Listener Feedback #74

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-212.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-212-lq.mp3

INTRO: Netcasts you love, from people you trust. This is TWiT.

Leo Laporte: Bandwidth for Security Now! is provided by AOL Music and Spinner.com, where you can get free MP3s, exclusive interviews, and more.

This is Security Now! with Steve Gibson, Episode 212 for September 3, 2009: Your questions, Steve's answers #74. This show is brought to you by listeners like you and your contributions. We couldn't do it without you. Thanks so much.

It's time for Security Now!, the show that covers all things security and privacy oriented, online and off. As long as it has to do with your computer, I guess, not with peeping Toms looking through the window. We don't help you with that.

Steve Gibson: No.

Leo: Steve Gibson is here. He's the man at GRC.com.

Steve: That's outside of our scope.

Leo: We haven't yet, anyway.

Steve: Yeah.

Leo: Never say never. Steve is the majordomo at GRC, the Gibson Research Corporation. They do that great SpinRite program and all sorts of great free - they. You.

Steve: They. I was going to say "they," the great unwashed masses of GRC. Yes.

Leo: Well, what was - how big? At one point GRC had...

Steve: Too many people. We had 23. And I just was pulling what little hair I have out. It was - I was just a babysitter. And I thought, this is not, I mean, I remember years later running across an outline that I had prepared. It was in the outlining program called Grandview, which I used to love. And it was an outline about a meeting that we had about how to have meetings. And I thought, my god, even our meetings were having meetings. That's all I did was have meetings. I hated it. You know, I like to argue with the bits and write the code and come up with solutions and things. And I was paying everybody to do that for me. And I thought, well, wait a minute. I've given up the thing I like doing the most, so...

Leo: Well, I watch, you know, I watch and learn, try to learn from that lesson. You know, very interested. I've never been an entrepreneur or businessman before. And so I pay very close attention to what you say about that because we're growing. You know we have six employees now. And I don't want to ever get to the point where I'm going to a lot of meetings.

Steve: Well, what I realized at one point, for a while it was exciting to have employees. It's like, whoa, look at all these people who are part of my team. And I definitely miss the camaraderie, the...

Leo: Yeah, we've got that here.

Steve: There was so - there was a really neat synergy among people. You know, people would come up with wacky things, I mean, it really wasn't aimed at clear productivity. It was just fun. It was just social. It was nice that, you know, I chose smart people, and so it's fun to be around a bunch of smart people. That's it's own reward. But I also at one point realized, hey, it's not how much money moves through, it's how much money stays behind.

Leo: [Laughing] I haven't learned that yet.

Steve: Yeah, yeah.

Leo: I'm moving a lot of money through.

Steve: And so...

Leo: Not keeping any of it.

Steve: And so I realized, yeah, it's really - it's heady to look at all the payables and receivables and look at the big numbers. But if it just kind of moves on past and you wave at it as it's going by, it's like, well, okay. That's exciting. That's a big stream. But you'd like to fill up your own reservoir sometime, so...

Leo: Well, but there's - and when you're building a business, revenue gives you more things you can do. And that's kind of where we're - I feel like we're at the building stage still.

Steve: Certainly have clout that way, yes.

Leo: And so the revenue helps us, you know, improve the studio, add more people. And as long as adding more people adds more revenue - see, really kind of it's - I feel like it's leveraged what we can do, what I can do personally by having extra - these people are so great. They help me. And then also we can play laser tag now. We have a team.

Steve: Well, for me the best thing that ever happened was the Internet because...

Leo: Yeah. You can virtualize it.

Steve: You know, receivables went away. I used to have two people who spent all their time just trying to get us paid. Because we were using big national and in some cases international distributors. And their contract said "Net 90," but they actually paid about 180 if we were lucky. So we were waiting half a year to get money from them, and then they'd send back, quote, "damaged," unquote, damaged goods. I mean, Egghead had a return policy where they'd take back any software. So people would buy a copy of SpinRite off the shelf and take it home, use it, and then say, well, you know, I think I'm done with it now. And they'd take it back and say, "I decided I don't want this." And so, you know, and then it was an opened box, which Egghead had return privileges for. So it would go back up the stream to the distributor. And then they would, like, send back huge boxes of destroyed product that it looked like elephants had had a party on them. And so I thought, okay, there's just got to be a better way.

Well, the Internet happened. And so by being able to automate this whole process, I've been able to - now I've got two people, Sue and Greg, who handle the accounting, bookkeeping, operations stuff on Sue's side, and dealing with all of our customers, the tech support needs, on Greg's side. And I just kind of get to move the technology forward. I just - this is it. I mean, this is just paradise for me. And by virtualizing, by shutting down offices - they both work out of offices at their home - I'm able to save some money after all of this. So it works.

Leo: Yeah. You're smart. You're smart. I hope I learn this lesson. I love it because I talk to you and people like Jason Calacanis, who've been there ahead of me. And I listen with great, you know, great interest to your lessons.

Steve: Well, and I'll tell you, it does take some discipline. I'm approached from time to time, as you might imagine, by people who have really interesting ideas. They offer seductive alternative lives. And I think, you know, what I have now is perfect. Do not mess with perfect. So I generally politely decline. I say, ah, well, that really sounds good, but I'm not your guy.

Leo: Saying no is more important, is a better skill than saying yes. I agree with you on that one. I have to learn a little bit of that. So we've got a Q&A episode today.

Steve: We do.

Leo: Questions from our audience that we're going to answer.

Steve: Some amazing news. Everybody, I'm sure you've had - you've seen it in the news, the news of this new WPA/TKIP hack.

Leo: Yes. I really want to know about that because I've talked - like you, whenever that happens, I get all the emails saying, is this true? Is it safe? And in the past these cracks have been less dangerous than the headline might imply. I don't know about this one, though. We'll find out.

Steve: Oh, even more so. This is so bizarrely theoretical, it's like, okay, just wander off to whatever hotspot you want and don't worry.

Leo: Okay, good.

Steve: Yeah. And we have a bunch of fun and interesting errata. And then, of course, our Q&A.

Leo: So I guess we should start, Steve, with errata.

Steve: Well, news.

Leo: News? Okay.

Steve: We've got not too much. VMware Workstation was updated to 6.5.3.

Leo: They had a big conference in San Francisco, so they probably announced it then, yeah?

Steve: Yeah, well, this was mostly just - it was, well, the main motivation was probably they had - that they had the libpng problem in their code. So they updated, you know, we talked about that a few weeks ago. There was some overflow problems, not surprisingly, in image processing, PNG images in libpng. They updated internally to 1.2.35 and incorporated that into their 6.5.3. They also now offer full support for the Ubuntu 9.04 client.

Leo: Good.

Steve: So people who are using that flavor of Linux will be glad that they've got additional support. And there's also a ton of just other stuff. As I read through the list of all the other things, it's like, oh, well, that would be handy. Things like NAT, the NAT translation mode doesn't work on newer Windows clients. Oh, well, so that - it does now, but it didn't in 6.5.2, which is what I had previously. So it's like, okay, well, that's handy to have. And all kinds of other things that affect a smaller subset of their total users. So definitely worth updating to 6.5.3, which is the current, as of couple days ago, version of VMware Workstation. Also we've got a new Chrome. The Chrome browser from Google is now at 2.0.172.43.

Leo: [Laughing] 2.0.172.43.

Steve: Stardate, yes.

Leo: That's ridiculous.

Steve: [Laughing] What I found most interesting about this is that a severe flaw was found in V8, which is the open-source JavaScript engine that Google has been developing and is excited about that makes Chrome run so fast. But the problem was there was a way that you could - that JavaScript could be used to access unauthorized memory and also potentially execute code that is in the user's system. What I loved about it was that it was found and reported by Mozilla security. And I thought, well, that was nice of the Mozilla people to let Google know, I mean, a competing browser. Not that Google's - not that Mozilla's in much danger.

Leo: They use different engines, though; don't they?

Steve: Well, in fact I wondered if this meant that maybe Mozilla was taking a look at...

Leo: WebKit.

Steve: ...the V8 engine, thinking, oh, you know, maybe we ought to move that over into Firefox. Who knows.

Leo: Coulda had a V8.

Steve: [Laughing] And then there are some other reasons to update Chrome. There were some flaws found in the xml2 library. Oh, and the other cool thing is Chrome decided to formally no longer allow MD2 and MD4 hashes, which we know are compromised, in SSL certificates. So with this update you will not be able, you will simply not...

Leo: Wow, that's good.

Steve: ...be able to connect - yeah, it is - not be able to connect...

Leo: Do other browsers allow that?

Steve: They're still allowing it.

Leo: Wow.

Steve: Yeah, I mean, it would be nice even to have a notification. But then your typical user, what are they going to do when they get some notice that says, well, the hash this site is using is technically not secure and could be compromised. We don't know that it is. We don't think it is. But then we wouldn't be able to tell if it were. What's a typical user going to do? Just their head's going to explode. So Chrome has just decided to say, no, we're going to go the security route. And the cool thing is, well, if Chrome had a bigger market share, I think they're about 3 percent right now of market share on the Internet, so still not a huge factor. But it will certainly put some pressure on sites to update their certificates, if they're still using an MD4 hash.

Leo: They just started putting out a version for the Mac that's at least somewhat stable.

Steve: You know, and I fired up a VM in order to go to Chrome and look at it. It's so pretty. It's just...

Leo: [Laughing] They have - I don't know if they - they must have this on PC's themes. Because they have themes on the Mac. They have, like, grass and all those sorts of different...

Steve: I just, you know, I look at it, I go, "That's so pretty." But I'm not using it. I'm just...

Leo: I like Chrome. It's fast, it's pretty. But Fire- you know, it's interesting because some people are starting to think that Firefox is the new Internet Explorer, that with 3.5 it started to get a little bogged down; you know?

Steve: Oh, you mean - oh, I see. You mean slowing down and getting...

