. And the Google guys say, yes, and we're protecting ourselves from that. And they say, well, wait, okay, what about timing attacks? You don't have any control over timing the way you do if you're in a lower level compiled language. And the Google guy says, yes, and we've done everything we can. And given the fact that it's submerged in a bunch of other stuff going on, we think timing attacks are impractical.

They've developed their own elliptic curve technology for this. It does not use RSA certificates, which is one of sort of the controversial points. And they've said, look, it's just - it's too time-consuming to generate an RSA set of keys in JavaScript. Elliptic curve technology is vastly faster and equally secure. So, and, let's see, right now GnuPG is at 2.1 beta, and it supports the elliptic curve keys that End-to-End generates. And I can't remember the other one, the other major - I think the Symantec version of PGP, the latest one also supports it. So older versions don't. Yet you can use their keys in End-to-End. You just can't create new ones in End-to-End.

So End-to-End, which again for clarification is Google's browser-based sort of built-in PGP, it'll only create elliptic curve public keys. But it will happily import and use traditional RSA keys with no problem. It just won't make its own. And the current versions of PGP, I guess it's when 2.1, when GnuPG gets out of beta, then they'll feel it's stable and safe and good to go, and it does support the elliptic curve crypto also.

So I see this as a move forward. For, I mean, there's probably a certain class of user where the barrier of just management, you know, the Glenn Greenwald effect, where it was so difficult for Snowden to get him to be using PGP that their interaction was delayed for a long time. If this existed then, it would have been trivial for him to securely decrypt messages from someone else. So even if it wasn't a full-time solution, if it was there for the times when you absolutely have to have secure, end-to-end encryption, then this is, I think, an interesting experiment at the very least. And again, it's one of the many things Google is doing that I think is just really terrific.

So SSL was back in the news this week. I think it was seven problems were found, disclosed, and patched. Several of them were DTLS. That's the UDP transport TLS. And I meant to do a search just to see who's using that currently. I mean, it makes sense to have a secure layer on top of UDP. But that's not been available traditionally. So it's been TLS on top of TCP, which is what all of our web browsers and web servers use.

So what was found was a couple different crashes, so-called denial-of-service attacks where you could put things in an infinite loop on the client. So, like, if you connected to an insecure DTLS server - and again, I meant to find out, like, what one would be because I don't know of any right now, but so this is sort of still in the theoretical end - there was a way that it could crash your client. It's like, well, okay, that's not good. But, you know, it's not the end of the world.

What was the end of the world was the major new problem found, which was, again, lots of caveats. Only if you had a vulnerable OpenSSL stack on each side of the connection, so you'd have to be connecting your OpenSSL-based client with an OpenSSL-based server, both vulnerable. Then a man in the middle who could intercept traffic could spoof some cipher change messages which, when inserted into the handshake at just the right time, could - and they sort of said coyly, they said, "create a key material downgrade." Well, yes, null keys, essentially. So that's quite a downgrade. You know, when I first read it, I thought, oh, okay. So they can, like, push you down in the security of the suite, the security suite that gets negotiated. No, it's worse than that. You can have null keys, essentially no encryption.

Now, the good news is no one is known to be exploiting it. The latest version has been updated and no longer has this problem. And it was Adam Langley's posting in his ImperialViolet.org blog. He blogged about it immediately. And looking back at the oldest source he could find, which was 10 years old, I'm sorry, 15 years old, he couldn't find any OpenSSL code older than that. But in the oldest one that was 15 years old, the problem was there. So it's always been there. And it's a very subtle protocol attack which has now been fixed.

So here's another example of subtle problems that exist in incredibly complex software which smart people find. And we hope smart people will find them and report them so they can be fixed and people can patch before the bad guys find out. I mean, this is the reality of today's model. Again, this is why, as I was saying last week, I'm so annoyed, distraught, really, over the way the legal system is imposing itself in this loop because it's not good if researchers cannot do this research. Having this process, this feedback loop where security systems can be examined for problems, where benign researchers then tell the people who are vulnerable about the problem in order to fix it, and not have them in danger of being sued as a consequence.

Because, I mean, all of our experience says, unfortunately, this is the way the system works. It's not the way we wish it worked, but it's the way it does work. And we need that feedback. And unfortunately, the legal system is really threatening. I mean, researchers can choose what they want to do. They don't have to do this. They don't have to expose themselves to legal attack. So if legal attack is there, if the potential is there, researchers will research something else. And our security system as it actually functions today will stop functioning in as good a way. And unfortunately, this leaves our systems a lot more vulnerable. So this is not a good direction that we're seeing.

However, an example of a really clever good direction, I thought - I don't know if you covered this in your last podcast, Leo, MacBreak - was the news that iOS v8 is deliberately randomizing its WiFi MAC addresses. Very cool thing. One of the things that has been known is that cell phone carriers with WiFi enabled - not the cell connection, the WiFi connection - are being tracked as they wander around. There have been department stores that purchased devices to identify customers from the MAC address which is broadcast by the WiFi in cell phones. There were some stories about weird things like recycling containers in the U.K. had hidden WiFi receivers that were being used to track people.

And of course, as we know, a MAC address is supposed to be a globally unique IP which is the - well, not IP, sorry, globally unique identifier. It's 48 bits divided into two 24-bit chunks. And one chunk is the manufacturer's ID, and the second chunk is a serial number unique within that ID. And so, for example, if you use a wire-sniffing tool, it'll show you the manufacturer of the ethernet adapter which is obtained from the first 24 bits of the MAC address. And the idea is that, on any Ethernet, it's the MAC addresses which are used for routing Ethernet packets, even if IP protocol is being carried by the

underlying Ethernet packets. The IP is like a higher level addressing. The actual physical address is the MAC address. And so those are traditionally fixed. Every device manufactured by every manufacturer has a unique MAC address.

So what a researcher discovered in the last couple days about v8 of iOS, a new feature in v8, is that when the device, when an iOS 8 device is not associated with an access point or a hotspot, that is, when it's just in that mode, for example, where it lists all of the ones that it can see and asks you if you want to connect, that requires WiFi Ethernet transactions. And for those, iOS 8 makes up a MAC address. It's not actually communicating.

Leo: I wonder if that's going to break anything.

Steve: Probably not. I don't...

Leo: It's not making an active connection, so it's not...

Steve: Correct.

Leo: I mean, we've known this for a long time, that all devices broadcast SSIDs; they broadcast MAC addresses. That's why MAC address filtering is ineffective.

Steve: Well, and that's what Google's positioning system that they famously got into such trouble for, I mean, when the little Google Bot is driving around your neighborhood, it's the MAC address of your router that it's logging because that's fixed, and that should be...

Leo: You could change the SSID.

Steve: Right. Right. So anyway, so the idea is that what this does is it just fogs your identity as you're casually walking around. All of those interchanges where it's just having a non-associated, sort of pre-association dialogue with the hotspot, it just uses a fake MAC address. It just randomizes them, which is just kind of a cool feature. I imagine that that'll be added to Android when the idea catches on because it's a sort of - it's a nice thing to do. And so it's nice, again, to see that Apple is thinking in this direction. Small change, but a privacy enhancement. And...

