Listener Feedback #155

Description: Steve and Tom 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-379.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-379-lq.mp3

SHOW TEASE: Coming up on Security Now!, I'm Tom Merritt filling in one last time for Leo Laporte after a couple of weeks of absence. But we've got some great stuff to talk about. We're going to be following the continuing saga of John McAfee. Also, NASA may be providing a good example to the rest of government. And, how to keep your nuclear facility safer. All that and more coming up.

TOM MERRITT: This is Security Now! with Steve Gibson, Episode 379, recorded November 21st, 2012: Your questions, Steve's answers, #155.

It's time for Security Now!, the show that tries to keep you safe online. I'm a little sad because it's my last week filling in for Leo Laporte. I'm Tom Merritt, and of course joined by Mr. Steve Gibson of GRC.com. Steve, it's great to have one last shot at hosting Security Now! with you this week.

Steve Gibson: Well, now, after you move down closer to me, but really still not close enough that it's practical for us to get together for a show, and Leo's gone in the future, would it still work for us to be connected together up at central headquarters?

TOM: I don't see why not, yeah.

Steve: Yeah, okay, cool. In that case...

TOM: All right. There's hope for the future.

Steve: Right.

TOM: We've got some great stuff to talk about today. Of course it's a question-and-answer show, so we've got lots of good questions. We'll also - we're going to dig deep into Mozilla Firefox v17. Got a lot of stuff to talk about there. I guess not dig deep, but there's a lot of updates there.

Steve: And somehow, coming to you direct from Belize, John McAfee is still managing to post a blog. And I tweeted a link to what he posted two days ago. But we have an audio podcast, and people are commuting and jogging and gardening and doing whatever they're doing. So I'm going to share, two days ago, the increasingly bizarre story of John McAfee, who everyone knows as the antivirus person, now not so much.

TOM: Yeah, it's a movie, maybe a made-for-TV movie, but it's a moving in the making, what's going on with that guy. It's kind of ridiculous. All right, Steve. Let's start off. What's good to know about in the latest of the faster and faster releases of Firefox?

Steve: It's funny, too, because it just was not that long ago that we were at v3 and v4. And now we have v17. So, I mean, and there were complaints from the industry, that is, the users of Firefox, that they seemed to have a much too slow release cycle. And so they changed that. And I think that's a good thing. We are getting improvements out faster. We know that having competition is good. And I think they probably now see Chrome from Google as their primary competition, rather than IE, even though IE technically still has greater than 50 percent of the overall market share. Firefox is the alternative, as is Chrome, of course; Chrome is continuing to come on very strong. One of the things that happened in v17, which was released, I think, like yesterday...

TOM: Yeah, I think I saw the alert pop up to get it. I haven't actually updated it yet.

Steve: Yeah, it just did. And when I went over, under Help, to check, like, Help About, it said, oh, we have a new update. It was 9.9MB, I think. And so it sucked it down. It was compatible with all of my add-ons, so it was a painless transfer. For me, at least, it wasn't happening automatically. That's promised. And I had to do a full restart, which meant that all of my 80 tabs - yes, eight zero tabs - needed to get flushed and reloaded. And I also saved a bunch of memory because it still seems to sort of accumulate memory over time. But that problem is way better than it used to be, like back in v4. They really have spent some time on memory.

In this case, they've moved their so-called "Click-to-Play" feature, which we have talked about before, which has been in beta. It's now at release. So it appeared finally in this current main release stream of Firefox. And what that does is it's their newest solution for dealing with vulnerable or outdated plugins. So plugins that the browser recognizes as no longer the most recent and/or vulnerable, the browser will block their execution, so that they don't just automatically run, until the user clicks on it. So the idea is just - it's not to shut them down completely. The Mozilla folks feel that would be a little overreaching. But not to have, for example, them just automatically take off and run.

They posted back on October 11th, when this was in beta, they said, "Plugins that are blocked with the click-to-play flag will not load" - that is, onto the page - "by default, but can be easily activated by users. This gives us a more user-friendly way to warn about plugins that should be updated, that give users better control over their browsing experience. The large-scale plugin update notification we deployed last week" - so this would have been earlier in October that they're referring to - "used the old update notification mechanism for Firefox 16 and prior, and the new click-to-play mechanism for Firefox 17 and above." So this is where they were beginning to roll this out and test this for about a month and a half. And they said, "If you have old versions of Flash, Adobe Reader or Silverlight, and you're on v17, you will now see the click-to-play block next time you visit a page that uses one of these plugins."

Now, again, I haven't been to Silverlight, and I'm not using Adobe Reader any longer. I moved away from that. So I have not seen this behavior in 17. And I don't know when the browser will update itself for people who don't go and get it. But if you're curious, by

all means, just go to the Help About box, and that'll sort of wake it up and say, oh, look, there's something new here for you.

TOM: I just tried that, actually, on this machine. And it didn't do it. This is on Ubuntu, though. So that may be a different kind of behavior. But I couldn't make it tell me...

Steve: Well, yeah. And also, if you go to the add-ons tab in Firefox, it'll list all of the add-ons that you've got. Oh, I'm sorry, plug-ins. The plug-ins tab. There's a link at the top that you can click to check to see if any of your plug-ins are old. I did that, and it gives you a - it does a very nice test, gives you a nice list sorted from vulnerable and bad to okay, and then unknown sort of are grayed out. The problem is, it was showing me that Adobe's Flash needed to be updated. So then I went over to Adobe.com/flash, where you can check the version, and the Flash player ran with a little bouncing red square or cube. So it's like, okay, well, this doesn't really seem to be working yet. Or maybe it needs to be turned on or something.

TOM: My Shockwave Flash says it's up to date when I'm looking at it right now. But again, that's under Linux. So that may be true. In fact, I think it is true.

Steve: So another new feature that I looked around to try to explain what this was to myself and to our audience, but they're saying that their first revision of the, quote, "Social API" and support for Facebook Messenger is now in Firefox 17. There's no sign of it in any of the user interface. So it's unclear exactly what it is. But it appears to be about allowing the browser to post links to the user's previously configured social media accounts.

TOM: Okay.

Steve: So you would somehow register yourself, like your Facebook page or your Facebook account, with Firefox. And then what they're trying to do is to create some sort of a more seamless, lower friction interconnection to make it, like, you right-click on a link and then say, like, "Post to my Facebook page," and it just does it through some sort of API of some sort.

TOM: Anytime Facebook comes up, though, I have a hundred questions about how it works. I'm interested to see that, too.

Steve: And then they have a handful of sort of just updates. The Awesome Bar gets larger icons. They dropped support for Mac OS X v10.5. They are now supporting an important attribute that's defined in the HTML5 spec, and it's a formal W3C spec. I think it's Chrome and Safari have already been supporting it, and Opera and Firefox and IE don't. But now Firefox does. And the page that I went to see where the support stood hadn't yet been updated because it was just yesterday. But that's a new attribute in the infamous iFrame tag in HTML.