Leo: Slowing down, buggy, issues. And I think that people are looking more and more toward Chrome as an alternative.

Steve: It's funny, isn't it, that, I mean, it can happen to anybody.

Leo: Sure.

Steve: We used to have lightweight personal firewalls. Now look what they've become. I mean, you dread loading one of those things on your machine. That's why I can't wait for Microsoft's forthcoming AV solution. It's like, oh, good, thank you very much. And we have one really bizarre story. I just thought this one - I saw this pass by on my radar, it's like, whoa, isn't that odd. U.S. state offices, typically governors' offices, have been receiving HP laptops they'd never ordered.

Leo: I wonder from whom?

Steve: Isn't that interesting? In one case, West Virginia's governor, looks like Joe Manchin, received five unasked for, never ordered, HP laptops.

Leo: Hmm, completely loaded with the latest software.

Steve: Well, precisely. I mean, and what's interesting is this has been - there are other, 10 other instances - four were delivered, six were intercepted - of other government, U.S. government offices spread around the country, just receiving an HP laptop. Here you go.

Leo: Wow.

Steve: And, I mean, no one at this point - they have not been analyzed. There's no news of what they contain.

Leo: But I see the FBI is investigating.

Steve: Oh, yeah, yeah. The FBI is on it because they're thinking, wait a minute, what is the story with this? And the assumption is that security is now enough of a factor, people

are trained not to click on links in email, the idea being that the traditional ways of getting inside the perimeter are no longer succeeding enough, that now people are saying, well, let's just send them a laptop, all set to phone home, and see if somebody who's inside the perimeter, the security perimeter, fires it up; and, if so, that gives us a foothold inside the network. So isn't that...

Leo: [Laughing] That's a trap, somebody said in the chatroom.

Steve: It's just amazing.

Leo: Admiral Ackbar says, "It's a trap."

Steve: Don't turn it on.

Leo: No.

Steve: Now, I did, I have to say that when I heard they were HP laptops I was thinking, well, how could you tell if there was any malware installed?

Leo: There's so much other crap in there.

Steve: Oh, my goodness, yes. I mean, HP is the worst laptop I know of in terms of gunk all being preinstalled.

Leo: They, like a lot of kind of consumer-grade systems - I think of Gateway - have those background downloader/uploader fix programs running all the time. So they're always phoning home.

Steve: Oh, there's stuff going on. I mean, and all of the trial ware. Things are expiring and telling you, well, after your two months of using Betty's Flower Shop program, don't you want to purchase it? Uh, what? No. How do I get this off of here?

Leo: Now, theoretical question. Could they reimage the drives, just format it and would be okay? Or is it possible to hide something? I guess you could hide something in the keyboards, a keystroke logger, or something in BIOS; right?

Steve: Yeah, that's a very good point. You could go, if you're physically delivering a computer, that's a very good point, Leo, you could do all kinds of extra sneaky things to the hardware that goes beyond what your typical drive-by malware download could do. That seems like, you know, a little overly sophisticated. But, I mean, yeah, for example, you could have an extra radio.

Leo: A transmitter, yeah.

Steve: An extra WiFi system or something. Anything could be in there. So it's definitely creepy. And I would argue probably that people who don't order equipment and receive stuff should be skeptical. So I thought this was really interesting. Also it appears, and I haven't looked at it closely yet, but that there's a lot more information has surfaced since we mentioned it tentatively last week. I was a little skeptical when only the Register.co.uk had talked about this crack of GSM, the cell phone network technology. That's the topic for next week. We will do - I said next week's TID, which is Topic In Depth.

Leo: Oh, good. I like that. That's a new - that's our new - and now [vocal fanfare], the Topic In Depth. It's like CNN. I'll get what's-his-name to record something for us.

Steve: Sort of a nice announcement.

Leo: This is Topic In Depth.

Steve: Okay. But...

Leo: We're going to talk about ESM cracking?

Steve: GSM, next week, cracking GSM, which of course is a big deal because it's one of the encrypted technologies which we're using for cell phones, the other one being CDMA. And I have said a number of times that I was not satisfied with the encryption of cell phone technology. If you search back in our transcripts you'll find me having said that a number of times. And whoops, sure enough, here it is. Apparently if you - someone with a laptop and a special receiver is able to do what we imagine only the NSA previous to this could do. And they certainly have been able to for some time because the technology just wasn't that good. So we will talk about that next week.

The big, big, big issue that has happened between last week and this week, which so many people wrote in about, is what happened with WPA, that is, the traditional WiFi technology. Remember that WPA has two different types of encryption. It has the old-style TKIP encryption based on the RC4 cipher, and then also state-of-the-art encryption using the AES cipher. We did a whole podcast a while back about an exploit of the TKIP encryption which allowed a 15-minute-long decryption of a short packet if the access point supported quality of service, QoS. And the idea was that QoS, a quality of service access point, had multiple packet queues. And the reason this took 15 minutes was that, if you upset the access point with wrong guesses too often, that is, more than two within a certain time window, like a minute, then the access point would decide, oh, I'm being hacked, and it would shut down and rekey everybody. Which would cause you to lose the work you had done up to that point. So you had to - you had to make sure you didn't guess twice within a one-minute sliding window. And so that's why it took 15 minutes, because you needed at most 15 or 16 guesses which would probably be wrong in order to perform this particular attack.

So the big news of this new attack is you - and what the authors state is that it is no longer necessary to have quality of service support on the access point. And it's like, okay, well, so what else? Well, it turns out that what they've done is even sort of more theoretical because what it requires is a condition of the radio reception, which seems unlikely to occur without, like, some tremendous amount of work, which is that, strangely enough, the access point and the user who you're attacking cannot be within radio range of each other.

Leo: [Laughing] Wait a minute.

Steve: It's like, okay. So the attacker has to literally be a man in the middle, meaning their radio has to be the link between the access point and the user.

Leo: So you're posing as the user.

Steve: So, well, I mean, but physically, I mean, physically the radios of the two endpoints, the access point and the user, cannot be within range of each other, or this won't work. So the attacker is literally the intermediary radio link passing the traffic, acting as a relay, passing the traffic back and forth. But this doesn't work if the endpoints are within radio range. So it's like, okay, fine. So now what?

So basically this is essentially the same attack, but you're using the fact that the endpoints can't hear each other, radio-wise, to obtain the equivalent of what you got with the quality of service queues, that is to say there's - and if listeners have questions and really want to know about this, we did discuss the way RC4 and WEP works when we talked about how badly broken WEP had become. One of the problems with WEP, with W-E-P, Wired Equivalent Privacy, is that's the original WiFi standard that you just really don't want to use anymore because it's over as far as its security goes because it was so poorly designed in the beginning. There's something called an IV, the Initialization Vector, which typically is just a counter. And every packet which is enciphered, that is, encrypted, uses the next larger IV, initialization vector. And that's required because the way the encryption works, it's not secure unless you have this initialization vector which is used sort of to seed the encryption for every single packet.

Well, one of the things that WEP never did was to insist on initialization vectors incrementing. That is, you could take a packet and hack it, and meanwhile the access point is spitting out more packets with incrementing initialization vectors. So you could then take the packet you had intercepted and decrypt it and then retransmit it to the access point, which even though its initialization vector was now old and technically expired, the access point didn't enforce the currency of initialization vectors. Well, that was fixed in WPA. So one of the nice things about WPA is that no access point will accept an initialization vector that is older than any that it has already accepted. So the initialization vectors must be monotonically increasing in value over time.

Leo: Have to be in sequence.

Steve: They have to be in sequence, thank you. That says it more easily.

Leo: Or monotonically in time.

Steve: An out-of-sequence initialization vector is just ignored. It's thrown away. It's like, okay, well, we're not sure where this came from, but we're ignoring it. So the multiple queues in a quality of service supporting access point have independent initialization vectors. And so you're able to use one of the queues that's, like, way behind in order to play games with the encryption and the key, which is uniform among all of the queues. So that - so the fact that you had multiple queues meant that you had desynchronized initialization vectors because they were - these initialization vector counters were per queue.

Okay. So that was the wedge which the first group cleverly exploited. Now this newer approach says, okay, if we're literally the man in the middle from a radio link standpoint, then we're not passively - we're not passively listening. We're able to get a packet and mess with it before we send it on, meaning that the other end hasn't seen the larger initialization vector yet. So they're able to literally, to use the fact that they have sort of preemptive access to the traffic in order to perform the exploit.

The problem is that after all of this you still have nothing any more worrisome than what we had before, which is that you could only decrypt very short packets where you know most of the content, that is, you know what the plaintext will be. Well, that pretty much limits you to ARP packets. And in their paper they only talk about ARP exploits. And what worried people was that the title was "Useful Decryption in Less Than a Minute." Well, 37 percent, I think it's 36.9 percent of the packets could be decrypted in about a minute because of some optimizations they found. For example, they were able to nail down some other bytes in the ARP packet as a consequence of knowing the IP of the access point, since they have to be a member of the access point to be doing this radio traffic transfer.

So they did some clever things. But all you end up being able to do is spoof a single very short ARP packet in that length of time because you already know most of the data. You do not - and this is the critical part - you do not get the key. This doesn't crack the WEP key - or, sorry, the WPA key. It only allows you to determine on, for a given short packet, you are able to determine the cipher bitstream which allows you to basically change the ARP packet to something else. Oh, and during this time the other end is blacked out. That is, you can't forward that packet while you're doing it.

So they've got this other fancy business where normal-sized packets that are carrying non-ARP payload, they have to let - they have to, like, bridge those through to keep the user from knowing what's going on while they suspend the ARP packet for a minute. So, I mean, there's all these criteria, and the fact that they have to be out of radio range from each other. That is, the endpoint that you're hacking can't be within radio range of the access point. They have to rely on you to be the forwarder. So it's extremely sort of theoretical and okay, good to know, but it doesn't mean anybody has to run away screaming and worrying that WPA has been hacked.