Leo: Of course their iBeacon's probably announcing their location pretty...

Steve: Exactly.

Leo: They have an alternative method.

Steve: Yeah, I saw some dialogue suggesting that, well, yes, but this is a way for Apple to say, ah, well, if you want to track our users, use iBeacon, which is our technology for doing that.

Leo: Yeah. That's good. I mean, it's nice to be able to turn that on.

Steve: Okay. So The Register. Just so full of it. The headline was "Redmond is patching Windows 8 but NOT [all caps] Windows 7, say security bods," they said, and then the subhead was "New tool checks differences, could lead to zero-day bonanza."

Leo: It's not just The Register. I mean, others are reporting this, too.

Steve: Well, yes. Well, is it others reporting it, or are they reporting what The Register reported?

Leo: Maybe they are. I mean, it was at a conference in Heidelberg, but nobody picked up on it until The Register did. So, yeah, I guess The Register gets credit.

Steve: Yeah. So The Register says: "Microsoft has left Windows 7 exposed by only applying patches to its newest operating systems." Okay, but they're not talking about XP here, they're talking about 7. "Researchers found the gaps after they scanned 900 Windows libraries and uncovered a variety of security functions that were updated in Windows 8 but not in 7. They said the shortcoming could lead to the discovery of zero-day vulnerabilities. The missing safe functions were part of Microsoft's dedicated libraries intsafe and strsafe [as in string safe] that help developers combat various attacks.

"Researcher Moti Joseph, formerly of Websense, speculated Microsoft had not applied fixes to Windows 7 to save money. 'Why is it,' he asks, 'that Microsoft inserted a safe function into Windows 8 but not Windows 7? The answer is money. Microsoft does not want to waste development time on older operating systems, and they want people to move to higher operating systems,' Joseph said in a presentation at the Troopers14 conference." And I was hoping that that was not his Boy Scout troop.

Okay. So here's what that is. Many people were worried and concerned about this and tweeted it. This is nothing. This is these 900 Windows libraries, okay, well, this is just the API Foundation, the function foundation offered by the operating system. And it's true that, as we know, many of the functions that programmers use can be used in an unsafe fashion. When we talk about buffer overrun, one of the common things that's done is a string copy, where you copy a string from, like, the URL into a buffer. Well, strings are typically terminated with a zero, a zero character, a so-called "null termination." So a simple-minded strcpy will copy every byte from the source to the destination, one after the other, until it hits a zero, the null terminator saying I'm at the end of the string.

So a programmer unaware of security will allocate buffer space often on the stack because in fact, in C, if you just declare variables in a function, they're allocated on the stack by default. That's how these problems occur. So then a hacker says, ooh, I'm going to give the guy a really, really long string, much longer than is reasonable, much longer than is, like, in the spec. And so what happens is the programmer, the insecure authoring programmer, allocates enough buffer for the string he expects. The bad guy

gives him a string that he doesn't expect and overwrites the amount of buffer allocated on the stack.

And unfortunately, since the stack is shared, not only with data on the stack, but return addresses, it's possible to put, when you overwrite the stack, to change the return addresses which have been stacked there, causing the function to go where it wasn't meant to go, thus buffer overrun and so-called "return-oriented programming," ROP exploits. How do you fix that? In the strcpy you add another term, which is size of destination buffer.

Leo: S-t-r-n copy [strncpy].

Steve: The original string copy didn't have it. The new one does. And so Microsoft has been adding new intrinsics, is what they're called, new low-level functions, and encouraging programmers to use them. So Windows 8 has more of them than Windows 7, which has more of them than XP, which has more of them than 2000. They're always adding them. And so they're not updating them because they didn't exist in Windows 7. And, yes, I mean, it's annoying that we're using a commercial operating system that Microsoft keeps obsoleting in order to generate upgrade revenue, rather than a non-commercial operating system, where they just fix everything and add things to it, and everyone gets the updates.

But that's not the world we're in. We're in a, you know, Microsoft is a for-profit organization. Consequently, if you want the latest and greatest goodies, you need to use the latest operating system. And programmers have to use those new tools. If you're just using Windows 7 code on Windows 8, then you're not using any of the new libraries. I'd be surprised, in fact, if that has any effect whatsoever because, if any programmers want their code to also run on the older OSes, to be backwards compatible, they can't use the new functions. And so in their libraries they say only use things from XP and earlier, which turns off Windows 7 and 8 improvements, because the developer wants it to run on all of the systems.

So, again, this was, first of all, complete nonsense from The Register, and also nothing to worry about. Nothing like they said, where Windows 7 was no longer being updated because Microsoft wants to save money. It's not being updated, true, because Microsoft wants to make money by selling Windows 8, which has more features, which are probably not being used. So big deal.

So one of the Twitter channels I follow in one of my accounts which I use for following, which of course is famously not SGgrc. People say, "You don't follow anybody." Well, I follow lots of people, just not there. The LACNIC is the Latin America and Caribbean, or Caribbean, depending upon where you live and how you pronounce it, registry. There was a mushroom cloud with a lot of red heat which is the icon for where the end of IPv4 address space is tweeted from. They posted an announcement that there were "No more IPv4 addresses in Latin America and the Caribbean." That's actually not quite true. The subhead said: "Latin America and the Caribbean have entered the IPv4 exhaustion phase." Turns out this is all phased. "The delay in deploying Internet Protocol v6 in our region is cause for concern."

So this was this morning. This was today that they said: "The Internet Address Registry for Latin America and the Caribbean, the organization responsible for assigning Internet resources in the region, announced the exhaustion of its IPv4 address pool and expressed its concern regarding the fact that operators and governments throughout the

region are delaying the deployment of Internet Protocol v6. LACNIC reported that its pool of available IPv4 addresses reached the" - and this is why it's not "none," but they are down to just - "4,194,302 mark, and that this has triggered stricter Internet resource assignment policies for the continent. In practice, this means that IPv4 addresses are now exhausted for Latin American and Caribbean operators."

The CEO of LACNIC said: "'This is an historic event; the fact that it was anticipated and announced doesn't make it any less significant. From now on, LACNIC and its National Registries will only be able to assign very small numbers of IPv4 addresses, and these will not be enough to satisfy our region's needs.'

"Since it began operating in 2002, the organization has assigned more than" - okay, so now 2002, now we're in 2014, so 12 years. So in the last 12 years the organization has assigned more than 182 million IPv4 addresses throughout Latin America and the Caribbean. So they've assigned in 12 years 182 million. They now have remaining four. And obviously the demand and the rate of assignment has skyrocketed with Internet use in the last 12 years, so it's not, I mean, even if it were linear, this would be a problem. But we know that it's not.

So just finishing this, they said: "As agreed by the regional community, now that only 4 million available IPv4 addresses remain, LACNIC's pool of IPv4 addresses is considered officially exhausted, and the [quote] Gradual Exhaustion and New Entrants policies have come into effect, introducing new procedures and requirements for those requesting resources."