To remind our listeners, iFrame is a widely exploited problem in HTML, only because it's so powerful. Essentially, an iFrame - "I" stands for inline. And so it essentially allows you to embed another web page inside of the current web page. You can create a region on the web page and load that with a URL using the iFrame tag. And the problem is it's powerful, but it's dangerous. There have been all kinds of ways that have been found to exploit this power. And so the new spec, or Firefox's new support of the so-called "sandbox" attribute, allows the browser to essentially sandbox the iFrame tag. If you simply put the phrase "sandbox=" and then a null string, open quote/closed quote with nothing in it [sandbox=""], then that enables a new set of extra restrictions for the

content of the inline frame. And then if you wish to deliberately, as the author of the page, if you wish to relax those protections, the sandboxing restrictions, you can then add some "allow" verbs in that double-quoted expression.

So, for example, if it's empty, then all restrictions are applied. And those restrictions are, for example, if you had "allow-same-origin" inside the double quotes, then that would allow the iFrame content to be treated as being from the same origin as the containing document. And of course that's risky because it's same-origin protection which protects us from all of the cross-site hacks which people have come up with over time. Or you can say "allow-top-navigation," which allows the iFrame content to load content from the containing document. So that sort of gives it visibility outside of its own frame, which, again, could be useful, but use that with caution. Then you can also say "allow-forms," another dangerous thing to do. So the point is, if you don't put "allow-forms" in, then forms will not function in the iFrame, and that's been another source of real exploitation in the past. And, finally, "allow-scripts," where you have to explicitly allow scripting. If you use the sandbox tag, then it will turn scripting interpretation on within that frame.

So this is a nice move forward. And now we just need to get Opera and IE to support it. And then, of course, since it doesn't just happen automatically, you don't get any of that protection unless you add the sandbox tag with a null string. So then we need everybody who's using iFrames to update their own code to put that in, and we'll get more safe pages.

TOM: So this is not a browser user protection. This is a page-creator protection, to say, look, you don't want your pages to get hacked through an iFrame.

Steve: Yes. So, yeah, exactly. This is an authoring protection. And there are so many pages now which are based on templates. Like, for example, most of the online forums have a very common, standardized look. It's very easy, then, to modify the templates wherever they use an iFrame, just to add the sandbox tag. And suddenly, throughout the entire forum site, that sandbox tag will appear. And you want to be careful how you use it because you don't want to break things that rely on iFrame features. But you also want to always restrict everything that you're not actually using. So this is just a good move forward. And as you say, it'll take some while for it to get adopted. But ultimately I expect it will be. So that's a nice thing.

And then they also, with v17, they've made it faster. And they say more than 20 performance improvements. And then there were apparently some problems they were having with page scrolling on sites with fixed headers, which v17 in Firefox fixes. So, all around, nice solid update from our friends...

TOM: These are exciting updates, but they are not nearly as exciting as the saga of John McAfee, which we've sort of been following now.

Steve: Okay. Now, if this were a podcast on April 1st, we might be accused of making this whole thing up.

TOM: Nobody'd believe us, you're right.

Steve: Yeah. So get a load of this, listeners. This is two days ago, posted on November 19, on John McAfee's blog. And, by the way, anyone can check it for prior and subsequent postings. Earlier ones exist and later ones exist. The site is www.whoismcafee.com. And if you just do www.whoismcafee.com, that takes you to sort of the main page of the blog, which will show everything current. If you scroll down, you

will find the posting that I'm going to share now, which John put up two days ago, on the 19th.

He wrote: "The first two days, Sam and I were on the run." And if you're curious, you can find all about Sam and who she is and where she came from and pictures of her and so forth. "The first two days Sam and I were on the run were far from our house. I felt helpless, especially given the fact that so many of our friends and workers were being arrested. I realized that, unless I knew moment by moment what was happening, my chances of coming out of this intact, both emotionally and physically, were slim. I needed to be close to area where" - I'm just going to read this with his typos in there. "I realized that, unless I knew moment by moment what was happening, my chances of coming out of this intact ... were slim. I needed to be close to area" - I guess he meant "the" area - "where the events occurred and needed to watch and hear the actions of the authorities. I also needed to do my own investigation, since the police only seem to be interested in my whereabouts. My safety is contingent on the truth being discovered. I today announced on NBC Television that I am offering a $25,000 reward for the capture of the person or persons responsible for Mr. Faul's murder." And that's the event that we shared when we opened this topic last week. "After two days we returned to the house in disguise, and I began my watch."

TOM: Now, wait a minute. He's posting to a blog. He's talking to a Wired reporter on the phone. He's on NBC TV, and he went back to his house?

Steve: Yeah. Get a load of the disguise. "The first day I colored my full beard and my hair light grey, almost white. I darkened the skin of my face, neck, and hands carefully with shoe polish and put on an LA Saints baseball cap with the brim facing backwards..."

TOM: Who are the LA Saints? Oh, Louisiana, okay, thank you.

Steve: Oh, good, yes, right, Louisiana, "and tufts of the front of my hair sticking out unkempt through the band. I stuffed my cheeks with chewed bubblegum stuck to the outside of my upper and lower molars, making my face appear much fatter. I darkened and browned my front teeth. I stuffed a shaved-down tampon deep into my right nostril and dyed the tip dark brown, giving my nose an awkward, lopsided, disgusting appearance."

TOM: This is like movie makeup.

Steve: "I put on a pair of ragged brown pants with holes patched and darned. I wore an old, ragged, long-sleeved shirt. I donned an old Guatemalan-style serape and toted a bag containing a variety of Guatemalan woven goods. I adjusted my posture so that I appeared a good six inches shorter than my actual height and slowly walked up and down the beach with a pronounced limp, pushing an old single-speed bicycle and peddling my wares to tourists and reporters using a broken English with a heavy Spanish accent."

TOM: Hey, he made some money.

Steve: "On my second day, while peddling small wooden carvings, I nearly sold a dolphin carving to an Associated Press reporter standing at the edge of my dock. He was pulled away from the enticement" - of, you know, John's sales pitch...

TOM: Of his amazing dolphins.

Steve: "...by an urgent phone call. Among the people I spoke with that day was the caretaker at Mr. Faul's house." And Mr. Faul is John McAfee's next-door neighbor, and Mr. Faul is the person who was found...

TOM: Who was found dead; right?

Steve: ...shot in the head with a 9mm Luger pistol. "The police had stated that Mr. Faul's housekeeper discovered the body. His caretaker told me that Faul did not have a housekeeper. He himself discovered the body, he said. I found this interesting and filed it away as a piece of data that might help at some point. Why would the police lie about this? Lies always have a reason."

TOM: Yeah. Are you going to read the whole piece here?

Steve: Nah, I think you're right, we're done. Anybody who's interested, you know where to find the rest. It goes on like this, talking about disguises and watching the police and their machinations and so forth.

TOM: He has other disguises, too, when you get down farther.

Steve: Oh, yes.

TOM: Yeah, it's - I don't want to say "insane," but...

Steve: So anyone who's curious about...

TOM: It's good reading. Entertaining.

Steve: ...about John McAfee, the antivirus pioneer of the PC industry who is now on the lam, can find it at whoismcafee.com.

TOM: I like the one where he's got a Hawaiian shirt speaking German. I think that may be...

Steve: Yes, yeah. And there was something about his pants, and then he said like a disgusting Hawaiian shirt, as he put it. And I was chatting with Mark Thompson, a friend of mine, of AnalogX.com...

TOM: Oh, yeah.

Steve: ...just yesterday, and he was like, "Did you see the McAfee blog?" I said, "Yes, I know." And he's like, "He's out in front of his house on the beach, and he's wanted for murder. Like what is he thinking?" It's like, well, he apparently has confidence in his disguise.