The lesson from this is what we've seen time and again - and I'm sure listeners, long-term listeners of the podcast have seen this - and that is that, as something begins to get weakened, additional sort of chinks are found in its armor. So this is a reason to move from the TKIP cipher, which is sometimes referred to as WPA, over to the use of AES, sometimes referred to as WPA2, although that's really not an official designation, as we've mentioned when we've talked about this stuff before.

So nothing big to worry about, in my opinion. But further pressure on abandoning the RC4 cipher, because being an XOR-based cipher it does allow the so-called key stream to be reverse engineered. For short packets we're beginning to see this more and more. Who knows what's next? Better just to be away from it in case anybody comes up with something really bad.

Leo: And just to reiterate, it's the same thing we said last time: Use AES.

Steve: Yeah.

Leo: Which I have already done on all my WPA systems.

Steve: Yeah. The only reason I could imagine someone might not is if they had some piece of equipment which didn't support AES, and they were saying, well, okay, TKIP is all I can really use because I have to have this piece of equipment on my network. And I would say, okay, just recognize that it's not as secure. It still seems pretty good. But it's not clear what we're going to have come along tomorrow. Certainly as long as anything that happens is public, we'll let our listeners know. But, yes, use AES unless you really can't. And maybe even consider give any devices that can't use AES their own access point running TKIP so that at least they're not part of the same network.

Leo: That's what I did when I had to use WEP for some older hardware, just have a kind of standalone WEP router out there.

Steve: Yup. And again, if you can, isolate it from your internal network because you don't want somebody to crack that and then have access to your LAN.

Leo: Right. If I put that outside a WPA router, in other words, have the WPA/AES router be my main router, and then attached to it have a WEP router, does that isolate it? Or a TKIP router? Does that isolate it?

Steve: I would, if you had one additional router...

Leo: You need three to do this.

Steve: Yeah. Because of ARP spoofing problems. You don't want to let somebody be able to spoof ARP in order to intercept all of your inner network's traffic, which they would be able to do if they were able to see ARP traffic. So routers do not transit ARP traffic. They don't bridge ARP. They don't have to because they're maintaining separate networks on either side. So that's the one thing you would need to do is to have one more router between that and your WEP, your untrusted WiFi, because of the possibility of ARP spoofing.

Leo: And is it, on your base station, your access point, is it usually pretty clear which is AES and which is TKIP? Does it say AES?

Steve: Yeah [tentatively]. The problem is it's really fuzzy nomenclature. In fact, we talked about a router that would allow both, it was TKIP plus AES, and then there was also just an AES. Well, that's what you'd want to use unless for some reason you needed both. And the idea was, it's like, hey, a feature that the router had was that you could use either. Well, you probably don't want that. You want to stay away from TKIP. So unfortunately there just is - the nomenclature used is not standardized, and it is fuzzy. But most users, I think, if you look at the settings, and maybe you look at the corresponding help guide, you just want - you want the strongest you can get. But if it allows you to have both, you do not - you want to stay away from TKIP and not have that also.

Leo: Got it.

Steve: In errata, there's been some interesting news about ultracapacitors. We talked, an episode that many of our listeners have said they enjoyed as one of our rare sort of off-topic episodes, we talked about the whole idea, which fascinated me from a physics standpoint, about the company EEStor, I think they're down in Texas, who have a technology that they've been working on for a number of years that's got some impressive venture capital folks behind it, venture capital folks who don't tend to make mistakes. It sort of leaked out that they had let, that EEStor had let a contract out to a company called Polarity, and that Polarity had been given the job of using Polarity's high voltage-to-low voltage converter technology to be integrated into the so-called EESU, which is the name for EEStor's capacitor-based battery.

Checking out, I did a little browsing around, and Polarity, the company, looks very legitimate. Their little news page says that in '09 they were awarded follow-on production contract for the Navy's SPS-49 radar upgrade. Also in '09 they were awarded development contract for next-generation TWT Test Sets for Teledyne MEC, whatever that is. And then on their little news page it says 2009 awarded contract from EEStor to integrate Polarity's high-power HV to LV, which is high voltage to low voltage, converter into EEStor's EESU, that will be used in Zenn Motor Company's small to medium-sized electric car.

And then there was a - they were also awarded a contract, a Varian contract for high-powered solid-state modulators and so forth. So look into the company, it looks very real. And associated with some comments that were made is the assumption that in September or October there's going to be a big announcement. So we may be close to seeing one of these things actually working, which would just be spectacular, in my opinion.

We'll just remind our users, or our listeners, that the whole technology here, the idea is that you have a capacitor that is extremely large in terms of its capacitance, and at the same time is able to store a tremendously high voltage without it shorting out or breaking down because the amount of power that is stored goes up with the square of the voltage that the capacitors can be charged to. So you need both extremely high voltage and high capacity. And potentially we end up with a tremendous breakthrough in energy storage. Which, you know, could affect all of our lives, since we're all carrying things around now that have batteries, and batteries are annoying in so many different

ways.

Leo: So this is exciting. You think maybe in the next couple of months we're going to see something.

Steve: I think we're going to see something, yeah. I mean, I glommed onto this originally, as we all know, because it was like, ooh, this is, I mean, this is like the answer, if they can build these. And we may be close to seeing production of this, which would be great.

One other little blurb I got a kick out of was someone pointed me, a friend actually, to - we've talked about RISCs and CISCs, RISCs being Reduced Instruction Set Computers and CISCs being Complex Instruction Set Computers. It turns out there's such a thing called an OISC.

Leo: What's that?

Steve: Thank you, Leo. That was your cue.

Leo: [Laughing] Wait a minute. Let me think if I can figure that out. OISC. Well, we know it's Instruction Set Computer. Origami? Optional? I don't know. What?

Steve: I love it. It's One Instruction Set Computer.

Leo: Well, that's about as reduced as you can get.

Steve: And what I love about it is, if you have one instruction, you don't need an opcode.

Leo: Yes. Just keep doing what you're doing. How can you have a One Instruction Set Computer?

Steve: And it's Turing complete.

Leo: No, it's not.

Steve: Remember we talked about the voting machine last week and how they had come up with a whole bunch of gadgets by using the code at the end of existing subroutines in order to execute their own code, and that they had enough of them that they had created a Turing complete computer such that they could do anything any other computer could do. Well, it turns out that you can have a one-instruction computer which is Turing complete.

Leo: No.

Steve: There are various choices of instruction. But the typical one, the instruction is subtract and conditional branch. So it's a three...

Leo: So that's two instructions, though.

Steve: Well, no, no, no. It's one instruction. So what it does is, you have three parameters.

Leo: Okay.

Steve: It subtracts the second parameter from the first. And then if the value is negative, it branches to where the third - to the address of the third parameter. So that's the instruction. Subtract and conditional branch. And so I see what you mean. Technically you could call it maybe two instructions.

Leo: It's one big instruction.

Steve: It's one big instruction.

Leo: Okay.

Steve: And the way you solve, I thought - there are several clever things about this. The way you solve the, well, what if I don't want to branch, is, well, the branch target is the next instruction. So if you don't want to branch, you just say, well, the branch is the following instruction.

Leo: Keep going.

Steve: So whether it branches or not, it ends up at the next instruction.

Leo: But it can't really do any work.

Steve: Well, it actually does. It's all - you can, from that one instruction, you can synthesize anything.

Leo: Really.

Steve: For example, because think about it, like you need subtraction as opposed to addition because, if you subtract in the right way, that is equivalent to addition. And you can perform logical operations. You can basically get this to do everything you want. For anyone who's curious, Wikipedia has a great page on OISC computers. People have built them. Someone actually built one out of hardware and programmed it to do things. And there's emulators and simulators and - so this is not a brand new concept. It's something that's been around for a while. I just got a big kick out of it. The one-instruction computer. And mostly what I loved is, like, well, you don't need an opcode. Starting off, right off the bat, no opcode because you don't need to tell it which instruction to execute. It executes the only one it's got, over and over and over.

Leo: Neat.

Steve: And I mentioned, we were talking about 3D technology, and a whole bunch of users wrote in because I was talking about how, well, one way or another you need to give each of our eyes a separate image. I talked about the red-green glasses and the notion of LCD, high-speed LCD shutters flickering. And of course the other technology is polarized glasses, where - and I don't know, probably people who listen to this podcast have messed around with traditional polarized glasses where, you know, as you rotate the lens against the other one, you can see it, like, black out and then come back? Well, those are linearly polarized glasses.

So one approach that had been used was that a special screen is needed, a silverized screen which will not scrambled the polarization as it reflects the projected light back to the viewer. So you have a projector which puts out the two images that are bound for people's left and right eyes with polarized light that is 90 degrees off axis from each other, like one is vertically polarized; the other is horizontally polarized. And the glasses are the same. The problem with that is you have to hold your head exactly...

Leo: Straight.

Steve: ...exactly straight, exactly.

Leo: Yeah.

Steve: Otherwise, as you tilt your head, you see bleed from the wrong polarity.

Leo: Although I imagine it's self-correcting because it starts to look weird as your head tilts. So you...

Steve: Yes. And in fact what has been found is that people quickly learn, I mean, as they turn their head they go, oh, whoops, you know, and so they quickly adapt. Okay, now, here's the part that hurts my head, is that I can understand, the idea of vertical and horizontal polarization. I mean, I've, in the old days of Polaroid sunglasses, I would take two lenses and rotate them, and you could see how it would black out, and then you'd rotate them 90 degrees, and now you can see through; okay? There's a different kind of polarization called "circular" polarization. And...

Leo: Yeah. I have that in filters, you know, in camera filters. You have circularly polarized lenses.

Steve: Okay. Well, you can have clockwise circular...

Leo: Right.

Steve: ...polarization and counterclockwise circular polarization. And they block each other out. Which blows my mind. Because the way the current 3D technology works is there's a digital projector with a special thing in front of it which at many times a second, 144 times per second, so also many times per frame that it's transmitting, this thing is flipping back and forth the polarization, the circular polarization between clockwise and counterclockwise at the same time that the image is being changed.