Now, the only thing I could find, because I wanted to do some more digging to figure out what that actually meant because it sounds kind of ominous, "the gradual exhaustion and new entrants policies," among other things, you now - individuals have to demonstrate a need for IPs. This is something, shoot, there's a term for it, an IP, like, request substantiation form or something like that. I have always had to fill one out when I've set up colocation relationships. When I originally had servers at Verio, I had to explain to them why I needed the block of IPs, and essentially how I was going to use them. And the same thing with Level 3. They have them, but even 10 years ago they were saying, you know, we'll give you some, but you just have to prove that you need them.

And so now what's happening is for the first time end-users are going to have to demonstrate a need. And no doubt they're going to have to demonstrate why they can't run behind NAT. And so what we'll start seeing is their answer for that is we need to run servers, because it's a server on an IP is still the way the Internet works. Whereas huge numbers of clients can all run with local addresses behind Network Address Translation. And so this is continuing this interesting story that we've been watching on this podcast for a surprisingly long time, you know, that the sky is falling, and we're running out of IPv4 space. Well, no.

And as we get closer to that, we start making, you know, suddenly IPv4 space becomes increasingly valuable. We have seen stories of people voluntarily relinquishing their huge /8 networks, like a whole first number, like a dot five. Dot five used to be unallocated, famously. That was what Hamachi was using. And that was actually not ever used by anyone. And so now it's in use. And it's still the case that there is a ton of non-routed IPv4 space. And I think, as IPv4 space gets tighter, several things will happen at once. There'll be more pressure to move to IPv6. And it may well be that at some point people will have no choice but for new allocations of IP space to be in v6 space, not v4.

But the truth is there's still a lot of unused v4. And so there will certainly be some pressure on people who are hoarding their current v4 space to prove their need, and

registries may start pulling it back. And the way that would be done is they will simply say, okay, look, you're sitting here squatting on v4, like on a /8 network, or a /24, depending upon the way you think about it. But normally the idea being that only the first digit is fixed, and you have all of the other three digits in your IP space. And so what someone could say would be you have six months to push all of your users to one end of that. And then we are going to stop routing three quarters or five eighths or who knows, we're going to stop routing what you're not using to you and put it back in the pool. And mark my words, a couple years from now we'll probably be having stories talking about basically people being forced to give up their space in order to bring that back into the pool and reissue it.

Leo: This is a great Wikipedia page on /8s that are still owned by various corporations like HP and DEC. Apple has its own 18-dot.

Steve: Doesn't HP have two, like 14 and 15, I think, is HP?

Leo: APNIC has 14. Well, Level 3 has at least two. But you can make a case for Level 3. It's the individual like Apple having - or AT&T or, well, I guess they're an ISP.

Steve: Well, yes. And, for example, when you say Level 3 has two, well, and I'm occupying...

Leo: You're one of them.

Steve: ...a little, yeah, I'm occupying - I have 16 IPs out of their...

Leo: Why does Merck, the big pharmacy company, have 54-dot? I mean, you know, that's crazy.

Steve: Yes. That is absolutely wrong. And believe me, there's conversations being held with them.

Leo: Oh, yeah.

Steve: Saying, okay, you know. And the idea is they have no, I mean, this is a public resource. And so the story I told of how this is going to happen is that they will be told that they're simply losing a chunk of their space. And so they'll be given time, but they're going to have to push all of their users - because no one needs, you know, that's 16 million IPs. They just don't need them. And so they just push their users to one side, one end of that block, and then the block will be chopped. They can still have their 56-dot, but not *.*.*, you know, it'll be 56.5 or 1. And, thank you, we're going to take back all the other ones because there's just - there's no need, just no need for them to have it.

I did note an interesting piece about Netflix talking about bandwidth stuff. We got one

interesting Q&A.

Leo: I actually do love this. What they're doing is so great.

Steve: Yup. So now what Netflix is beginning to do is they're showing poor connection notices to their users through their Netflix apps at their end. So Verizon's users first started seeing these, and Verizon was very unhappy. However, it became clear that it wasn't just Verizon, that is, Verizon wasn't being singled out by Netflix. It was bandwidth-based. And so in Netflix official policy they said their goal in doing this is to help their subscribers understand when their experience is degraded based on their network provider as opposed to their home WiFi, et cetera.

Leo: See, Verizon, Comcast, and the rest count on the fact that customers are going to blame Netflix. They're not going to blame their ISP. So this is Netflix's way of saying, eh, maybe not. Maybe it isn't our fault.

Steve: Right. Right.

Leo: I think this is great. And I don't, you know, Verizon's threatened to sue. I don't see how, on what grounds.

Steve: No. And actually Netflix has covered themselves. They said when Netflix feels that many clients are experiencing congestion on a certain segment of a certain ISP's network, they will display the message for clients who are experiencing degradation. So what Netflix is doing is they're looking at the traffic flow. They've got a client that is receiving at their end. And so it's able to have, in the same way that any interactive application does, it's able to say, hey, the stream is jerky. The stream is having gaps. And so they're able, the client is able to close the loop back to Netflix. And then Netflix is able to use that as instrumentation, endpoint user point instrumentation, in order to notice that a certain set of clients all in a certain region are having this problem and then say to the client, tell the user, we've got a problem with your ISP. So just FYI.

Leo: Verizon got them to stop by saying you can't prove that it's us. There's other things it could be. And they have a point there; right?

Steve: Yup, yup.

Leo: I mean, there's lots of things that could cause that.

Steve: Well, okay. So you referred to, I think it was two weeks ago, probably during our Net Neutrality, the Level 3 blogs?

Leo: What a great post, yeah.

Steve: Yes. And so I had never - I had them in my notes. I didn't transfer them. But they're in today's show notes. And there's two really good blogs by Level 3 that I wanted to share with our audience. And in fact I'll, of course, I'm in a Level 3 datacenter, and Level 3 is a Tier 1 bandwidth provider. One of the blogs, there are some charts in them which are really telling because they demonstrate the saturation of a so-called "peering point."

And Leo's got them on the screen right now. The lower one shows no saturation. That is, there's a daily cycle, so you can see this daily cycle by date, where in the evenings, when there's more call for bandwidth, the bandwidth utilization goes up. And it just, just touches the total carrying capacity of that point. And this is a 100Gb fiber or switch or peering point. And then it goes back down again, and it goes back up.

The point is that, because even at peak that particular point - in the diagram this is a Washington, D.C. locator point. Thanks to a little bit of buffering, it never actually drops packets. But the first chart is a completely different example. And this one is an interconnect in Dallas showing the week of April 13th where unfortunately, and this is also, this is a 100Gb interconnect, saturated. And it is saturated for, like, three quarters of the day. So only at the minimum point in the day is there no packet loss.

And remember, I mean, we've talked about this extensively in prior podcasts, the way routers have buffers. And so packets are inherently sort of coming in sporadically from all kinds of different directions. And so you need to have some buffering because the output of the router is going to be a connection with a fixed rate. So what you want is you want the average bitrate coming in. It must be, the average bitrate coming in must be less than the fixed bitrate going out. Clearly.