TOM: I guess so.

Steve: So we have some Yubico news.

TOM: Oh, yeah. What's up with Yubico?

Steve: And this actually relates to - this is one of the things that I knew was coming, which is why I did a few weeks ago a podcast focusing on near field communications,

NFC, because I wanted to lay some groundwork for what NFC is, exactly how it works, how it relates to WiFi and RFID and so forth. There is a new YubiKey from Yubico that they call NEO. And I don't know if they've actually started to use NEO the character from "The Matrix" or not. But that's the NEO...

TOM: Mr. Anderson.

Steve: Right. That's the character that they sort of refer to. And I chatted with Stina Ehrensvrd, who's the founder and original inspiration for the YubiKey, a few months back when she was down here in Southern California. We got together at Starbucks and had some coffee, and she put me under NDA for a number of things happening. This is one of them.

TOM: Okay. So now it can be told, finally.

Steve: Yes, now I can talk about this. What they've done is they've added a next generation to the YubiKey, which supports NFC, which allows you to do things, for example, with all smartphones except, of course, the iPhone 5, which curiously doesn't support NFC. But my BlackBerry does, and the Galaxies and all these...

TOM: A lot of different Android phones do, definitely, right.

Steve: Yeah. And so this allows you to use, for example, you could use LastPass Mobile on your smartphone with the NEO YubiKey in order to authenticate, to verify to LastPass on an instance-by-instance basis, that you are in physical possession of this YubiKey NEO. And so it gives you the same one-time password. It's got the two slots. Now, it does include something that they call NXP. NXP is the renamed Philips Corporation, Philips Semiconductor, something called SmartMX security technology, which is some public key crypto. They say that it supports CCID compliant USB token behavior. And CCID is the chip or smartcard interface device, which is a standard that's been established that does use some public key/private key stuff.

TOM: Is that like chip-and-PIN? Or is it different?

Steve: I don't know. I know that it's CCID. And also there's something called Mifare. Mifare Classic is a standard technology for physical access control systems, where you wave something near a reader, and that unlocks the door, access control. And so this supports that standard, too. So I'm really pleased because all along the hope has been that we wouldn't have the so-called keychain or key ring full of individual tokens where it's like, okay, sort of looking like a janitor key ring of different active devices. And so they're really sticking to a standards-based approach, which means that, if you were to upgrade your YubiKey to this one - and I don't actually know if there's an upgrade planned. You have to check. I do know that this is not cheap. This little sucker is 50 bucks.

TOM: Just for the device.

Steve: Just for the new YubiKey NEO. And I imagine there are, if you did a corporate deal, there's no doubt quantity discounts and so forth. But I'm glad to see this. And there's more news coming. But probably not till - actually, it's going to be a while, probably not till next summer. So other things are in the works and being tested and so forth.

TOM: Does this let me then just tap the YubiKey against the phone?

Steve: Yes.

TOM: And that unlocks my LastPass, and I'm on my way?

Steve: Yes.

TOM: Wow. I know 50 bucks is not chicken change, but that's pretty compelling. That's pretty cool.

Steve: It's pretty cool, the fact that every phone except, conspicuously, the Apples, would - that supports NFC. What this does is it uses, again, one of the NFC standards that will send your phone a browser link and then launch the browser and provide the browser - the URL contains the one-time password over a secure connection. So SSL protects it, and then you're able to perform the whole transaction securely. So it's very cool.

TOM: Lepton is taking me to task for saying "chicken change." And I think he would have preferred "chump change." I mixed "chicken feed" and "chump change" to form a new - it's an evolving language, Lepton. All right. Let's talk about Facebook.

Steve: Speaking of evolving.

TOM: Yeah, I was happy to hear this news.

Steve: Yes. Facebook announced on their blog, maybe it was just yesterday, it just happened, that they were going to be rolling out, first for all North American users, and then following that for the rest of the world, HTTPS is on by default. So we've talked about the evolution of this, that some time ago, maybe it was - feels like it was maybe as much as a year ago, it was some time ago that Facebook responded to the threat which we talked about a lot which is, for example, represented by - boy, I'm blanking. It was that crazy little gizmo, the add-on for Firefox. Firesheep.

TOM: Oh, right, right, right, right. The one that stopped you from getting your Facebook spied on when you're at the WiFi...

Steve: Oh, and, I mean, it was horrifying. When Firesheep first came out, as an experiment, I took a laptop over to Starbucks, and the entire column populated with the Facebook pictures of all of the people sitting around me at Starbucks. And the reason being that Facebook, when you log into your Facebook account, you have to authenticate. But then to stay active in your Facebook account, it would then drop you back out of SSL, so no more encryption. And the cookie being sent with every query was in the clear. So Firesheep just sniffed the air because Starbucks is famously an open WiFi hotspot like any open WiFi hotspot, where everything is available in the air. So this thing just saw these transactions, grabbed the profiles, and by double-clicking on the person's listing in Firesheep, you could impersonate them. You were logged in as them.

TOM: And that was the big thing. Because some people were like, oh, well, I don't post anything on Facebook that I don't care if people see. That wasn't really the point, although Facebook does give you a presumption that you're restricting who can see your stuff. It was that they could then be you. They had your authentication. It was awful.

Steve: Full impersonation, yes. So the fix to this is, especially in a wireless network environment, where there is no encryption, is to maintain an SSL connection persistently. So Facebook initially made it an option so that security-aware users like all of our

listeners could go into their Facebook profile and turn on "always use secure." And there's really no reason not to, except that all of those SSL negotiations, that connection setup, does represent some overhead. So in Facebook's blog they were boasting about that they've done really cool things on the technology side to minimize the impact. They're a little concerned that users may notice some slowdown. So you can turn this off, if you're crazy.

TOM: Yeah, right.

Steve: I wouldn't turn it off. I'm sure everybody who's listening and using Facebook already has it turned on. But the significant piece here is that, if you had never specified it either way, it will start being on by default. So Facebook is taking proactive responsibility to bring up the cone of silence for their entire billion users, ultimately across the world. So that's really good news. And in bizarre timing, just today, the HTTP transport security standard has been ratified and formalized by the IETF.

TOM: Oh, good news.

Steve: So that is now, again, we're moving forward continuously to increase the security across the board. Google, of course, has been supporting persistent security on their site and searches now for a while. Facebook joins them. And this, of course, just adds pressure for everyone to do this. So it's just 100 percent good news.

TOM: And honestly, the performance hit is not that bad, I don't think. It's almost a psychological thing. I've gotten into arguments with some developers about this before. And if everybody just did HTTPS from the beginning, and nobody did everything else, nobody would notice because it isn't so bad that people are like, wow, these websites just don't work.

Steve: Well, yes.

TOM: It's all a relative thing, I guess is what I'm saying.

Steve: And also that there's so much streamlining and caching. We have covered SSL, or TLS as it's now referred to, transport layer security, the exact protocol, carefully in the past. And the negotiation that the client and the server go through, and the more time-expensive SSL portion, the whole public key stuff, that only really has to happen once within a long period of time because the client and server will agree on credentials, which are cached and then reused.