So this light goes down and hits the silvered screen, which has to have a special screen. You can't use the old-style, glass-beaded screen because those screens do not respect the polarization. They scramble the polarization upon reflection. So you have to have a special screen. Then what comes back to the user is the two separate images, circularly polarized with different spins, however that works. And they're wearing glasses with the matching circular polarized filters.

Leo: I love it.

Steve: So now they can turn their heads, and it doesn't matter. And I cannot conceptually get how that works. But it does.

Leo: You know, photographers are familiar with this because you can buy filters that are circularly polarized filters. And I've even seen demos, you can look on YouTube, where somebody has a polarized filter on a screen, and then as - and one of these filters, puts it in front of a camera and rotates it, and it darkens or lightens, depending on the rotation. So...

Steve: Well, but it wouldn't. That's the problem. What you described, I'm familiar with, a linearly polarized filter. But this circularly polarized filter wouldn't lighten or darken with rotation.

Leo: Well, it only does if it's canceling waves. So if you have a - it takes two filters; right? So it's canceling because, as the rotation goes, it's canceling the other filter. So you're putting them out of sync.

Steve: I don't know. This sounds to me like it would not change as you rotate the filter.

Leo: Okay. You're right. Because now that I think, it'd be a linearly polarized filter

that would because it would go out of sync with the waves.

Steve: Yes. And so this is somehow all on and all off. It wants circularly polarized clockwise light, and all of the counterclockwise light it blocks out.

Leo: Yup.

Steve: Like, that's just so cool. But again, I can't get my brain around how it does that. But, I mean, I can't think of an analogy, I guess, that fits.

Leo: I think a circular polarizer on a camera is different. You're right. I think it's got linear polarization in it.

Steve: Two last things in our errata. There's news about one of my favorite sci-fi authors, our friend Michael McCollum at Scifi-AZ.com, whose books I love. Leo, I respect your decision not to read unfinished multivolume series because it's extremely frustrating.

Leo: It's so hard.

Steve: Oh, and especially, for example, when we did "Pandora's Star," and talk about a cliffhanger, I mean, you were just left thinking, oh, my god, when are we going to get Part 2? And of course none of Peter - I don't think Peter Hamilton's ever written a short book. So, you know, a huge investment. And then you're left hanging. Well, the news is that Michael will have the proof copy of the final book in the Gibraltar Series to me on Saturday. He sent me email, I think it was Friday of last week. And he said, "Hey, Steve, I've heard through the podcast listeners that you're still interested in having a look at editing this text when I'm finished with it." And remember that I had mentioned a couple weeks ago that one of our listeners had noted that Michael had updated his page on his website, saying a few months from now. Well, he said that his reread went much faster than expected. And so upon hearing that on Friday, I said, "Yes, I'm absolutely interested in editing book three. I'd like to get a copy of book two for the Kindle," because I know that his site now offers eBooks in virtually every format ever known. I mean, it's just amazing how many different formats. And he said, "Well, I've got news. It's on Amazon."

Leo: Wow.

Steve: And so he said, "I'd be happy to reimburse you for the cost. But you can get it instantly for your Kindle from Amazon." And I said, "I don't want reimbursement. I'm happy to do this." So I ordered it from Amazon and - both the first and second book, which were instantly delivered.

Then I learned something very cool, which I had never had occasion to learn. It only affects people who have multiple Kindles on a single account. But I've always sort of wondered about the synchronized deal. And it really works beautifully. Because what I

learned is, I can read on my stair climber, which is a breakthrough. So I have the DX, the big-screen Kindle, rubber-banded now to my stair climber's console, which I can no longer see. It covers it up, but I really don't need to see. And so I've been having my hour-plus-long workouts just gleefully reading. And so this is going to allow me to get to the Kim Stanley Robinson, is that the guy, Kim Stanley - the Mars...

Leo: Yes, the Mars - Red, Blue, and "Green Mars," yeah.

Steve: Yes. It'll allow me, because now I have time to read because instead of...

Leo: On the stair climber.

Steve: On the stair climber. And what's cool is that Amazon beautifully synchronizes. They don't have any problem with the same novel being loaded into multiple Kindles on a single account.

Leo: Oh, that's good to know. And when you get to page X on one, the other jumps to page X.

Steve: Well, so what happens is, I'll be reading along for 66 minutes, typically, or plus that, because I normally like to finish whatever aspect I'm in on the stair climber, and I just stop. Then I take my little Kindle, K2, to dinner with me. When I turn it on, up pops a little, I mean, all by itself, it pops up a notice, says, oh, you are currently at the following location in Steve's DX Kindle.

Leo: That's neat.

Steve: Would you like to move there in this Kindle? And so, I mean, so they're formally saying we have no problem with a person having multiple Kindles and the same books on multiple Kindles, which I never really...

Leo: I can't see why they would.

Steve: Yeah.

Leo: I mean, it's just more money for them. But, no, the iPhone app does that, too. And so I kind of knew this because they have a Kindle iPhone app, and it will also synchronize with your Kindle standalone. So I guess they're just extending that feature across all Kindle platforms. That makes sense, yeah.

Steve: Yeah, it's very neat. So I did want to tell our listeners, anyone who has enjoyed Michael's books, anyone who is waiting for Book 3, I don't know from the time I'm through with it how long it'll take him to publish it electronically. I wouldn't think very

long. If anyone, like Leo, has been abstaining from getting into the Gibraltar series for fear that they'd hit the end of the books that were in print or available, yet - and then not have the story ended, fear not because I don't think it's more than a few weeks before number three is done. And, oh. And so I did read number two again to sort of remember where we had left our hero.

Leo: That's my problem. That's my problem. I have to reread it to catch up.

Steve: Yes. And so I reread number one when number two came out. And I was tempted to reread them both, but I wasn't sure if I had time. Turns out I did have time. But it's just a - I will say again, the Gibraltar series is an intriguing plot. I just - I like his work. It's not insanely long and infinitely detailed the way Peter Hamilton's are, where you end up with a massively complex world you're holding in your head. Michael's tends to be more directed toward a plot line. So it's a little thin on unnecessary characterization and unnecessary detail, but it's hard sci-fi at its best.

I'm rereading "The Sails of Tau Ceti" at the moment because I just - now I have to wait for Book 3 to get ready. And I was just - I was loving his description of light sails and the use of light sails for braking and how we use electrostatic fields to gather the hydrogen, interstellar hydrogen and funnel it in. And, I mean, it's just - it's great sci-fi. So I'm looking forward to the Kim Stanley Robinson stuff because you have said that it's very much that way, too.

Leo: Yeah, oh, yeah, yeah. It's all about technology and a lot of hard science in it. I love hard sci-fi.

Steve: And lastly, I just got the ThinkGeek email today, and I always just browse through it to see if there's anything that grabs me. Well, they were announcing a bunch of new T-shirts. And ThinkGeek.com has a bunch of T-shirts. I ordered some of this one that I just loved. And I thought I've got to tell our listeners because I know there are geeks like us who are kind of curmudgeon-y, who for this T-shirt, there's just never been a more perfect T-shirt. It's black, and it has one word, in big, uppercase letters with a period, in white. The word is "NO."

Leo: We were just talking about that at the beginning of the show, learning to say no.

Steve: Just N-O.

Leo: No.

Steve: I just love a black T-shirt that just says no, period.

Leo: No. I will not.

Steve: Just don't, you know...

Leo: Don't ask me.

Steve: And people, of course, will ask - it'll be a conversation starter, too, because people will say, "No what?" Ask me a question.

Leo: Anything. No anything.

Steve: Anyway, I loved it, so.

Leo: I'm going to have to find that one. I see a "No Comment." I see "There's no place like 127.0.0.1," which is one of my favorites.

Steve: Yup, that's an old one.

Leo: "No, I will not fix your computer." But one that just says no, no to everything.

Steve: It's just perfect.

Leo: They're great. I love them, really nice people.

Steve: Yeah.

Leo: Do you want to do a SpinRite letter or...

Steve: I don't see the need. Everyone listening...

Leo: Just buy it.

Steve: ...knows that SpinRite solves problems, fixes drives. A number of our Q&A people mention SpinRite appreciatively. So I thought, ah, that's fine.

Leo: Oh, I found it. I'm going to put a link in the show notes.

Steve: N-O. Isn't it perfect?

Leo: They call it "the shirt of ultimate disambiguation" [laughing]. Yes, that's true. There's just no ambiguity.

Steve: It's big. It's just NO.

Leo: No.

Steve: I just think it's perfect. It's a perfect geek shirt.

Leo: All right, Steve. Are you ready?

Steve: Ready.

Leo: For questions from the audience [trumpeting]. Starting with Craig in Chicago, who is sounding and seeming rather desperate. He says: Hi, Steve. I sent this a few times now that iPig is no longer offering its service, and I need to use my computer in hotspots. I guess iPig was a hotspot VPN type service. Listening to you and Leo, I know I need a VPN, but I can't afford a server. I'm not up to speed on running it. I'm waiting for your VPN service, Steve.

But in the meantime I need a service. I came across this: HotspotShield.com. They offer SSL connect like iPig, but I only trust it if you, Steve, give it an okay. I've been a paid user of SpinRite since the mid-'80s. It's been a lifesaver. That's a long time. It's been a lifesaver. It's 25 years. Please tell us if using Hotspot Shield is okay, and then at least I can relax until your VPN is out. And of course I can't wait until you're offering it. Thanks; and Leo, thanks for your great service. Please respond to this. I think there's others needing a VPN solution. Hotspot Shield. Have you - are you familiar with these guys?

Steve: Well, first of all, Craig, I hope you're listening to this because, Leo, I don't know how many times he has submitted this.

Leo: Oh, dear.

Steve: But every time I look there's this message from Craig. And he hasn't provided me with his email address, so I haven't been able to say, okay, message received, we're going to take care of this. And so it's just over and over and over. So Craig, got it. Here's your answer.