And so the buffer takes, it smoothes out short-term variations so that it gets maximum value from the output by keeping it going. So the buffer is good for utilizing your outbound bandwidth. And it's necessary because you might have bursts where more is coming in briefly than is able to go out, so the buffer holds it and is bleeding it out at the constant rate of its output connection. So this is one of the fascinating aspects of this packet-based Internet where, due to this web of connectivity and users clicking on links and browsers downloading resources, I mean, just everything is just sort of kind of happening at random times. The inbound buffers collect that and then keep the links busy.

But, and here's the problem. I mean, this is the entire problem. At any given point the router has fixed bandwidth output. And if more is coming in than is able to go out, it has to drop some. The buffer fills up, and more comes in. And a huge amount of science, amazing science has gone into the optimal buffering policies. And in fact this notion of quality of service, packets can be tagged. They can carry a QoS tag which essentially gives them priority. It says I'm VoIP. I am extremely time sensitive. Move me to the front of the buffer. That doesn't hurt anything because it means that that packet won't get delayed as it waits in line in the queue for its turn to be transmitted. It would have occupied buffer space anyway, so the QoS just moves it to the front in order to minimize delay. So this particular - so the idea is that this flow, as it's called, tagged with this quality of service, is saying I am delay-sensitive. Please don't make me wait.

Other types of quality of service might be, I'm not important. If you have to drop something, choose me. And other packets can say, I am really important. Please, if other packets don't say anything, drop them, don't drop me. And so obviously this is all subject to abuse. And a lot of this is essentially everybody behaving themselves and playing by the rules. And at the top level, generally that happens. And what we're seeing is the breakdown of some assumptions, which we'll get to in our Q&A here, that is, the

assumptions which are breaking down as the 'Net is getting stressed and what that means.

Really quickly, in miscellaneous stuff, I wanted to correct something. I've been talking about Chrome being so bloated. And it was during my playing with Chrome a lot during the certificate revocation work that I fired up Chrome on a different machine, and it didn't occupy nearly as much memory because I didn't have any add-ons installed. And I thought, what? And so I removed a bunch, which I had just sort of accumulated, and it went right back down. So, and I didn't have, like, the kitchen sink in there. I just had a couple things. But, boy, they were big. And so...

Leo: Which ones in particular?

Steve: I don't remember. I think I still have - yeah, I don't remember. But I've been saying that, and I wanted to just say, hey, turns out it was the add-ons. Chrome with no add-ons, it does launch a bunch of processes, but that's actually done to get process isolation for security purposes. So I can forgive them that. And I'm beginning to see people talk about add-on bloat. That is, there was an interesting - called HTTP Switchboard that allows you huge amount of control, much more so, for example, than NoScript allows over on Firefox. And the developer was specifically saying it's much less large than most Chrome add-ons. So people are beginning to be aware that that's where some of this bloat is coming from. So for what it's worth, if anybody else - I have had other people say, oh, yeah, god, my Chrome is just huge. It's like, well, it's probably your add-ons. So I wanted to mention that. Also...

Leo: Do people really, I mean, that is more an aesthetic thing because, really, don't we have enough memory nowadays? They got eight gigs of RAM on a lot of these machines.

Steve: Well, for example, Jenny is unable to run Chrome and a tool that she uses for editing screenplays at the same time. She always had been, but she stopped being able to use both. And because I know she's a heavy Gmail user and she uses Chrome, I said try closing - because she was getting errors. And she of course said, "Hey, what's going on here?" And I said, "Try closing Chrome." And, yep, no problem at all. So it is the case that we can still be running out of memory. And it's just wrong, Leo, you know? Says the assembly language programmer.

Leo: That's really it. It's just unaesthetic that it's taking so much RAM. Because modern operating - I have not seen an operating system run out of memory in a long time. Have you? I know Jenny's did, but I wonder about that screenwriting app. Frankly. OSes don't run out of memory anymore.

Steve: And it's no doubt that it's a small vertical app that uses a lot of memory, yup. But still...

Leo: I mean, seriously, operating systems handle this well. Usually. The worst that will happen is you go to swap, and it slows down.

Steve: Correct. Correct.

Leo: I haven't seen an OS or an app saying I don't have enough memory to run in literally 15 years.

Steve: Well, it happens to me all the time.

Leo: Really?

Steve: That I'm running out of virtual memory, yeah.

Leo: Virtual memory.

Steve: Well, you know, like RAM. But I'm still on a 32-bit OS, so I've got a 4GB limit. But again, 4GB...

Leo: Well, what happens? When you're running an app that runs out of memory, it says I don't have enough memory to run?

Steve: Yeah, I get a balloon pops up on the desktop and says OS is running low on memory.

Leo: You need to get a modern operating system, my friend. That's really sad. It doesn't go to the VM, huh? It doesn't - rather, the swap? I can see it would slow it down.

Steve: I'm just telling you what happens.

Leo: All right, all right. We'll catch you on Windows Vista. That'll fix this.

Steve: No, no, not Vista. I'll skip that. So we saw "Edge of Tomorrow" and absolutely loved it. Just, for what it's worth, for sci-fi people, Tom Cruise did it again. I thought it was absolutely a fabulous ride, wonderful movie. My favorite movie of last summer was "Oblivion," and this was another great piece of science fiction.

So unfortunately I can't say the same about "Halt and Catch Fire." I am tolerating it and watching it; but, ooh, boy. It's a little annoying. We've had two episodes now; and, yikes.

And lastly, a SpinRite update. I got a nice note, and actually he saw the show notes and tweeted me, thanking me for sharing his testimonial in today's podcast, and I thanked him for the testimonial. This is Erik Ellsinger in Sweden, and his subject was "SpinRite saved me from going crazy." He said: "For a while now my computer has freezed up from

time to time. It completely stopped responding to input, and the hard drive light was on constantly. It usually lasted for 10 to 30 seconds, and then all was fine again. But lately it started to do that at least once every five minutes, which was driving me crazy because I got interrupted all the time with whatever I was doing.

"I looked in the Win 8 task manager" - so not an old system, or I guess he maybe could have upgraded an old system. But he says: "...the Win 8 task manager and saw that during the times the computer freezes, the operating system hard drive is at 100% use. But the response time is 0ms, and it's not reading or writing. But the hard drive has to be doing something, since the hard drive light on the computer is fully on. Has the hard drive encountered a sector it can't read, and it's just retrying it until it gives up?

"I've had SpinRite for a while now, and I thought I would give it a try. But I needed my computer for work, so I waited to run it until school was over. After SpinRite was finished, I booted Windows. And now, after using my computer for a day, I can say that the freezing problem is completely gone and cured, and SpinRite saved me from going crazy. Erik." And I'll just say that we hear this all the time. People complain that their computers are booting more slowly. They're just not operating as quickly as they used to. Drives, as we have discussed, can have serious problems which really go unnoticed because they try to manage it, try to manage what's happening inside them as best they can. And ultimately what happens is you'll turn your machine on, and it'll say no operating system, operating system not found, missing operating system. Or it'll go into a boot loop or something. Running SpinRite on systems which aren't obviously screaming that their hard drive has died often fixes problems that are less obviously about the hard drive, but actually are. So again, valuable preventative maintenance. And thanks for sharing that, Erik.