So I think probably the only concern is that, if you have an SSL page, then you want all of the assets which the page loads also to be over SSL. Otherwise you get that scary sort of mixed security warning. And so, if there were other servers supplying assets with which you had not negotiated SSL connections, then those subsidiary assets could, like, slow down the page load. And they may not be as readily prone to caching as the main page and all of its content. But as you say, it's just like this is something that's only going to get faster as our machines get faster, as the servers improve their performance, as hardware accelerate of public key stuff improves.

And remember that we can move to ECC encryption. The elliptic curve crypto is much faster than traditional crypto. And as clients support it and servers support it, they will automatically notice that each other supports it, and they will choose that preferentially over the slower, older, and even maybe less secure traditional factor the product of primes - I said it right this time - problem that is used to create the security in traditional

public key crypto. So I just think we're going to see this drop away as being a problem before long. And this just represents - this is just good, to have all of our connections being secured for us.

TOM: Now, our next story is a problem that I did not expect to ever see. Did the Naval Observatory go back in time?

Steve: Well, something weird happened. Simon Zerafa, who often sends me interesting tweets and observations - he must spend his entire life just scraping the Internet, looking for interesting things, because he comes up with a lot of interesting stuff. He found a note that that the SANS Internet Storm Center Diary had posted, explaining why people were reporting that the year on their computer clocks was wrong. And what SANS posted on their Internet Storm Center Diary was, "A few people have written in within the past 18 hours about their NTP servers and clients getting set back to the year 2000. The cause of this behavior is that an NTP server at the U.S. Naval Observatory" - which is pretty much the authoritative time source throughout the U.S. - "was rebooted, and somehow its clock reverted to the year 2000."

And then, because of the way NTP protocol works, it's very much like DNS, where you have a hierarchy of time servers. The idea is they don't want to load down the central time server, so there's a whole bunch of second-level time servers that synchronize themselves to the master NTP time server, and then they serve time and so forth down the hierarchy. And so they said, so this mistake of being 12 years in the past then propagated out for a limited time to downstream time sources that also obtained this wrong value. They said it's a transient problem and should already be rectified. "Not much really to report except an error at the top of the food chain," as they put it, "causing problems to the layers below. If you have a problem, just fix the year or resync your NTP server." So that was bizarre.

TOM: So we had the Y2K problem after all. It just wasn't what we expected, and it was 12 years late.

Steve: Exactly.

TOM: I heard about this NASA issue where they lost another laptop. It's not the first time that they've lost laptops, but they're finally taking some definitive action, it sounds like. What's going on there?

Steve: Well, yeah. And I thought I would report it today because today is the deadline, November 21st. And I had meant to talk about it last week. And again, it's one of those things where I had pulled these notes together, and somehow it didn't make the final cut. But on Halloween, October 31st, an unencrypted NASA laptop was stolen from an employee's locked vehicle which contained unencrypted personal data, so-called "PII," Personally Identifiable Information, on a large number of people. It was password-protected, but not encrypted. And it's funny because that was the actual line that was used. And I thought I would note to people that password-protecting the laptop prevents the OS from booting. But all you have to do is attach the drive, put it in a USB case or attach it as an external drive to an already booted system, and it'll go, oh, look, a new drive, and you have full access to the file system.

TOM: Yes. Here's what's on the drive. Would you like to look through it?

Steve: Yeah. Look at all these names and addresses. So this turns out to have been the second time this year - not even in history, but just this year, I mean, historical, NASA's

had a bigger problem, but just this year - that a NASA laptop containing sensitive, in-the-clear information was stolen. About six months ago, in March, a laptop containing names, Social Security numbers, phone numbers, email addresses, dates of birth, college GPAs and other sensitive personal data of NASA employees at NASA's Kennedy Space Center was stolen from the car of a worker at the facility.

And you've got to wonder why NASA seems to be having these problems. Maybe there's a little bit of foreign intrigue going on here because you never know what you might find on a NASA laptop. So what finally happened after this is that NASA said, okay, enough of this. They issued a directive prohibiting the removal of computers from any NASA facility unless whole disk encryption is enabled or all sensitive files are individually encrypted.

TOM: Now, why wouldn't they just say, "All of our hard drives have to be whole-disk encrypted?" Why leave any exceptions?

Steve: That's where they're headed. But bureaucracy, and who knows what logistical problems they might encounter. So CIOs, the chief information officers at all NASA facilities have been instructed to complete disk encryption on the maximum possible number of laptops. Again, that's your typical sort of soft bureaucratic-ese speak.

TOM: Little CYA there, yeah.

Steve: Yes, by November 21st, today, and to add encryption capabilities to all laptops by December 21st. So who knows. Maybe laptops are in employees' homes. Maybe they're in shipment, or they're traveling, so you can't get to them.

TOM: And it's a large enterprise, so it may take a while to make sure you've got everything covered. I get that, yeah.

Steve: And to give you a sense for how creepy this can be, the NASA Inspector General, whose name is Paul Martin, giving testimony before the U.S. House of Representatives, the Committee on Science, Space and Technology, the Subcommittee on Investigations and Oversight, noted...

TOM: Snappy name.

Steve: Yes, oh, yeah. Well, it's, yeah, yeah - noted that in March of 2011 the theft of an unencrypted notebook computer resulted in the exposure of algorithms used to command and control the International Space Station. So if you see anything falling out of the sky...

TOM: Oh, man.

Steve: Oh, goodness.

TOM: Don't even joke. That's horrible.

Steve: And then in another incident, sensitive data regarding NASA's Constellation and Orion programs were similarly compromised when a laptop containing the data was stolen. So hopefully we'll still have a space organization by the time they get this locked down. But they're doing it now. And this has made the news. And this is just - this ends up being good because soon it will be, well, other CEOs will say to their CIOs, "Hey, NASA has all of their laptops encrypted. Why don't we?" So that's good.

TOM: That's a good way of looking at it. They hopefully will set a good trend. Pretty much every laptop should have whole-disk encryption. Every laptop in existence, used by anybody, should have - I'm sure that's the way all of the people in the audience think, too.

The cloud is always a good topic for a security discussion, and there's some news about Google Docs not just being for documents anymore.

Steve: That's right. Now Google Documents can also be used by trojans as their command-and-control server.

TOM: How convenient.

Steve: [Laughing] Hey, it's free. Researchers at Symantec have detected trojan horse malware named Backdoor.Makadocs that is using Google Docs to communicate with its command-and-control infrastructure. The malware appears to be targeting users in Brazil, harvesting and collecting specific data, including the infected computers' host names and operating system types. The malware uses Google Drive's viewer as a proxy to receive instructions from a master command-and-control server. And apparently this trick disguises, well, naturally this trick disguises the communication as encrypted connections within a trusted service. So it's like, hey, well, we trust Google. And we have an SSL secure encrypted communication using Google Drive to Google Docs. And so not surprisingly, a Google representative was quoted saying that the company is investigating because "using any Google product to conduct this kind of activity is a violation of our product policies."

TOM: So this is isn't infecting somebody who's using Google Drive. It's using Google Drive to mask how it's sending command-and-control?

Steve: Well, yeah. The idea is, of course, Google Docs is in the cloud. And so you would set up a Google Docs account and give that account public access. And then anybody who wants to can see that. And so the idea would be that the trojan itself would incorporate a copy of Google Drive, or maybe download it once it infects a new computer. Essentially it would set up a local instance of Google Drive, mapping the drive to that Google Docs folder.