I don't know them, so I went to take a look. And the first thing I see is some raves from CNN and PC Mag and a couple other publications, and say, okay, well, that's something. And all there is is just press this to download. It's like, okay, wait a minute, what's - who's this company?

Leo: Press this to download.

Steve: Why is this a free service? So down at the bottom is AnchorFree. So I go to AnchorFree and find AnchorFree. And it's like, okay, here's a little more information. This is where this Hotspot Shield is coming from. And poke around a little bit. And then I sort of see, okay, somehow this is advertiser supported.

Leo: Oh, boy.

Steve: Hmm. How does that work? So then I go to the advertiser page. And I will share with our listeners, because this is a perfect example of a little bit of simple security research anyone can do. And so the headline says, "Advertise on the AnchorFree Media Network."

Leo: No no no no no no no.

Steve: "Advertise to AnchorFree users." Okay. "Advertise to AnchorFree users who are always connected while on the go. They seek out WiFi, shop for the latest technologies, use VoIP for making calls, and look for mobile connectivity, all while using AnchorFree for security and privacy" - [clearing throat] privacy, privacy - "while surfing the 'Net."

Leo: And you can have access to them.

Steve: Oh, exactly.

Leo: And they're yours.

Steve: Oh, wait, it gets better. It gets better. This is just the warm-up. They are - because we've got contextual advertising happening here in a minute. We know what that means. "They are today's broadband ber-user, and we can touch them."

Leo: Oh, dear.

Steve: Not only, yeah, touch them whether...

Leo: Don't touch me.

Steve: ...they want to be touched there or not.

Leo: Don't touch me there.

Steve: Don't touch me there. "Not only that, we offer something truly defining in online advertising. We provide our users security and privacy while enabling brands to target contextually relevant advertising campaigns to some of the most tech-savvy users online. AnchorFree's technology enables ad placements across any one or more of the domains that are visited by our users."

Okay. That's enough of this. I mean, I've answered my question. Our listeners now understand that what this means is that what you do and where you go is being tracked and monitored and contextualized so that they can choose - apparently they're doing interstitial, actually it somewhere does talk about interstitial advertising. So they're monitoring the websites you visit, and they are changing in some fashion the content of the pages you download to insert their own ads. And that's why this "VPN," unquote, solution is free. That's their model for making it free. So...

Leo: Now, I think, though, to be fair, we do a lot of ad-supported free stuff. Our stuff is ad-supported free. They do disclose; right?

Steve: I think it's probably very clear that anyone using this service...

Leo: Is going to see ads.

Steve: ...will quickly realize that this is what they're doing. So you're right. My outrage is, I guess, at the idea that anything is changing the data in my link. That is, in order to be displaying their ads from their advertisers, then my web browsing is being filtered by them. So, yes, Leo, I think it's entirely fair to say, hey, but the service is free. So that's...

Leo: Right. I mean, Google's free. There's a lot of free stuff.

Steve: Yes.

Leo: That is ad supported. So...

Steve: Yes.

Leo: I guess the most important thing is disclosure, a strong privacy policy that you can read, and it is - this is something I wish...

Steve: And I have to say, I did read the fine print on their privacy policy, and it is entirely one-sided. I mean...

Leo: Yeah. Well, there you go.

Steve: In fact, it didn't even seem to really apply to them that much. It's like they got somebody else's privacy policy.

Leo: And that's the other thing. An independent third party auditing it would be nice, like a trustee or somebody. I mean, we're going to be - I think more and more you're going to see this kind of thing. I mean, this is a model for all broadcasting. It's ad-supported free broadcasting. And as it migrates to the Internet, I mean, we obviously don't collect any information about our users. It would be hard for us to do so, and I certainly have no desire to do so. But I can see, you know, I can see that that's happening. So in this case you're saying stay away.

Steve: Well, no. I guess I'm saying Craig wanted my approval, and I can't give it to a VPN product which is doing this.

Leo: Right.

Steve: I guess I would, if it were some sort of a - I mean, to me a VPN almost seems sacrosanct. It's like...

Leo: I agree, yeah.

Steve: It's like, "Do not mess with my data." And here's a company that says, "We're messing with your data." You're using a VPN. We're securing you until you get to our servers. So we are securing your hotspot connection through an SSL link. We understand that technology. That's probably bulletproof. But once we have your data, and we've decrypted it, we're going to modify it to suit our needs in return for the service that we're providing you. And to me it's like, eh, don't think I like that so much.

Leo: So no approval.

Steve: Yeah.

Leo: No seal of approval.

Steve: No. There are - there's HotSpotVPN. That's a service that we know and like. It's not free, but it's not expensive. And whereas iPig was free, this HotSpotVPN is not free, but very good. It's based on the OpenVPN client and server technology. It's one we've looked at and known about for years. So if someone wants something to use, HotSpotVPN is a service that we've looked at and that is not very expensive. And we've talked about it in the past. You could go to - if you went to GRC, the Security Now! page, and did a search for HotSpotVPN, search for that string, you'll find that we've mentioned

it and talked about it in a great deal of detail in the past.

Leo: Yeah. And both you and I use it.

Steve: Have.

Leo: Or have, yeah. I don't travel anymore. I don't need it.

Steve: Right.

Leo: Question number two, Flash cookies from Bob Carneim in Oak Ridge, Tennessee: Hi, guys. I try to stay ahead of the curve. I've been deleting Flash cookies for, well, probably a couple of years now using the Advanced settings panel of the Flash plug-in. And you can do that if you go to any Flash video, YouTube, for instance. Right-click on it, select Settings, click the Advanced tab, and you can see right there you can modify that. He says: But what's next? What tracking method is out there I don't know about yet? What comes after Flash cookies?

Steve: Well, I just sort of liked the question because it evidenced a maturity of sort of recognition that, if it's not one thing, it's another. I mean, we've gone from browser cookies, now there's Flash cookies. Of course the problem is that we're - over in the Internet Explorer world, we know that ActiveX controls are provided for all kinds of purposes and could easily have their own tracking technology embedded in. In fact, there is something called user persistence objects, or something like that. It's something I've got on my list of tracking technologies to track down and haven't yet. But there very well may be other things coming along.

My hope is that the outrage caused by this kind of undisclosed opt-out approach - we know, for example, when we were talking about the report that came out a few weeks back, the researchers, I think they were at UC Berkeley, who noted that more than half of the most popular sites on the Internet were now using Flash cookies because browser cookies had proven too easy to disable. So they were being deliberately sneaky and using something else to hold onto people.

Now, you could also argue that, for example, a bank wants to be able to maintain log-in information, and that users might naively disable their browser cookies and then no longer be able to use the banking site. So the banking site is just, like, trying to help users to have an experience that they need because they have to have cookies enabled in order to use the site. And they get, like, more tech support problems because people have disabled cookies, and now the site doesn't work. Well, it's like, okay, I mean, there's a dilemma. There's tension between what users want and what the web server-side services want. But this is all sort of part of the immaturity of this technology.

What I hope is happening is that our legislators are beginning to wake up to this issue, and there's signs that they are, such that we'll be protected legally from whatever comes next by having some sort of dialogue where what's going on is explained to us. The Flash problem is that it's something no one expects. It's sort of out of the blue. When you hear about Flash also spying on you, it's like, uh, what? What are you talking about? I turned cookies off. No, you didn't turn Flash cookies off. Those are a whole different cookie.

Leo: It's a whole different kind of cookie.

Steve: Whole different kind of cookie.

Leo: They're always working. I mean, you know, you're exactly right. You know they're working on something. They're always going to be working on something.

Steve: Yeah.

Leo: You know, this is - and again, I mean, I would just say that there are good reasons to - my bank wants to preserve information about me so I don't have to always jump through a lot of hoops to log in. You've used this computer before? Okay. We'll let you get in.

Steve: Right. And it's funny, too, because the sad thing might be that we wouldn't have a unified solution. That is, it might be that a bank would be reduced to requiring their own plug-in in order - that would be run by your browser in order to provide static state information for your browsing session. Well, that would be sad because your bank would need one, or all of your banks would need one. Then eBay would need one, and Amazon would need one, and you'd end up with this big mess of individual plug-ins that your browser ends up lugging around because we were never able to agree upon...

Leo: Right, right.

Steve: ...a clean, uniform, opt-in solution. So instead you were opting in to individual plug-ins, which just was causing a real problem. And it would be sad if that's where we end up with. But I could see us heading there.

Leo: Yeah. Yeah, we need some mechanism.

Steve: Yes.

Leo: Not completely unreasonable to have something that we can do.

Steve: No, it's absolutely required because, as we know, browsing is a stateless act. You ask for this page, and then that's the last the server knows about you. You then click on a link, well, it needs to know that that's you clicking on the link, not somebody else clicking on the link. And there's no persistent connection. Normally it's a stateless event. So something has to provide some state information on a per-browser page transaction basis. So cookies used to do that. Flash cookies are doing that now. But we're incrementally disabling these things, which is going to end up causing a problem for the very real need we have to maintain state.

I mean, I like the fact now that I've got cookies flushing in Firefox, and a simple little cookie manager where I just say, no, the site I'm on, eBay or Amazon or whatever, where I've logged in, I want this site to be able to leave persistent cookies so that when I come back the next day it says, oh, Steve, hi, just log in anyway to make sure that it's still you, not somebody else using your machine, but we assume it is you. And, for example, with eBay you can say keep me logged in for all day, in which case you're able to use it without having to continually reverify who you are.

Leo: Question three from Bill Barnes in Charlotte. I guess that's Charlotte, North Carolina. He says: I'm wondering about punching holes in the wall. I frequently need to get to a computer from the Internet. Okay. I figure, all right, I know who needs access, I'll let them in. Then I open a port in my router, point it at one computer, and open the same port in the computer's firewall. Port forwarding, it's called.

Steve: Yes.