Leo: All right. Are you ready, Steverino, for some questions?

Steve: You betcha.

Leo: Well, I've got them. You sent me a bunch of questions, and we're ready to answer them here, starting with Question #1. Let me switch out of the show notes.

Steve: That's a good place to start.

Leo: Well, I could start with Question 2, but then we'd have to backtrack, and I don't like to.

Steve: Yeah.

Leo: Samuel Johnson, famous diarist, follower of Boswell, writes...

Steve: I don't think that's our listener. But...

Leo: You don't think it's him? Could be somebody else, maybe? He says he's in New York City. Oh, yeah, you're right. Samuel Johnson's from London. He says he's confused by your bandwidth billing: You recently mentioned something about "95/5" and "percentiles" with the bandwidth you purchase from Level 5. I'm sorry, Level 3. But what about megabits? What about gigabits? Don't you just have a connection? Huh?

Steve: Okay. So I did sort of breeze through that, and this is interesting.

Leo: This is because you don't buy bandwidth the same way normal people do.

Steve: Well, actually, I buy bandwidth the way datacenters sell it. And datacenters sell it the way it should be sold.

Leo: Ah.

Steve: Normal people buy bandwidth based on dreams.

Leo: Fantasies. I give - so, for instance, I just got Comcast business class. And they say up to 100Mb down, up to 10 or 20Mb up. And it's $200 a month. It's expensive. That's how normal people buy it.

Steve: Right.

Leo: You don't get that.

Steve: Now, but remember - no. And remember that what we found was that when users, end-users who had agreements like that actually tried to use it all...

Leo: They get in trouble.

Steve: Yeah, the ISP said, whoa, wait a minute, uh, that's not what we meant. So when you're in a real datacenter, I mean, like, any datacenter, this 95/5 is industry standard, is the way bandwidth is billed in datacenters. And the way it works is interesting because there are two things that I pay for, or two rates. There's something called the "committed data rate," or CDR. And that's sort of like, in general, it's what my bill will be. That is, I will never be billed less than the so-called committed data rate. And I am free to use that data rate 24/7. Literally. And, for example, for GRC that's 10Mb. That's - I used to be 15, but I just wasn't using nearly 15, so we brought it down to 10 to save me some money.

But there's also something called "burstable data rate." And I have a 100Mb connection to the Level 3 switch. So although my committed data rate is 10Mb, I am burstable to

100. Now, that turns out to come in handy because, say that somebody somewhere wants to quickly get a Security Now! audio file from us, or some other large file, like one of the videos we have on the site. Well, if they have a high download rate, they can ask for the file, and they'll spike my bandwidth up above 10Mb, but get the file quickly, and then they're done. So the tradeoff would be that they were clamped at 10Mb, and they still have to get the same total number of bits for the file. It would just take them longer.

And so since I'm in a big datacenter, and all of the connections running around are a gigabit or a hundred megabits, it's like, well, you know, someone's going to get the file. They might as well have it quickly. Then they'll go away, rather than make them hang around and so forth. But then the question is, how does Level 3 support the infrastructure required for this burstiness? That is, if the entire infrastructure was just 10Mb, that would be much less expensive for them. Yet you wouldn't have the convenience of being burstable, of being able to sort of blast out a lot of data in a short time and then be done.

Leo: This underscores why this is something that normal people don't do. You're serving. We're not serving. We're down - we're the other end.

Steve: Correct.

Leo: This is you pay for, and I pay for it this way on our servers, as well, you pay for the price of having a server, of delivering, of offering data as opposed to just consuming it.

Steve: Right. Right. So what if I went over this 10Mb so-called committed data rate for a certain amount of time? That is, I'm allowed to have bursts that go high. But how many? For how long? And how do you bill that? So the answer is as follows: They take one month, which is the billing interval, and they divide that into five-minute slices, and they count the number of bytes and multiply by eight to get the number of bits. So they count the amount of data interchanged in each five-minute slice. And so essentially that's the average data transfer in a five-minute window for every five-minute window in the month. And that's going to give you a set of numbers of the bandwidth used in each of those windows. Then they sort them from maximum to minimum. And they take the highest 5% and throw it away.

So the highest 5% of the sorted array of five-minute windows is just forgiven, ignored. But the next highest one, that is, the 95th percentile, is you are billed as if that was your data for the entire month. So one of the consequences of this is that, in the old days, when people were unaware of this, sometimes they'd get hit with an unbelievable bandwidth charge. You may have heard of this sort of anecdotally, of this happening to people where, like, a company got hit with an insane bandwidth charge because the 5% of the month turns out to be, I don't remember, I think it's 18 hours. And so if, in an entire 30-day month, you had 18 hours scattered through the month at really high bandwidth, essentially you would get charged as if that was the bandwidth you'd used the entire month. And if your burstable data rate fee is higher than your committed rate, and it typically is, and if it was, like, way higher than, for example, if it was three times what your committed rate was, you could easily be hit with a bill that was four or 500%, four or five times what you're normally charged.

But this is the way all datacenters bill. It's called 95/5 billing. Wikipedia's got an article.

And, I mean, it's like that's the way it's done. And this has been worked out years ago. And it turns out that it's the right formula for creating a facility which needs to be able, as Leo said, to serve data to customers, to be able to handle spikes and peaks which are natural. There's a daily cycle. There's event things that happen. But this is the way the big boys pay for their data. Not, as Leo said, the way end-users do.

Leo: Although it raises an interesting question because this goes back two weeks ago to our conversation with Brett Glass, who said that his customers didn't want to pay what the bandwidth actually cost. I wonder if ISPs at some point might start charging this way. I guess not. It doesn't make sense to charge this way to end-users.

Steve: Well...

Leo: But then they could get paid for - in other words, it's tying how much you use to how much you pay as opposed to the flat-rate pricing we currently have.

Steve: And my example when we were talking was electricity. That's the way we pay for electricity.

Leo: Right, pay by the minute.

Steve: People, like, yeah. Well, people know that if they run their AC during the summer...

Leo: You're going to pay more.

Steve: ...that their bill's going to be higher. And it does, now, the nice thing about that is it does shape behavior. Suddenly, air conditioning is not free.

Leo: Right, so you use it less, exactly.

Steve: It's, you know, you open your windows more. You try to do without it. You minimize your use. And so anyway, this sort of leads us right into our next question.

Leo: Well, let's go to it, then. Question 2 comes to us from Bill Sherwood of Goldendale, Washington, which I always thought Chicago was the Windy City. Turns out Goldendale is the Windy City. Who knew?

Steve: So Bill says.

Leo: I don't understand what Brett and Leo were fussing about concerning Net Neutrality. A content provider pays to get on the 'Net, and a user pays to access it. Each pays according to usage. Simple. Done.