TOM: Gotcha.

Steve: And so all the trojans would be watching that folder. And then the trojan master, the bot master, would then post instructions about what they wanted the trojan fleet to do to that folder. All of the connections, all of the trojans monitoring the folder would see the update, get the instructions, and then go about doing whatever they've been told to do.

TOM: I know Google probably doesn't feel this way, but it's actually a fairly good recommendation for using Google Docs as a convenient way to share information.

Steve: Well, and of course traditionally we were using IRC Chat. But IRC Chat didn't have the reputation that Google has and didn't just come with free SSL connectivity in order to protect and make secret and thus non-packet-sniffable the transactions going on. And so here you just bootstrap yourself and use all of this nice, well-working, high-availability Google infrastructure in order to communicate to all your trojans. How convenient.

TOM: How very convenient. Finally, this is somewhat upsetting, especially to fans of FreeBSD. They had a pretty severe server compromise a couple months back.

Steve: Yup. Well, and they didn't discover it until nearly two months later. What they found on November 11th, so that's 10 days ago, was that back on September 19th, so nearly two months ago, actually a little more than two months ago, but after nearly two months of the breach, back on September 19th, two of their servers had been breached. What the security team believes is that the intruders gained access to the servers using an SSH authentication key that was stolen from a developer.

So this is a problem. You've got secure access to servers, but not just one person has that. Your whole development team needs to be able to post to that central repository. So every single one of those developers has highly privileged access to the server; and, if they lose their credentials that authenticate them, then this sort of thing can happen.

So the consequence is, and this is what FreeBSD has formally stated, is that users who've installed FreeBSD since September 19th are being advised to completely reinstall their systems because third-party application software packages - and there's been no enumeration of them. Maybe we'll get that in the future. From what I've heard, just as of this podcast time, they're still looking into this to figure out exactly what happened. But they know it's not the kernel, not the system libraries, the compiler, or the command-line tools. So it's just the FreeBSD application repository. They think that may be where some mischief occurred.

So they're saying, if you did a recent, since September 19th, install of FreeBSD - and normally the FreeBSD install is a network install that sucks everything down from the repository on the fly - that you are advised, and they apologize, of course, but you're advised to do it again. So it happens that FreeBSD is my UNIX of choice. It's the one I have gotten to know and like a lot, so...

TOM: It's a good one. Am I remember this correctly? That's the one that Darwin is based on for OS X?

Steve: Yes. It was FreeBSD. And actually it was Brett Glass who recommended it to me a long time ago. And that was a good recommendation.

TOM: Well, what's going on in the Twitterverse? You got any good twitters to talk about this week?

Steve: Oh, just - I've mentioned this before, and I keep having people thank me for this recommendation, so I just thought I would share the fact that people are liking this. And this was Mark Grennan, who tweeted, he says, "I'm a MySQL database administrator, and I love the outdoors," he said. Oh, actually, that was in his profile on Twitter. So he said, "@SGgrc Just a thanks for the tip on the show 'Homeland,' the most well-written dramatic show I've ever seen. Thanks."

TOM: Man, she can be hard to watch sometimes. She's so freaking nervous all the time.

Steve: She's a little twitchy.

TOM: She's so good at that.

Steve: Yes. And but she's acting. Claire Danes is a serious actress, and...

TOM: No, Claire Danes is not like that in normal life. I've seen interviews of her.

Steve: Yes. Yeah. She is doing a great job. It's in its second season now. And you really need to, if you're going to get onto it - and I keep hearing people talk about it. I mean, it is wonderful. It's on Showtime Network. And you can get, like, the first season. You can - I'm sure the various archives and repositories, I don't know where Showtime officially has theirs, probably just on the 'Net you could watch it from them.

TOM: I think you can get it from iTunes. And I hope I'm not misremembering this, but I thought Netflix may have had them, or at least did have them for a while. Stuff sort of comes and goes there.

Steve: Yeah. Anyway, I do recommend it. I'm enjoying the second season as much as the first, which was terrific. And I have two quick little notes about SpinRite. Greg Kurts, he also tweeted because I have "@gkurts," he said, in Cleveland. He described himself as a "software developer who works from home, family tech support guy, dad, husband, and giver of treats to dogs."

TOM: Good man.

Steve: And he said, "@SGgrc This is twice that SpinRite has helped me salvage a drive that was 'gone.' I can't brag about it enough. Money well spent." So thank you, Greg, for publicly tweeting that.

And then Ray McGill, who's a listener in New Hampshire, said, "I want to buy SpinRite for USB drives. A long-time SN listener. I want to use" - okay, now, he said SN, but he meant SP or, oh, no, SR, SpinRite. "I want to use SpinRite to scan two USB external drives. The only answer I've seen is that there's no support for USB, and some people have gotten it to work through magic. Can I use SpinRite for USB?" So I just thought I would answer Ray's question, and to anybody else who's wondering, the answer is yes.

TOM: Do you need to be a wizard, though?

Steve: You do not need to be.

TOM: Oh, good, all right.

Steve: But you need to have a more recent motherboard. But "more recent," of course, is relative. What you need is a motherboard that understands USB. And all motherboards today do. They let you boot from USB. They see USB. You've got USB drives in boot options and so forth. So as long as the motherboard that you're running it on does recognize a USB drive plugged into it - and that's easy to do. Plug the USB drive in, reboot, go into the BIOS, and see if it's listed there. If the motherboard sees it, SpinRite will see it, and then it will run on your external USB drive with no problems at all.

TOM: I'd forgotten about having to install drivers to be able to see USB drives. But the very first time I ever used a thumb drive, Robert Heron at COMDEX, I want to say 2002, 2003...

Steve: Yeah, probably 10 years ago.

TOM: ...one of the last COMDEXes, brought his stories for Fresh Gear to me to post on the web. And he's like, hey, you've got to try out this cool thing. It's a thumb drive. And I had to install the drivers in Windows '98 to be able to read the files.

Steve: Yup.

TOM: But you don't have to do that anymore.

Steve: Nope. As long as the BIOS knows about it, SpinRite will be able to see it.

TOM: Excellent. Well, let's move on into our Listener Feedback #155, shall we?

Steve: Wow. Yup.

TOM: All right. Paul White in Portland kicks this off with some good news about Verizon. He says: Longtime Security Now! listener, I have heard/watched them all, and a SpinRite user, blah blah blah. He wrote "blah blah blah." I just wanted to pass this along to you. Changed my email address for my online My Verizon web account and got this message upon completing validating the new email: "Your new email address has been validated. Once the new email address has been in our records for 30 days, you can use it to reset your password." Am I wrong, or is a major player doing something right? Thanks to you and Leo. Keep up the fine work. Paul.

Steve: Yeah. He is exactly right. This is the first I had heard of that. I'm also a My Verizon account user.

TOM: As am I.

Steve: And it's like, wow, that's really good. Because of course this prevents a number of potential exploits where somebody might somehow get onto your web account. For example, like if there was a - I don't know that Verizon doesn't use SSL, but we were just talking about the problems that you have of having your session sniffed and impersonating someone. So you might be impersonated, but what you want to prevent is them being able to then leverage that into getting your password. Or not getting it, but being able to reset it, claiming to be you and to have forgotten it. So putting up a 30-day waiting period, essentially, that represents a really nice step forward in security.