Leo: Then someone challenged me about opening that port in the firewall: "Isn't that an access point the bad guys can get through?" Without considering it deeply, I figured an open port was like a CIA phone number. Someone randomly dials a number, the lady answers with a challenge - 41357. If the caller is unable to provide the correct reply code word, she hangs up. That's the end of the attack. Am I wrong?

Steve: Well, this was an interesting question because what Bill is saying is that he's opening services to the Internet, so that he, wherever he is, and sadly any hackers, wherever they are, are able to access the service running on the machine on his network. So the router that would normally block and provide good security for unsolicited incoming traffic, the problem is, it's blocking Bill, who wants to connect to a service running on his computer. So Bill says no to the router, open this port, and if anything comes into this port, send it to this IP behind the router to this computer.

Then, as we know, for example, your typical personal computer today, whether it's a Mac or Windows or Linux, has a firewall. Well, that's going to stop it again because it's going to be an unsolicited incoming packet. So again you need to say to the firewall running on the computer, no, allow something coming inbound on this port to come on in. So now what's happened is, any traffic out on the internet is able to get all the way into the service running on the computer.

The problem is, everyone on the 'Net has access. So now you've got the issue of, okay, so I have to log in to get to this computer. However, in the best case that's true. History has not demonstrated, unfortunately, that that is the case. For example, just this week there is news of an IIS, that's Microsoft's Internet server, the FTP service running in IIS is vulnerable to attack. So, and IIS is the so-called personal web server that you can turn on and configure in Windows XP Professional, for example. And you can use it locally. But in Bill's scenario he would have mapped the FTP ports through to his machine.

So normally you have to log into an FTP server. Unfortunately, it turns out you don't have to log into Microsoft's. So this happens to be a perfect example of, yes, you'd like to have your CIA lady challenging you, a challenge-response or some sort of log-in. In the best world, that's what you've got. The problem is that time and again we run across mistakes in the coding of that log-in or challenge-response password system that can allow somebody unauthorized to connect to you.

In my opinion, the only safe way to handle this is if, for example, Bill was at work, and he knows his work's network, then you allow a selective opening of the port. You say, allow incoming connections only from this IP range into my local network. The beauty of that for, for example, TCP connections, like FTP uses in this example, is you cannot spoof the IP of an incoming TCP packet and have it succeed. You can make up an IP address, but then in the connection-establishing TCP handshake, the responding packet will go to the IP you spoofed, not back to the spoofer. So that's extremely good security.

But it does narrow Bill's freedom. For example, he would only be able to access the service running on his machine at home from prespecified IP ranges, which, again, the huge security is random hackers scanning the 'Net. For example, you can imagine right now, with this known vulnerability in IIS, there's an uptick of people scanning for FTP ports, hoping to find exposed IIS FTP services that they can immediately use to compromise the service behind. You don't want to be exposed to that on an ongoing basis.

So my feeling is punching holes in the wall indiscriminately can be a very dangerous thing to do. And maybe what you want to do, if it fits your need, is not have that hole punched through to a main machine on your LAN. Have it go to the so-called DMZ, to some machine which is isolated from the network, that might be able to give you some of the freedom that you want or that you need, but if somebody compromised it you'd be less damaged. Although even that is creepy even to say.

Leo: Never creepy to say.

Steve: Yeah. So anyway, Bill ends up saying, "Am I wrong?" It's like, well, it's dangerous. You absolutely need to recognize that it's dangerous. Because vulnerabilities are being found in these sorts of services all the time. And unless you can restrict the port range from which you're making a connection, which does give you very good security, I just think it's too freaky. It's just too frightening. That's really where you want a VPN.

Leo: Opening a server, anytime you open a server on your system you're opening yourself up.

Steve: Yes.

Leo: And you rely on the security that the server provides.

Steve: Yes. And history teaches us that that's not a good thing.

Leo: But nevertheless, people do it all the time. I mean, I have FTP servers running on my NAS, and we port-forward over to it. But we have a log-in. And we just hope that the FTP daemon is secure, and you keep an eye on the updates and all sorts of stuff. Because otherwise, I mean, look. I have a web server running. Not on my local network, but it's running on my - it may be even a more critical network, if I think

about it. So we just trust that we're locking it down. Bear's always looking at security holes, and you harden it, and you do the best you can.

Steve: Yup.

Leo: Joe Dorward in Bracknell Forest, England echoes a common question: Steve, going back over the parents and passwords issue, where people who just don't get it - like elderly parents or whatever - can't be induced to care and take greater precautions. We had a question about that a few weeks ago. I remember you've mentioned in the past that writing down a really good password is better than memorizing a poor one, and something about assessing the threat vector before devising a solution.

Seems to me, then, that parents, or people in general accessing the Internet from their homes, are safe enough writing down a very good password, even putting a post-it note on the screen, on the assumption there's a higher risk of somebody guessing a weak password over the Internet than there is of someone seeing the good password written down in their home. We can't expect most people to adopt a more secure way of life if it's not easy. So getting them to write down good passwords, unless you give them YubiKeys or something, is a pretty good solution, and better than having them use "password" as their password. He makes a good point.

Steve: Yeah. I think it's a very good point because you do have the danger over the 'Net of somebody having the opportunity to guess a weak password. And you could argue that, okay, written down on a post-it note, there's just not that much physical exposure to their written-down password. I would add, and the reason I wanted to bring this up, is that Joe made a very good point, and I wanted to add one more thing, though, that I did say before, just to make sure it's heard. And that is, make a change to the password you write down. Deliberately do something to it. Swap the first and last character. Add something of your own, either to the end or to the beginning. Because you can combine the thing that you've written down, which is bizarre, with special characters and punctuation marks and things, and then always remember to do something custom to it that's easy to remember. And in fact, if you forget, and you put in the password exactly as it is, it'll get rejected, and you go, oh, that's right, I have to add my special incantation to it.

Leo: Right.

Steve: Then you've really got the best of both worlds. You've got something nonguessable, but the part you remember can be guessable because the concatenation of those won't be. And that's a perfect, I think, compromise for people like people's parents who don't want to do separate passwords for everything and so forth.

Leo: Yeah, that's kind of what I do. I don't want to talk too much about how I do it.

Steve: No, not supposed to.

Leo: Yeah, that's kind of what I do. And I have - and this is what parents won't do - secure password stores. And so I wouldn't write them on post-it notes. I keep a secure password store which has its own master password. And that's pretty - and that's kind of the same idea. It's just a hard one to get Mom and Dad to do.

Steve: Yeah.

Leo: Or whoever. And Mom and Dad, if you're the tech literate person in your family, I apologize. To get the kids to do, let's put it that way.

Steve: No disrespect intended.

Leo: No disrespect intended at all. Paul Bye in Rochester, Minnesota is being annoyed by DNS results being altered by ISPs. We've talked about this many times before. He says: Dear Steve, I thought I'd pass this along as you and your listeners of Security Now! might find it interesting. You've discussed this topic on previous shows. I don't remember you ever mentioning this one in particular.

My ISP, Charter Communications, it's a cable company, recently made changes to their DNS so that, like many other ISP-hosted DNS servers, if you put in a hostname, and there's a DNS miss, the hostname isn't in their database, they return an address to their own, quote, "helpful" search page, saying did you mean so and so? I could live with being rerouted to the search page when I'm browsing. It's annoying, but it doesn't cause any major harm. The problem I and so many other people probably have is that this is altering the fundamental way DNS is supposed to work and causing all TCP/IP-based programs that depend on DNS to fail.

The first major problem I hit after this change is trying to connect to my company's network with their VPN client, something I depend on heavily to do my work. I can still connect okay. But now when I type in a hostname that is on the internal company network, but not visible to the outside Internet world, Charter's DNS server doesn't find the hostname, and then returns their search page IP address. Now, obviously this breaks the VPN. So far the only way I've been able to get around this is either manually editing my connection settings to put the VPN-specific DNS IP addresses, which I then have to switch back after I disconnect, so that's inconvenient, or go find some DNS servers that don't behave this way.

You've recommended OpenDNS in the past. I'm sad to say they're doing the same sort of thing. They are. I'll vouch for that because I use OpenDNS. I finally found a post indicating the existence of some publicly available, strong, stable servers that do DNS correctly. And he found this on the DonationCoder forums. We have a longer link which I'll put in the show notes. He says he's switched to those for now. I plan on setting up my own DNS server after hearing that you do this as well, but I thought listeners might want to know about the ones mentioned in this forum post as I found them to be really good and high performing.

Thanks so much to you and Leo for the show. I've been a listener since day one. Look forward to every Thursday when a new episode is available. So what's the deal on this? I didn't realize it could break a VPN.

Steve: Sure. And in fact mostly it's there are problems with non-web protocols because some of these are not very selective. I happen to - I don't know if I would call myself the world's foremost expert on this at the moment. But the DNS benchmark, which is ever so close to being released, and we'll be talking about it before long, has explicit handling and detection of DNS servers that do this. And we've discovered that this is a - "we" meaning myself and the people who are in the GRC.dns newsgroup at GRC. We've discovered that this is something which is becoming more and more common.

I did want to mention, just for the sake of completeness, to Paul that you can create an account with OpenDNS, and you can turn that behavior off. The other thing that's happening is ISPs who are generally moving towards this are also generally providing an opt-out option. So that I might suggest that he check with Charter and see if there's a way to turn this off. Because most ISPs are being made sensitive to this because savvy users are saying, wait a minute, I don't want your darn search page coming up. I want an error, for whatever reason. And so it might very well be, I don't know in the case of Charter directly. There's one big ISP, as you were reading this I was trying to remember the name, I think the name begins with C-o-m, can't remember the name of it, though. Anyway...

Leo: You're joking; right?

Steve: Comcast.

Leo: There you go.

Steve: That's the one. No, I was just drawing a blank.

Leo: Okay.

Steve: You know, I'm getting old.

Leo: I wasn't sure if you were just being cagey or actually had forgotten.