Steve: Okay. So this comes back...

Leo: That's kind of what I was saying. But okay.

Steve: This comes back to what I was saying about assumptions.

Leo: Right.

Steve: And what's happened is, as the 'Net has evolved, some original assumptions have begun to break down. So the ISP made an assumption which has turned out not to be true. And that is, users surf the web, they send and receive email, and from time to time they download files. Probably for most users that's the case. But then, of course, famously, torrenting happened, and people were connecting in a way that was new, that wasn't part of the original what-a-user-does model, that caused the breakdown of those assumptions.

And then something like Netflix happened, where television that was traditionally delivered over the air or by cable, which is a shared medium - remember that cable works because it's got Channel 5 on it. And anyone who wants to can tune to Channel 5 and suck it off of the common cable. But switching to an Internet model really changes things, as we were discussing two weeks ago. So now users are not surfing the web, click and look at a page, click and look at a page, or sending and receiving email. They're streaming video in real-time, which is a huge change in behavior.

Now, on the transit provider side, this is something that bears on this because, for example, say that on one side we've got Netflix, who's - I think they use Cogent as their datacenter provider. And say on the other side we have AT&T. And in the middle is Level 3. Remember that the Internet is a network of interconnected private networks. So the original assumption of the providers, the bandwidth providers, was we would do, everyone would do free peering. That is, they would peer their networks with each other. And the assumption, again, because this is all about assumptions, the assumption was that the value each provider receives from other providers carrying their customers' bandwidth would roughly equal the cost they incurred from carrying that other provider's.

Now, imagine, though, that Level 3 is sitting here with Cogent on one side and AT&T on the other. And there's just bandwidth going through, passing through Level 3's network. And it's like, well, okay, you know, fine. Except that at nighttime this bandwidth rises to a level where Level 3's routers, the actual routers on the edges are no longer able to keep up, such that the buffers inbound are spilling over and dropping packets. So first of all, I mean, I don't want to - there's really no fault here. So this is not about fault. It's that a leg of the Internet is being saturated because, as a coincidence of geography and this new model of usage, which is users are watching television on the Internet, the 'Net is being strained in a way that it never was before.

So here's Level 3 that, like, wants to let transit happen, yet its buffers are overflowing. Now, not only are packets being dropped for the TV watchers, but unfortunately, since packet-dropping is generally nondiscriminatory, Level 3's customers are complaining. That is, their traffic is being interfered with for no fault of theirs, and no fault of Level 3's, because it's all trying to get through this pitch point which is being saturated by traffic that isn't originating in Level 3 or coming to Level 3, just trying to pass through. But it's got to use that connection.

So one of the things that can be done, and I don't mean to single out Level 3, and they're not doing this as far as I know, but there have been ISPs, carriers that have said we need our traffic, our customers' traffic, to have priority. At that overflowing buffer, our customers' traffic gets through. And so that's where suddenly people start getting upset because they're saying, okay, wait a minute, you're not treating all traffic the same now. You're discarding traffic preferentially, and that's not fair.

And so here's Level 3, that is sort of as a courtesy letting this torrent go through. And that's fine, as long as there's room for it. But if there isn't room for it, then something has to give. And the people on either end say, well, install more routers. Make room. And Level 3 says, why? This is not our traffic. This is just stuff going past. We're getting no value from it. So the point is, without any side-taking, without any what-should-we-do, although I think everybody agrees we need to keep Washington out of this because lord knows legislators are not going to improve this, this is a breakdown of assumptions. It used to be that the way the Internet was being used, a set of original assumptions held. And they really don't any longer. They've been, at the best, these original assumptions are being strained. And we'll see how it shakes out. It's not clear how it's going to.

Leo: Question 3 comes from Charles Woods in Katy, Texas. He brings us a blast from the past: Steve, could you provide an updated answer to my question way back in Episode 24, Q&A #3, knowing what we now know from the Snowden information? At the time he asked: With U.S. government NSA eavesdropping and spying so much in the news - oh, those were good times way back then - do you really think that SSL, SSH and other things we think of as safe are truly safe from the folks who, you know, have this high-end stuff and big computers, like the NSA? Can't they just crack through strong encryption?

Steve: Well, we know a lot more now than we did then. And I would say we still have confidence in things like SSL, SSH and other things. That is, we have confidence in the technology.

Leo: In fact, you could make the case that everything that's been released so far enhances our confidence because there's no mention at all about being able to break through those things.

Steve: Yes. Everything we're seeing supports this concept of the weakest link. And we know that systems are porous, not because crypto breaks. And in fact, I think it was Shamir, he was famously quoted as saying, "Crypto almost never breaks. It's everything else." And so it's the glue. It's the connections. It's the fibers. It's the human factor. The actual math is still secure, always has been. And even where we now think the NSA may have been trying to soften some algorithms, the math itself was good. It's just where it came from that was suspect, or who chose those numbers that you're using. Where did those come from?

So we've lost our innocence in the last year. And I think we, as an industry, we have far more appreciation for how large a budget the NSA has for how much they want to monitor the population. And, wow, look at the change in SSL over this period of time. Like now all these services that Firesheep was once able to penetrate, only a few years ago, are now HTTPS always, and users are much safer.

Leo: So trust the math.

Steve: Yeah.

Leo: Question 5 comes from Robert Osorio, Lady Lake, Florida: Steve, I guess I teach my clients well, or maybe I'm just lucky. None of them had been hit by CryptoLocker, well, until the other day. I teach my clients not to rely on their antivirus and to use common sense when dealing with email links and email attachments. Unfortunately, one of them got bit the other day, but it could have been worse. I don't like fancy antivirus solutions, and I steer WAY clear of Norton and McAfee's bloatware. Usually I recommend free solutions for my residential clients. I guess he teaches, I would guess from his location, seniors how to use computers. Security Essentials - that's the Microsoft solution - AVG, AVAST, and Kaspersky or ESET NOD32 for businesses, although he, like us, prefers the basic AV versions, not the bloated suites.

This particular client just had Microsoft Security Essentials installed on a Windows 7 PC. MSE actually detected CryptoLocker and repeatedly removed it and killed the process repeatedly for 10 days, according to the logs. Wow. Like most trojans, CryptoLocker has a second hidden encrypted component to restore itself after an antivirus app deletes the primary file and process. So while the antivirus could detect it, it couldn't permanently remove it. The client ignored the repeated detection and removal alerts. I had to lecture her about that afterwards.

However, Microsoft Security Essentials did slow down the progress of the trojan. CryptoLocker hadn't gotten very far in the encryption process. It had only encrypted 500 files in the Documents folder in 10 days, like half a file at a time or something. There's a utility on BleepingComputer that can list the files that have been encrypted. Apparently MSE kept interrupting the trojan and forcing it to restart. A pity my client decided to ignore the continual removal alerts MSE was giving (sigh), or I could have stopped it sooner. She only called me when she started having problems opening some Word documents. Turns out when Word opens or tries to open an encrypted .doc file, it opens it as a text file that's full of gibberish.