TOM: I think Mat Honan would have liked this sort of policy to have been in place.

Steve: Precisely. Precisely.

TOM: Second one comes from Daniel in Colorado, posing a question regarding cascading ciphers. He says: Hoyo, Steve and Leo. TrueCrypt offers the ability to use more than one type of encryption when building a file container or when encrypting a hard drive. Say we are using a paired AES-Twofish configuration. TrueCrypt will encrypt each block with Twofish and then encrypt that intermediate result with AES. Each cipher gets its own 256-bit key, which are mutually independent. I am wondering if this merely doubles the encryption, like adding a second lock to a door, or if it somehow creates a stronger encryption by reencrypting the first cipher text with a second, different key and algorithm.

Steve: Great question. Okay. So the idea is, the reason this multi, what did he call it, paired or cascade, he used the word "cascading ciphers," the reason that's there is, I mean, this is really belt-and-suspenders stuff. This is the TrueCrypt guys saying there's a non-zero probability, but probably vanishingly small, but still maybe not absolutely zero, that a flaw could be found in a cipher. And if a flaw were found in the cipher that you had chosen to use for encrypting your volume or whatever in TrueCrypt, then that would be a problem. So the reason they allow you to use multiple ciphers is that the chance of all of

those multiple ciphers being simultaneously found wanting, I mean, that probably really is about as close to zero as you need to get.

So it actually isn't for greater security, like from the standpoint of making it harder to crack your drive, although we'll address that in one second. The main reason was just, again, belt and suspenders. If you encrypt under one cipher and then reencrypt the encrypted results from the first cipher with a second different cipher, then you get the encryption goodness of both. So that, if one were completely destroyed cryptographically, like overnight, you're fine.

TOM: So it is like two locks on a door in that sense because you have to - even if you break through the first lock, there's another lock you'd still have to break through?

Steve: Yes. Maybe it's like two doors in succession.

TOM: Okay, yeah, sure. I got the first door. Oh, crap, there's another door.

Steve: Like "Get Smart" when he was trying to get into his base. So it doesn't double the encryption because, now, if we switch from why they did this, his question is, does it double it, or is it stronger? And the fact is it is stronger in the sense that you have two 256-bit keys, both of which you would have to brute-force in order to crack each of the ciphers in order to decrypt the drive. Now, I would argue that one 256-bit key is already all the suspenders anyone needs. And so, yes, this could give you the equivalent of a 512-bit key because they are independent of each other, and you'd need to brute-force both in order to decrypt the drive.

Now, it's not double the strength because that would be a 257-bit key, rather than a 256-bit key. You'd be adding one bit to the key. That doubles the strength. This is 2^256 times stronger. So it's just ridiculous. Which is why it's like, okay, TrueCrypt does a good job at generating a very high-quality, extremely pseudorandom 256-bit key. The danger in TrueCrypt is that you don't choose a strong enough password; that as far as we know, the only known problem comes from the user using a weak password and then somebody brute-force guessing passwords, looking at the header on the TrueCrypt volume and decrypting it. So the double cipher doesn't protect you from the actual probable means of attacking a TrueCrypt volume. It protects you from either of the ciphers being found to be insecure suddenly through some breakthrough that we're not expecting to have happen.

TOM: Well, this is similar to what we were talking about with quantum key distribution before the show, which is, yeah, it's more secure, but it doesn't actually address the likely security vulnerability.

Steve: Well, and we've talked about this often. The perfect model is the chain, with security being the strength of the chain, which depends upon the strength of every link in the chain. So the weakest link is the problem. That determines how strong the entire chain is. And so you may have one quantum crypto key exchange link that just, I mean, is made out of titanium. But if the other one is a plastic straw, that's going to be the one that gives. So you still have a problem with your chain.

TOM: We've encased the titanium link in another case of titanium. We still have a plastic straw link, though, so...

Steve: Yeah.

TOM: Steven Metz in Highland, New York responds about secure nuclear site data monitoring. He says: Hi, Steve and Leo, or in this case Tom. Did you write that in, or did he write that?

Steve: No, no, I did because I kept seeing Leo referred to. I thought we've got to give Tom some. Now, this is a long post, but I want you to share it with our listeners, Tom, because it is cool. And we'll talk about it afterwards. But it's just nice to know that somebody's thinking about this, and there actually is a solution.

TOM: All right. He says: I love the show and listen to it on my daily commute to work. I was just listening to Episode 375 and wanted you and the listeners to know that there is a solution to the issue of securely monitoring nuclear power plants remotely. I'm referring to this portion of the episode where Leo said, "Do not manage nuclear plants remotely, please." Steve, you said yeah. Leo said, "If you don't want to be onsite, then don't do the plant at all, just don't. If you don't feel safe enough to sit next to the core, don't build it in the first place." Steve, you said, "Yeah, so what's done is pure convenience and absolute sheer stupidity." Leo said, "I have some files which I would not wish anyone to tamper with. I store them on a computer never connected to the 'Net, never had been. If I can take such precautions, why is it that infrastructure computers are connected to the Internet? For the sake of convenience? As an example, nuclear power stations have been around since the '60s, and they weren't connected to the Internet then. They worked. Why connect them now?"

So here's what Steven Metz says. Disclaimer: I work at a company that sells the solution I'm speaking of, Owl Computing Technologies. As you said, it would be convenient to be able to monitor the status of our infrastructure with pieces of it geographically distributed from a central place in real time. Also, many facilities have databases that archive real-time data from all sensors and devices. Many industries have regulatory requirements for such databases and must replicate them offsite. So the issue is, can data be exported in real time without providing a path that hackers can use to probe and attack the source of that data? As your listeners know, firewalls and antivirus programs are no guarantee for home PCs, much less nuclear facilities.

The company I work for has solutions in place at more than half of the nuclear facilities in the United States. These solutions consist of a pair of machines. One is connected to a secure internal network or device, while the other is on a separate, less secure network, perhaps with Internet connectivity. A pair of fiber optic cards connects these two servers, but one card only has the optical hardware to send data through the fiber, while the other only has hardware to receive data from the fiber. Thus it is physically impossible for a signal to be sent or received in the opposite direction. We call this a "DualDiode" configuration.

If you wish to know more details, we refer to the machines as the "send-only" and "receive-only" servers. These names are based on how information is sent relative to the fiber optic connection between them. The send-only server is on an internal protected network, while the receive-only server is on a less secure network, perhaps with Internet connectivity. There is a protocol break between the servers, meaning that you cannot specify the final IP destination address for data from the send-only side. A computer on the protected network can only specify the IP address and a port number on the send-only server and has no knowledge where the data goes after the DualDiode.

If we're dealing with files, then the whole file is received on the send-only server. The file is then sent over the one-way fiber optic link through a protocol similar to ATM. No TCP header information is retained or passed to the receive-only server. Instead, the send-only server assigns the file a channel ID based on the port number the file came in on.

This channel ID goes with the file to the receive-only server and is then mapped to a destination IP address/port. This is like a TNO double-blind setup that you speak of, Steve, in that the send-only side has no knowledge of the destination IP or port number, and the receive-only side has no knowledge of the source computer IP address.