Steve: No, Comcast is now doing this, too. And they're beginning to spread this across the country.

Leo: Of course you know why they do this. They make money. They put ads on that page.

Steve: Absolutely. It's like, whoops, sorry, you've made a typo, heh heh heh, but look at this. Maybe you want to buy one of these. It's like, no, thank you, I just want my error, please. Now, the link that he provided, I followed the link, curious whether he knew about DNS servers I didn't know about. Because one of the cool things that the benchmark, the forthcoming benchmark from me does, is it has, I've forgotten now, like a bunch, maybe a hundred? Maybe it's not a hundred. It's a lot. We have a big - I have a

- it knows about a huge number of publicly available DNS servers. And it compares their performance to your DNS server's. And the idea being that it may be that your ISP has slower DNS servers than are available publicly. So by changing to these publicly available servers, you get better performance for all your Internet stuff.

It's going to be a very cool app. It also tests for and warns you if your ISP servers or any of the public ones, like OpenDNS, are doing this DNS redirection, because that's something you would probably want to be made aware of. So it is the case, for example, with Comcast that you can configure theirs not to do this. So I would suggest that maybe Charter's is the same. But I was saying that on this forum, the IPs they gave was very familiar to people who have been using non-ISP DNS servers: 4.2.2.1 through 4.2.2.6. Those are Level 3 servers, which...

Leo: Oh, I thought it was Verizon. That's Level 3, okay.

Steve: Yeah. Maybe it is. I think there was some - there might have been some ownership change because now you say that I...

Leo: Let me ping it and see.

Steve: Well, actually, and not surprisingly, my DNS benchmark determines the ownership of all of the DNS servers that it finds.

Leo: It's still Level 3.

Steve: Okay, still Level 3. I thought so. Anyway, so...

Leo: And they don't mind if you do that?

Steve: Well, they're open to the public. And, I mean, I've seen the IPs given around a lot. The problem is they're not quite as stable and reliable as you might think. And once again, the DNS benchmark, which will be free when I get it documented, basically it's finished, I just need to get the documentation done because it's got so many bells and whistles in it. It's even able to determine the reliability of all of the DNS servers from your vantage point and in addition to ranking their performance in a number of different parameters. It's turned out to be very cool. And it's where months of my time has gone because I think it's going to end up being so important.

But our users have found that these Level 3 servers are not as reliable as they thought. And specifically some of them don't - it uses a technology called "anycast," where you have fixed IPs that hopefully find a DNS server that's local to you. But it turns out that some, I think .3, for example, is problematical for, like, a lot of people. But anyway, we will be talking about this in more detail in the future, soon as I get the benchmark documentation finished. In the meantime, I would suggest that people who are seeing this behavior check on their ISP's pages to see if there's a way they can configure this behavior off, if they don't see it as a benefit.

Leo: Is it only VPNs that have a problem with this? Are there other...

Steve: No. It would be any technology. See, the idea is your system on your behalf looks up an IP. If it is not found, then you receive an IP instead for a web page which is the search results. Well, what that means is that any programs, other than a web browser, are going to get really confused by this. A web browser will show you the page that the search engine provides. However, other things, like an FTP client, a chat client, I mean, anything essentially else, will get this IP which is bogus, and they'll try to connect to it.

Leo: You get an HTTP page, yeah, so it does - I mean an HTML page which you can't interpret. What's this?

Steve: Yeah, exactly. And so they don't get - if instead they received an error message, a DNS error, then they could present you with a dialogue box, which they probably are configured to do so, saying hey, you just typed in a chat URL that's incorrect. Instead they assume that the IP is correctly mapped to the URL, I mean, it really does break DNS. The purists, the old curmudgeons among us, are upset with this because it breaks DNS. This is not the way DNS works. Yet it's a creeping, as you said, it is a revenue source that ISPs are increasingly waking up to and going, hey, we can do that, too. Let's get our little piece of the pie.

Leo: And, you know, I think it's not completely disingenuous of them to say it also is better for users because they don't get a 404, they get something more useful back. You know, the mom-and-pop thing. But, now, if it breaks everything but HTTP, that's not a good thing.

Steve: And it's interesting because my probe in the DNS benchmark, it was originally doing a test for www.subdomain.something.com, and looking to verify that an error was returned. It turned out that there were some smarter ISPs which would look to see whether you had www or not. But that was causing me to miss some redirections. So then I changed my code to remove the www and just look for domain.com, like bogusdomain.com, and verify that we got back an error. And so that's the way it had been for a couple months. Then somebody - we realized that there were - yet another iteration was that that would not return the error, but the www would. So now the benchmark is doing both. So anyway, we'll be talking about this in more detail when I'm ready to unveil the benchmark to our listeners, which is just a matter of me having a little time to get it documented. But the technology is in place.

Leo: Rorx points out that this is really something more appropriate for the web browser to do than the DNS to do. The web browser, like Internet Explorer, can come back with an MSN page and say, oh, no, you meant this.

Steve: That's a beautiful, a beautiful example. That puts the responsibility where it ought to be so that you're not crippling all the other applications on your machine that don't know how to interpret an IP address coming back that's wrong to a question they ask.

Leo: Problem is Comcast makes no money on that.

Steve: Right.

Leo: Here's a question, I guess for me, from Rod Duckworth in Sydney, Australia. He wants to know about old shows. He says: Steve, you won't remember me, but I have spoken to you on the phone a couple of times some years back re SpinRite. I'm a long-time user. Always loved SpinRite. Thank you. I have owned an IT company called Hi-Speed Networking in Sydney, Australia for some 20 years now. I employ about 15 people and have been using SpinRite since the first available version. I'm licensed up to v6. I've also been listening to Security Now! since Episode 1, as I'm always on your site, keeping abreast of what's happening, using ShieldsUP! when I'm on-site at clients, as I, too, specialize in IT security and ethical hacking.

By the way, not now, but I have some ripper SpinRite testimonials and stories that I'll send you at some stage for Security Now! that'll be fantastic for you when I get around to writing them down. However, what I wanted to know was, although I can download all the old versions of podcasts via your website in MP3 mode, iTunes only lets me go back 20 or so. I'll explain why that is, by the way. Although I can - it's not iTunes, it's us. I can and have got these audio files from your site. They won't load into iTunes as a podcast as such. Okay.

See, I had some time off from Security Now! some time back due to personal issues that occupied my life for a while. I haven't actually missed listening, it's just I don't have as many of them as - I don't have all the shows as podcasts. He also says: I occasionally get over to Long Beach. That's a long trip from Australia.

Steve: Yeah.

Leo: So next time I'm in there I'm going to pop into Starbucks and have a coffee with you and thank you personally for the podcast and SpinRite. Meantime, how can I point my iTunes at a podcast source that has all the episodes in podcast format for me to download via iTunes? Is there any way? Also, you mind if I put the MP3 on my site in a secure members area?

So let me explain what goes on. iTunes is dumb. It doesn't know anything. We don't make our feed longer than 20 shows on any of the shows because, as the feed - we could have every show in there. But the feed, the RSS for the feed would be hundreds and hundreds of kilobytes. Probably a megabyte if I put all the shows. Then that means every time you check the RSS, which could be several times a day...

Steve: Ah, you download the entire file.

Leo: Download the entire file. Which is considerable bandwidth for us and for you, especially if you're in Australia and you have bandwidth caps. There's no - I don't see any reason to offer a megabyte or two megabyte RSS feed. So the RSS feed,

and I think this was really the intent of RSS, you know, you look at RSS feeds from websites, for instance, they don't have everything ever published on the website. It's just the most recent X articles. In our case it's the most recent 20 podcasts. And that's what iTunes is using. Now, if you keep iTunes running all the time, it will update that, and you'll, you know, if you've been running iTunes and had the subscription to Security Now! since day one, your iTunes will contain a listing of all the shows. Because it just updates the listing. But the most recent RSS feed only contains 20 shows.

Now, he has a separate issue, which is when he downloads them from - I'm not sure why this is happening - from your site - try it from our site. I'm not sure if there's a difference. But those shows are not, he says, showing up as podcasts. You can go into iTunes and say in the info setting this is a podcast, check a box. What that does is it puts it in the podcast folder. It changes how synchronization occurs. It also makes it a spoken word file, which means instead of starting at the beginning each time it bookmarks where you left off and went back to that point.

Steve: Ooh, that sounds like a good option to have.

Leo: Yeah. I mean, it's the same with audio books. It's just a checkbox on iTunes, I'm pretty sure. They used to make it a very obscure tag, kind of an iTunes-specific tag. But now it's just a checkbox. So that's all you need to do. Download them, import them into iTunes, select them all, get info, check the box that says these are podcasts. It'll treat them properly from then on.

You know, you can get any show. Every show on our network is available through the TWiT website directly, if you know the naming scheme. It's always TWiT.tv slash the initials of the show, in this case "sn," followed by the show number. So this show, which is, what, Episode 212, is at TWiT.tv/sn212. And it goes all the way back to sn1, sn2, sn3. So they're all there. They're all - and in fact, if you look at the naming scheme from our server, you could even do this automatically with a little script because we don't change - the file naming scheme is always the same. So if you look at our file naming scheme, which is a little longer because it goes to CacheFly, or no, I'm sorry, it's from AOL, isn't it. So it's AOL, and then there's a redirect in there, blah blah blah. But if you look at that, the only thing that ever changes is the show number. So you could just manipulate the show numbers, and you can get any file directly. You could write a cURL script that would step through them all and download them all at once. So they're all there on the server, but the RSS feed never contains all the shows. It just would be horrendous.

Steve: Yeah, that makes absolute sense. I wanted to make one comment. He asks if we mind if he puts the MP3 on his site in a secure members area. And the only downside for us of that is that we would not get credit for the count of those downloads.

Leo: Yeah, we prefer you didn't.

Steve: It's better for us if you copy the links because then the link does route through Podtrac so that they're able to count the number, and that way they know how many people are listening. And then our sponsors are able to say, oh, this podcast has this

number of listeners, it's worth this much to us.