Fortunately, she had an offsite backup, Carbonite, yay. So after I had assured myself the trojan was removed, it was just a matter of restoring all the encrypted files, working from the list I had generated. The virus works through the Documents folder in alphabetical order - first the root, then the folders. It had only gotten as far as the folders starting with "C." I thought it was interesting, even though MSE hadn't completely eliminated the trojan, it did put up a serious roadblock to its progress that slowed it down dramatically. That's good news, yeah.

Steve: Yeah. And interesting feedback. I'm glad that MSE is able to catch it. I thought it was interesting that this person had so many documents that 500 was a small fraction of what was in her My Documents folder. And the alphabetic list, alphabetical order I

thought was interesting. And that's why I wanted to share this specific information with our listeners because I know from our feedback that they have encountered, if not CryptoLocker on their own site or on their own machine, on others. And so I imagine some of that could help.

Leo: I have to say parenthetically, you know, I've used Gmail for ages. I think we might even have talked about this, that the Gmail IMAP implementation is nonstandard and not very compatible with a lot of the stuff I use. Never worked very well with Apple Mail. They really - Gmail wants you to use it in the web, period, in the web interface. So I went back to my old IMAP provider, who I dearly love, and I've had an account with them for 10 years now, and I'm using them for primary email. But what it's done is it's surfaced how much crap I get that Gmail had been filtering because Gmail's antispam is really superb.

Steve: Ah, yup.

Leo: So even, you know, and I have SpamSieve and SpamAssassin running, SpamAssassin running on the server, SpamSieve locally. I have very good solutions. But this stuff still gets through. And a lot of it, a huge amount of it, is phishing scams. I get four or five credit card offers every hour, and they all are HTML emails with a big button that says, "Thank you for being such a loyal American Express customer. We'd like to give you an American Express card," and things like that. And they are so convincing that I am amazed that any normal person doesn't just fall for them every day. I just despair because, if you don't have really good antispam, and pretty much that's Gmail, then you're seeing this every day. And at some point you're going to click on one.

Steve: And remember, the major breaches that we've heard about, for example, famously, the huge RSA breach, that was traced back to one secretary who opened one file.

Leo: Email.

Steve: And that's all it took, yeah.

Leo: Yup. My mom sent me a note. She said, oh, send that fax again. I couldn't open it. And I went, oh, that is not good. Fortunately, she's on a Mac as a limited user; right? So I know she's going to be safe. So I looked, and the email came from an email address called fax@leoville.com. And it was a link to a Dropbox folder. Now, fortunately, Dropbox I'm sure keeps on this, and they see this happen all the time, and they had already killed the folder or killed the file. So she couldn't get it. But it looked like it came from me. It looked like it was a fax. It was a file in a folder in Dropbox, and she tried to open it. Had she not been running Macintosh as a limited user, and had Dropbox not deleted the file - so this stuff must happen all the time.

All right. Moving along, Mr. Steve, to Question 6 from Matt Reyes. He wants to know about our podcasting microphones: Evening, Steve. I'm curious about the mic you

use for Security Now!. Can you hear the difference? I don't know.

Steve: Well, and I just thought I'd give you an opportunity, Leo. Looks like you're still using the Heil. I see it in front of you.

Leo: Oh, yeah. You were with me when I won that Heil.

Steve: Yup.

Leo: That was at the first Podcast Expo in Ontario, California. You were there. You won Best Security Podcast. TWiT won Podcast of the Year. And the prize for the Podcast of the Year was this microphone, a Heil.

Steve: In a beautiful box.

Leo: Oh, yeah. Wood box.

Steve: Beautiful wood box.

Leo: Heil Sound PR 40 microphone. I'd never heard of it, frankly. Never heard of Bob Heil. And as a radio guy I'd used the RE20 and the RE27, the Electro-Voice mics. I'd used Audio-Technica mics, very high-end Sennheiser mics, Shure mics, lot of radio stations use those. So I was familiar with kind of the big brand-name radio mics. And I sat down, I think we used it on a podcast that day in Ontario, and I was playing with it.

Steve: I was like...

Leo: Remember that?

Steve: I remember the amazing bass response.

Leo: Wow, this sounds good.

Steve: Really amazing. Yup.

Leo: And I fell for it. I think they're really good mics. At this point we, of course, we know Bob Heil. I got to know Bob Heil as a result.

Steve: And then remember I was traveling to Canada with them. I had a pair with me, and goosenecks, and when...

Leo: You came up with Heils?

Steve: Yes, I had brought Heils up with me, and we set up our little mini studio to do the podcast when we were still doing the Call for Help show in Toronto.

Leo: Wow.

Steve: Yeah.

Leo: Yeah. So we're fans. We've been fans for a long time. Must be almost 10 years now. And they are actually pretty affordable. A lot of podcasters use them. We send them to all our regulars. It's HeilSound.com. He's an old-time rock-and-roll audio guy. And he also was a ham for many years. He said, "These ham mics suck." So he started making ham amateur radio mics, and then he started making mics for broadcasters, a lot of rock-and-roll guys. Charlie Daniels Band, all the mics are Heil mics. Lot of big bands do it. He's really good. So, yeah, Heil Sound. And I think it's around 300 bucks. They're not very expensive.

Steve: What would you say the mic is known for? That is, if it had, like, a personality or a reputation?

Leo: Well, the thing that you kind of - these are designed much like the Electro-Voice RE20s and 27s. They're dynamic. That means they're unpowered big-coil mics.

Steve: Right, so it's actually a generator.

Leo: Yeah.

Steve: It's a coil moving over a magnet.

Leo: So the sound pressure creates the electrical impulses in that coil as opposed to a condenser mic. They tend to be very small. We still use condensers on lapel mics and so forth. But they're powered, and they tend to be oversensitive, frankly. So the reason I use these, A, because they sound good. With male voices particularly I think they sound very good. But also they have really great what we call off-axis projection. So if I turn the mic, instead of speaking into it, if I turn it sideways, I immediately go off mic. And so what that means is noise in the room is 20 or 30 dB down. So you don't hear it. So I don't need a soundproof room. You don't need it. Our hosts don't have to sit in a radio studio anymore. And I really - I'm a fan. And they're very affordable. I mean, that's the point, 327 bucks.

Steve: And you really do get - you get no room noise from these. I mean, you hear me, but none of the echo that you would normally get if you had a microphone in a room.

Leo: None of that room tone. Of course, if there's garbage men emptying your trash, we hear that. It's not impervious to sound.

Steve: Well, they're kind of over there, too.

Leo: They're that way. Yeah, maybe you should aim it that way next time they come and see how it works. Heil is also the host of our show Ham Nation, which is an excellent show and a great success, by the way, on our network. It's a show for hams.

Steve: Nice. And those are people who like to talk.

Leo: A little bit.

Steve: They like to be in front of the camera.?

Leo: They love to talk.

Steve: They're hams?

Leo: Here's the package we send people. John just handed me this printout from BSW, Broadcast Supply Worldwide. You get the Heil microphone; the shock mount, they have a very good spider mount, which I see you're using, as well; a boom, which you're also using, which keeps it off the table; and the pop filter and the cable. And all of that 369. That is...