Also, files sent through the DualDiode can be scanned for viruses and/or filtered on file type, whitelist and/or blacklist keywords, XML schemas, or any other criteria. With the protocol break in scanning, even if malware were already inside a customer network, it could not dial home through this one-way outgoing connection. Similarly, streams of data can be sent. They are subject to the same protocol break, where all TCP/UDP header information is stripped. The streams can also be scanned in real time to verify conformance to different formats - MPEG-2, RTP, et cetera.

With this DualDiode type of solution, a company can share information from one network to another, while maintaining network separation. And the DualDiode can be installed in the opposite direction to only allow data to go into a private network. This is common in government use cases such as feeding live surveillance video streams into top-secret networks with assurance that no data can leak out of the top secret network through that connection. If more information is needed, feel free to visit our website, and he gives the website owlcti.com. Thanks, Steven Metz.

Steve: Isn't that neat?

TOM: Does that work?

Steve: Yeah. So the idea is you create a hardware-enforced one-way connection. You need a protocol which inherently doesn't require acknowledgment. So we're used to, for example, TCP or even UDP protocols where the consequences of our sending is acknowledged in order to verify that the data was received. But that's not secure. There's no way to do that safely. So what they do is they use a fiber optic link, and there's a transmitting laser diode at one end, but no receiver. And at the other end is a receiving photo diode and no transmitter. So it's hardware-enforced that the sender is just blindly sending data out, and the receiver is receiving it. And the trick is it has to be a protocol that will work without acknowledgement because there can be no acknowledgement of the data being received.

TOM: It's all SYN, no ACK?

Steve: Yes [laughing]. Exactly. And so it's going out, just assuming that the other end is connected and listening. And in that way, as he says, I guess they called it a DualDiode. I was thinking, when I first read this, I thought, well, why just not one diode? And I guess it's that they visually see it as one at each end. There's a diode, the electrical component diode, only allows currently to flow through it in one direction. And so each card, each interface card, is like that. One is a transmitter only; one is a receiver only.

And so their whole company, at least this aspect of what they're offering, is built around one-way connections. And then they've developed a protocol, a carrier protocol which doesn't require acknowledgment. So essentially the emitter, the sender, is sort of like it's broadcasting blindly into this fiber optic; and the receiver captures it and figures out what's been sent. Maybe there's error correction stuff so that...

TOM: Yeah, I was going to say, what do you do about packet loss or stuff like that? It all has to happen on one end.

Steve: Right. So they've added the technology to make that work robustly. And I thought it was just nice that such a thing exists. That's great.

TOM: Yeah, very interesting. Thanks, Steven, for that. Question #4, Gary McCleery in Oamaru, New Zealand - kia ora, Gary - says rubbish, rubbish, rubbish. Uh-oh. Hi, Steve. I just don't understand why you're so excited about Microsoft turning DNT on. That's Do Not Track. It's simply a request not to be tracked. It doesn't stop these people tracking us. It's not as if, when we set this setting, people can't track us. They make too much money from tracking us for them to stop. I admit it is an extremely small step in the right direction, but it is meaningless in respect to any actual benefit to the user. So please, please explain why you are so excited about what Microsoft has done. Cheers, Gary McCleery.

Steve: So I picked this out of the mailbag because it's something that a number of users have been confused by. And I said it before, but this is a perfect time to just say, okay. I get that. I mean, I'm Mr. Turn-Off-Your-Third-Party-Cookies, which arguably is a proactive step that a user can use to configure their browser to actually change the behavior of the browser. And I get it. In fact, Microsoft, as we discussed last week, just calls this a "signal," the Do Not Track signal. So it's just a signal that the browser is sending out expressing user intent.

And so many of these things take time. The keynote that we shared from the Microsoft executive last week understood that we're at the beginning of a discussion; that the world has changed. The issue of privacy and secrecy has changed as a function of the technology which we've generated, and this is creating new problems. And one of the things that Microsoft recognized is that users want control. They don't necessarily need the secrecy of their data because, for example, they're posting all kinds of stuff on their Facebook pages. But they want control over who has access to it.

And so I recognize we've got a long way to go. I absolutely get it that this doesn't change anything. But this represents a step forward. And this is the way it's going to happen is that this will happen; users will understand it. Microsoft, reading between the lines, it has to be that Microsoft sees a marketing advantage for them to have this signal turned on by default. Which I hope works. Where Microsoft, you know, the idea would be that IE would get the reputation of requesting not to be tracked more than the competing browsers. And the case is, it's easy to turn this on in Firefox. It's almost impossible to turn it on, I mean, you can, you can find it in Chrome, but you've got to dig.

So Microsoft surveyed their users. Three quarters of them said I would rather not be tracked. So Microsoft set the default that way and explains what they've done when you install IE10, which either comes with Windows 8, or you're able to install it on Windows 7. So, yeah, Gary and everyone else who, like, thinks I'm crazy, just be patient. I'm patient. There's no - it's just going to take time. And ultimately I'm sure, if the industry does not regulate itself, it will be regulated. And that's the tension that exists here. There is tension between what users want and what the tracking companies want. And if users' desires not to be tracked are not voluntarily honored, then they will be honored by law, and companies that break the law will suffer the consequences in reputation and treasure. So we're just at the beginning. This is not the end. This is just the beginning.

TOM: It's always easy to criticize a voluntary system like this; to say, well, people don't have to follow it. And especially in the early times, people aren't following it. The idea of any kind of system like this is that, if you get enough big companies to follow it, most of the reputable companies will follow it. And if only the reputable companies are following it, it's the disreputable companies that people will avoid anyway because they won't want to use that stuff.

Now, things are a little different on the Internet because it's so easy to go from one place to another. But I think about the rating system for movies. There's no law that says you are restricted to being younger than 17 to see movies. The movie studios and the movie theaters agreed to self-regulate, and it's as if it were a law. And you could say, well, the theaters don't have to follow that, and they don't. But that is an example of it working. Now, it doesn't always work. You can find examples of this sort of thing not working, and that's where it gets to what you were saying, Steve, which is the government will often step in in those cases if the voluntary regulation doesn't work.

Steve: Yeah. I think that the other thing, also, is there's a real question, which has never been answered, about what's the true value of this tracking. My feeling is they track because they can. And this is where Leo and I have gone around and around because it's not at all clear that anything bad will happen if this data aggregation stopped. It's not necessary for actually Google, now that they own a major advertising supplier, it's not clear that it's necessary for them to aggregate all of the locations that I go to and try to figure out who I am. People feel that it's a little creepy if they are, like, shopping for some things on one website and then go somewhere else and ads sort of follow them around the Internet. It's like, oh, that's a little creepy for them. So when people see evidence of tracking, it's not something that they think is inuring to their benefit.

So it's not at all clear that, if all of this aggregation was just outlawed, that anything would change. I think the advertising industry would go right along. Ads would be on websites. Everything would be fine. I believe they track and aggregate only because they can. And if they couldn't, if this had never happened, we would still have a rich, ad-based ecosystem. And I'll bet that's where we're headed.

TOM: Yeah. And in fact the increase in tracking has not made Internet advertising more valuable than television advertising yet, and television advertising is notoriously badly tracked, as far as who's watching it. So it's not as simple as what you've got out there. There's a lot of complex factors, a lot of moving parts.