Leo: Right.

Steve: And that makes it worth that much to us.

Leo: Technically our license, we use a Creative Commons noncommercial attribution share alike, allows you really to do anything you want with the podcast. Because we're really more interested in people getting it out there.

Steve: And sharing it.

Leo: Yeah, and sharing it. But we'd ask as a courtesy, if you are going to put it on your web page, use the link that we use, which you'll see starts with Podtrac.com. And all that happens is Podtrac, every time the show is downloaded from any source at all, a little counter increments. Podtrac does - I should say in the interests of, since you guys are all smart and understand this stuff, in the interests of full disclosure, Podtrac has a IP database of unique IP addresses. And we are trying, for advertisers purposes, you could download it 20 times. We only count that as once. So it's counting unique addresses from the IP. But we don't in fact save your IP address. There's no log. We're not saying who's downloading it. We just do that, compare it to an IP database...

Steve: It's only to get an honest count.

Leo: It's to get an honest, unique count. So when we say, for instance, 80,000 people listen to this show, it's not 80,000 downloads. In fact the download number is probably three or 400,000. It's 80,000 unique listeners. So, and that's something advertisers of course want us to do. A lot of shows will give out their download numbers because they're way inflated. We don't do that. So that explains it all. We don't and can't stop you from doing that, but we'd ask you if you would to please just put the link. We don't mind the bandwidth. We've got the bandwidth.

Steve: Well, and in fact there's a nice compromise, too. I would say if for some reason you're worried that the MP3 files will ever go away, you could certainly keep a local copy of them, but use our links for users to access them. So that if they ever went away, I don't know why they ever would, but then you've got your own backup copy. But the live access uses our links so that the counters get incremented.

Leo: Yeah. I don't think they're ever going to go away.

Steve: I don't think so.

Leo: I guess AOL, which provides our bandwidth, could at some point stop hosting, and they might have to move to another location. But I will do my best to make sure that the podcasts continue.

Steve: Well, and I'm insulated from that. My technology that I've got, I actually have a switch in the registry of the server. I just turn it off, and everything gets hosted locally. So it's just a matter of, if that ever happens, I just change a switch, and my own redirection gets shut down. So I'm prepared for that eventuality, as well.

Leo: Great.

Steve: Yeah.

Leo: Question seven, from John Prince in Somerset, UK. He had a disquieting dialogue with Netgear about this WPA crack we were talking about?

Steve: Yeah.

Leo: Hello, Steve. As an avid listener to Security Now! since the beginning, Episode 1, I know you'll be discussing this matter on the show. I thought you might be interested in the reply I received from Netgear about my own router, which is only a couple of years old. I realize that this vulnerability is not in the wild just yet, but I thought I'd make inquiries now rather than get into panic mode at a later stage. Seems a shame that they have to wash their hands of responsibility for their product, and that I may have to fork out hard-earned cash to replace a piece of equipment which is in otherwise good order. If you have any advice in the matter, I'm sure I and all your other listeners would be most interested in hearing it. P.S.: If I have to replace my router, I doubt this company will now be my first choice as a supplier.

Here's the reply from Netgear: "Thank you for choosing Netgear. My name is Amandeep, and I will be your support engineer. I appreciate the opportunity to assist you. Regarding your concern, please note that the DG834Gv2 is an end-of-life product, and there are no further updates planned for the router." He said it was only two years old, so that's pretty quick.

"Therefore, in order to get the WPA2 security on your network, I would request you to please upgrade your router with a new one." Then he gave some newer model numbers. "I believe this answers your query. If you need any further help, please email us back, et cetera, et cetera." Well, that's interesting. Two years old, you'd think they'd have WPA2 in it.

Steve: That's what I'm curious about. I mean, that's one of the things is that, if it's really only a few years old, that's certainly long enough for them to have, I mean, for them to have before now updated. I mean, even if it were a couple years old and it were - say that it was four years old, well, WPA2's been around, like when the router was a couple years old, to get itself updated. One possibility, strange as it seems, is that maybe

it's unable, the hardware is unable to handle AES encryption. It is the case that TKIP encryption of the less capable WPA is easier on hardware.

John asked for some advice. And the one thing you could do is use one of the routers that has an open alternative firmware. And that's going to inherently give you, I think, really good longevity because the open software community will keep it current and keep it patched and keep adding features to it. As the state of the art moves forward, I would tend to think you're much better off with that approach, if you're unwilling to continually march forward with the obsolete hardware problem. I mean, it's annoying that they're saying, well, it's end of life. At the same point, you can understand that they're spitting out so many different routers constantly, which is another problem that's sort of annoying, is that they're having to just discontinue support for the older ones.

Leo: Right. Yeah, firmware would be easy enough to do. But, look, every company has to make this decision at some point. It's expensive.

Steve: Exactly.

Leo: You can't support everything you've ever made.

Steve: Exactly.

Leo: But two years does seem a little short.

Steve: Yeah.

Leo: Last question. Grant McMillan in Brisbane, Australia wonders about decrypting today's data trivially in the future. Now, this is a good question. I like this. Hi, Steve. Given the exponential growth in computing power these days it seems that any encryption method used today could eventually be cracked quite easily. Will it be possible for a person to record packets today with the intention of cracking them once it's trivial in years or even decades from now? Thanks for the great podcast. That's a very good point. I mean, computer power doubles every 18 months.

Steve: Yeah, it's a great point. And what I liked about the question was notice that he talks about recording packets today. That is to say that, rather than ignoring the packets today because we can't decrypt them, save them. They're encrypted, but save them. Because it may very well be that, as we move forward - and we're seeing examples of this all the time. For example, we just talked about how we're no longer going to use MD4, or Google's Chrome is no longer going to allow MD4 signatures on SSL certificates for browser surfing because now we know that there's some chance in some situations that it's possible to forge certificates in order to forge the identities of the endpoints that you're connecting to with SSL. So the idea is that we're obsoleting those and replacing those with newer technology.

Well, there's no way to obsolete encrypted data which you have stored because you've saved it. And so at some point in the future it may be that AES encryption is weakened,

and that the world, for example, moves away from it to something else. Well, any traffic that is current at the time will follow that migration to the stronger cipher. But old traffic, which may not any longer be used in real time, but if you saved it in a safe somewhere, in an archive, at this point in the future it's like, ah. Now, finally, I can essentially turn back the clock. It's like having a time machine going back to stuff you had saved and for some reason felt or had reason to believe was very valuable, and now you can decrypt it. So it's a really great question.

One of the things that came up in last week's episode about voting machines was the notion, and in the researcher's mind they were exactly focused on this problem, that is, this design of this voting machine was 20 years old. How had it survived the evolution of technology during that 20 years? And, for example, this whole concept of return-oriented programming that we talked about, that had been in - that was invented relatively recently. So they couldn't really protect against something that they couldn't anticipate.

And similarly, the point was made in their article that RAM at the time, these RAM cartridges, these voting cartridges, had static RAM with batteries in them and didn't have very much RAM because physically RAM was much bigger then than it is now. But that now you could imagine an entirely different technology. In something the size of that cartridge which was barely able to hold RAM, you could put a whole Cray supercomputer in a cartridge of that size using today's technology.

So the question is, did the technology then, was it designed to be robust enough to survive during its application, during its use lifetime, can it survive all the forward motion of technological progress during that time. And that's really not something we've talked about before, but that is, it's a really great question. And it is absolutely foremost in the minds of people who think about what's the vulnerability. It's not just today, but it's either - it's technology which is locked in place today needing to survive during the technology's youthful lifetime. But the notion of saving encrypted data today on the off chance that some point in the future it will be decryptable, that's also very important.

Leo: Yeah.

Steve: Great question.

Leo: Yeah. So assume that at some point - but maybe you won't care in a hundred years.

Steve: Well, the downside, or the counter factor to that is he says, "Hi, Steve. Given the exponential growth in computing power...." Well, I will remind our listeners again that, even though it's so easy to add bits - oh, look we just added some bits to this key - every bit you add is exponential growth. Every bit doubles the number of possible combinations of the key. And so that's exponential. And we're saying that, okay, today a 64-bit key is probably not safe. A 128-bit key, okay. So we go, wait, we only doubled the length. We only made it - we only added 64 bits. But oh, my God, that is so much stronger than 64 bits. So it's deceptive how little we need to lengthen symmetric cipher keys in order to dramatically, that is, exponentially increase their strength. So when you go to 256 bits, forget about it.

Leo: So that's - I've often, you know, people have said should you - 1024 is plenty. Why should you use 2048? There's a good reason. Not for now.

Steve: Yes. It's not for today.

Leo: Yes, 1024 is plenty.

Steve: Yes. But it's not for today, it's for tomorrow.

Leo: Yeah. That's really interesting. Very, very interesting. Well, Steve, we've come to the end of this fabulation of fabulous questions. I just made that word up.

Steve: Uh, yeah. Confabulation would not be a made-up word.

Leo: Confabulation. Conflagration. Next week, GSM cracking.

Steve: Yes.

Leo: Should be a lot of fun. Our Topic In Depth.

Steve: TID, Topic In Depth.

Leo: Meanwhile, if you want this show or any of the past 212 episodes, you can get those from Steve's site. He has 16KB versions available for quick download. Quality goes down, but at least they're small. He also has transcripts, which are even quicker downloads, and a great way to kind of search and follow and figure out what's going on. Those are all at GRC.com. If you want to ask a question for future feedback episodes, it's GRC.com/feedback. And of course don't forget SpinRite's there, and all of the great software that Steve does. Most of it is free. And certainly SpinRite is well worth the money. If you've got a hard drive, you really should have SpinRite to keep it running in tiptop shape.

Steve: Sooner or later you're probably going to need it.

Leo: You bet. GRC.com, the Gibson Research Corporation. And Steve, we'll see you next week for another great episode of Security Now!.

Steve: Absolutely. Thanks, Leo.