Steve: Mine is actually dusty, Leo. Now that I'm looking at it, it's like, oh.

Leo: And that's the beauty of Heil. You never need to dust it. You never need to dust it. Now, I've used Neumanns. I've used some of the best microphones in the world. And I just, for day-to-day use, these things are rock-solid, and the whole network uses them, and many of our hosts, like Steve. We like them to sound as good as they can. So thank you, Bob Heil.

Steve: We need all the help we can get, and Heil provides it.

Leo: Yes. We're going to Heil. Thank you, Matt, for asking about that. Jeff Leavey in Poughkeepsie, New York wonders about drive spin-down: I use a laptop docking

station to back up my internal drives to an external hard disk. My concern is I may be putting an external drive at risk if I power down the docking station and then withdraw the drive too quickly. How long should I wait, if at all, to handle the drive after powering down? You know, that's actually a great question. When I unplug a USB drive and pick it up, there's still centrifugal force, or centripetal force.

Steve: Yeah.

Leo: And it's like a gyroscope. It's hard to handle. He says, like a gyro, torque is created if I move the drive with the platters still spinning. Is that bad for the drive?

Steve: So, okay. Drives never leave their heads on the normal data surface. In the old days, drives would retract the heads. And there was actually - there were ramps out on the edge so that the actuator would pull the heads, and they would hit these wedge-shaped ramps that would lift the heads off the platter as they came away. And then the reverse would happen when the drive was powered up. First the platters would spin, and then the heads would be rather gently lowered down onto the platter where the air bearing would keep the heads from ever touching the platter.

What's now being done, because that's an awful lot of mechanical and extra cost, is that the heads are just moved into the middle. They're moved into the hub. And the idea is that you wouldn't want the heads on the outer edge because if the drive does receive any shock, the platter, which is anchored in the center, would tend to ring like a bell, and it would vibrate. And of course the center is anchored, so the outer edge would have the largest displacement as this thing is ringing. And so the heads would be bouncing up and down. Consequently, the heads are moved into the very center, where there's the least movement.

So, now, relative to torque and that gyroscopic action, the good news is the minute power is removed, the heads are safe. They're either retracted and off the disk or put into the center, where they have the least opportunity for damage. The other reason they're put into the center, I should mention, is the torque moment which exists when the disk is spinning up, once the disk spins down, the drives do come into contact with the disk platter. There can be just sort of a type of welding which occurs, just due to the fact that you've got two super smooth surfaces in contact. They can just sort of do a spot weld.

So by having the heads in as close to the center, the torque of the motor has a much greater chance to break that weld than if the heads were way out on the perimeter of the surfaces, where they'd have a much stronger ability to keep the disk from spinning up. So as soon as the power's cut off, the heads are moved into the middle, and essentially it's safe to move the drive. Certainly doesn't hurt to give it 30 seconds to spin down. That's probably enough. But immediately after, as Leo has sensed, and Jeff, as you have noticed, when the disks are still spinning, it's a little gyro in your hand.

Leo: So cool.

Steve: And you are putting some torque on the bearings, too.

Leo: It's not good for it. But it's not as bad. It's good to know.

Steve: Right.

Leo: You're not crashing the heads.

Steve: No.

Leo: That's really good. Huh. I'm glad he asked that. I've been meaning to ask you that about eight years. Been meaning to ask you that. Our last question, I'm sad to say, comes from Charles - not, no, I'm sad it is our last question. I'm happy that it comes from Charles. Charles Miller in San Miguel de Allende, Guanajuato, Mexico, beautiful town, closes this week with something fun: Steve, the cute receptionist at the gym where I work out asked me what I listen to and could she put it on the sound system. I showed her the screen of my MP3 player. Later I noted her searching online for a band by the name of Harvesting Entropy - the title of our Security Now! episode a couple of weeks ago. Wow.

Steve: Yup. Got a kick out of that. So thanks for that.

Leo: That's very funny. It would be fun to play Security Now! through the gym and see how long it takes before people pass out.

Steve: See how long their membership lasts, yes.

Leo: Yeah, what's going on here?

Steve: Turn the station.

Leo: Although Harvesting Entropy would be a good band name. I think I like it.

Steve: Actually, it's a great band name, yeah.

Leo: Yeah, isn't it? Yeah. Steve is at GRC.com. We talk about it all the time. It's where you can get the greatest hard drive maintenance utility in the world, SpinRite. You can also get lots of free stuff, including Steve's password stuff, his security stuff, ShieldsUP!, which has been in use now for, what, 15 years? How long? For a long time. Ten years, anyway.

Steve: About 15 years, and I think we're at 59 million, if I remember.

Leo: That's how many people have used it. That's incredible. Great way to test your router before you put it online. I did the other day. I got a new router from Comcast, and they blocked, it's interesting, they were blocking 139, the NetBIOS ports that were how I first met you.

Steve: Yup.

Leo: So it all came around. You can also go there to get 16Kb audio versions of the show, if your bandwidth is limited, and full, beautifully written, human-written transcripts from Elaine Farris, who writes them for Steve. Those are all at GRC.com. That's also where you can go to ask questions for future Q&As: GRC.com/feedback. And you can find out more about, oh, about SQRL. You can go to the - I was at the forums the other day, looking at the forums. You saw the guy who posted what he thought was an Easter egg in the CryptoLocker announcement. Did you see that? A Latin phrase?

Steve: Oh, yeah, yeah. I don't know what the...

Leo: Talk about conspiracy theories.

Steve: I know.

Leo: If you take the first, I don't know, first letters of the TrueCrypt, not CryptoLocker, TrueCrypt announcement...

Steve: The TrueCrypt warning. And then you translate it, or no, I think you...

Leo: Latin.

Steve: You take it, you translate it to Latin, and then you take the first letters of the Latin version of it?

Leo: Anyway, it works except that you can't take the whole word "security." You just take the "se." And it says something like, what, I don't know, don't use this thing or something. I can't remember what it was.

Steve: Oh, no, it somehow has the initials NSA in it, which is what freaked everybody out.

Leo: Right, oh, yeah, yeah, yeah. Like it was a warning.

Steve: Yeah.

Leo: Anyway, I don't think it was real credible, but that's the kind of things you learn on the security forums at GRC.com. If you want full audio, full bandwidth audio or HD video, SD video, we have that on TWiT, TWiT.tv/sn. And of course you can subscribe because it's a podcast, and you can get a copy of it anywhere you get your podcasts, or use our great apps on the Roku or the iPhone, the Android devices, Windows, everywhere you want to be. The apps are great. Thank you so much, Steve. We'll see - what are we going to do next week?

Steve: Don't know. I'll have something fun. But the universe always manages to provide something, you know? So apparently it senses our need and provides a security disaster for us to cover.

Leo: So the translation is - let's see if I can find this: "If you wish, use the NSA." Okay. I guess. It could be. Why would they do it in Latin? All right, Steve. We'll see you next week on Security Now!.

Steve: Thanks, my friend.