Steve: It is funny, too, because you know that dollars spent on TV are being spent as wisely as possible. So when you're looking at the commercials of a given show or a given time, you can tell, like, oh, this is the market segment that these guys are going after because of the nature of the commercials that occur on that channel at that time or during that show. So you can sort of see the effect of that in aggregate.

TOM: And there's an argument to be made, not to carry this on too long, but there's an argument to be made that you can persuade people to want to be tracked because, if you're watching a video, and it's got a bunch of advertisements for men's cologne, and you're a woman, you'd rather, you know what, I don't want to see that. That's not relevant to me. I want to see ads that are relevant to me. So I'm not saying that everyone will go for that, but there's an element of that in Do Not Track, too, to say I would like to be tracked. As long as you're the one in control of that, I'm fine with that.

Steve: I think one of the things we may get to, and it will be interesting to see how this experiment plays out, is the notion of opt-in tracking, where you go to a site with a browser that is saying Do Not Track. The site you go to sees that, and they get enough more revenue if they allow their advertisers to track us who visit there than if they don't, that they say, hey, we're ad-supported. We need you to enable tracking of you while you're here because it makes that big a difference to our revenue. And that's the experiment I'd like to see.

TOM: Yeah.

Steve: What will people do? Will they go, hey, no, it's not worth it to me to be tracked, that's creepy, I'm going to go somewhere else where I'm not being asked? But that kind of informed consent is - that would be a nice thing to see.

TOM: Yeah. Question #5 comes from Tali Sherman-Hall in Seattle, Washington, wanting to know if we can help her get the word out. Tali writes: I recently started getting unsolicited and unexpected email from Digital Lifeboat, an online PC backup service. It appears that somebody set up an account using my email address by mistake, so Digital Lifeboat have been sending me updates about the PC being backed up on the account, with links to their website. If I wanted to access this person's account, with all of their personal information and all of the data in their backup files, I simply have to go to the Digital Lifeboat website and request a password reset using my email address. It is just the customer's dumb luck that I am too scrupulous to do so. Even if I don't do anything, though, they are still not getting any notification emails - I am - and they will never be able to reset their own password because everything goes to my email address.

I tried unsuccessfully to contact Digital Lifeboat by email and their webpage so I could warn them about the risks they are taking by not verifying customers' email addresses during account creation, and possibly to find a way to contact the customer without invading their privacy. But none of that worked. So today I finally posted on their Facebook page and did receive a response from a man who initially did not understand the issue and then proceeded to tell me about an even bigger problem.

He wrote, quote, "I forgot that our two-factor cell phone authentication is currently disabled. When we disabled it, we left this gaping hole in our security. Wow, that's our screw-up. Thanks for bringing it to our attention. I'd like to delete this thread so as not to advertise this issue, and we'll address it internally." That thread was deleted, but not before I took a screenshot of it.

It worries me to think of all the people trusting this company with their personal information and backup data because they're not being told the truth. Digital Lifeboat touts "256-bit AES encryption to insure complete privacy and security of your data," but what good is that encryption when they don't bother to protect the customer's account? It's the same thing we were talking about earlier. They clearly have awful security practices, and I'm still not sure they even understand the initial problem I brought to their attention. At least now all Security Now! listeners will know what's going on with this worrisome PC backup solution company. Thanks for allowing me to vent. So that's Tali Sherman-Hall saying this about Digital Lifeboat.

Steve: Isn't that something? So if Tali is right, and from everything she said it sure sounds like it is, this particular backup solution company doesn't have an email verification loop. So when you set up an account with them, a typo in your email address goes unnoticed by the person setting up the account or them. And of course, as we know, email is sort of the one common denominator that all Internet users have. And we're assumed that we have exclusive access to our own email account. So that's used as sort of a universal password recovery, prove-that-you-are-who-you-say-you-are solution.

So in this instance, what Tali is having is the person who set up the account had a typo. Who knows what his email address is that could have been wrong. So that the company believes Tali is this other guy. And thanks to the fact that email control is so important, she has control over this person's, unwitting and undesired control over this person's backup solution. So anyway, I thought this was a great little case study in the importance of setting things up right when you do PC backup.

TOM: 256-bit AES encryption doesn't mean anything in that case.

Steve: Exactly. The weakest link again.

TOM: It's that plastic straw that got you. It does sound like they actually had a two-factor authentication system, and it sounds like they at least have the idea of what should be done. They just had it off.

Steve: Well, and I'm not sure actually...

TOM: Didn't have any kind of backup.

Steve: Yeah, I'm not sure how their cell phone authentication being disabled factors into this because I don't know why that would have...

TOM: Depends on how they implemented it, yeah.

Steve: Yeah, why that would have fixed a lack of email verification. But certainly here's a good reason why email verification is really important.

TOM: Absolutely. You know, Nintendo Wii U, couple of good things about that yesterday. One, when they asked if I wanted to be opted into a mailing list, it was not an opt-in or an opt-out. It said, "We have a mailing list. Opt-in, opt-out." Forced opt situation. Ideal for everybody.

Steve: Nice.

TOM: And the second thing is they did email verification when I set up a new account. They said we're not going to do anything with your Wii account until you go to your email, and we make sure that that's really you. So, good...

Steve: Yeah, and then you click a link that verifies that you have control of that, yup.

TOM: Question #6 from Vincent in Pittsburgh updates us about Hamachi's private network IP allocation. He writes: Steve, I thought this was interesting. Hamachi must have retired their use of the 5.0.0.0/8 IP allocation. I recently deployed a new Hamachi installation, and it grabbed an IP from 25.0.0.0/8, which is IP space registered to Great Britain's Ministry of Defence. I guess they assume most users would have no legitimate reason for accessing this IP space. Thanks.

Steve: Okay. So this bears on a number of things from different directions that we've talked about in the past. First of all, we discovered, our podcast listeners discovered Hamachi in its infancy, and people just went wild. And in fact it wasn't very well known at the time, and this podcast is probably responsible for putting it on the map. We've had a great dialogue with the original designer of Hamachi. And so it's been a resource that a lot of our listeners have used.

Now, the problem was that Hamachi cleverly - the way Hamachi works is, when you installed the client in different machines, it was using the 5-dot network which was a chunk of IP space, Internet allocation space that had never been allocated. It was dark. But it wasn't dead. It was just not used yet. So we have private networks. Everyone is used to, who uses a little home router, will see 192.168.something.something, typically .0.something or 1.something. So the 192.168 is formally reserved for private networking. Then the larger one, of course, the biggest one, is the 10-dot network, which

is formally reserved for private networking. Five wasn't reserved; it just wasn't used.

But as the Internet, as we've discussed recently, famously began to run out of IPv4 space, the 5-dot network got allocated. So the idea was that Hamachi was very cleverly, five years ago or more, using the 5-dot network because it was, A, not in use publicly, meaning that there was never an IP address of 5-dot anything out in the world. Nor was it a private network, that is, it wasn't 192.168 or 10-dot anything. So the idea was it could safely coexist in your system. It wouldn't be public. There would never be a chance of it colliding with a public IP. But neither was it private. It couldn't collide with a 10-dot IP or a 192.168. It was, like, just completely off on its own. Until we began to run out of IP space, at which time the central authority said, okay, we're going to start chopping up the No. 5 space.

Well, that immediately raised a question, uh-oh, what's Hamachi going to do, because they couldn't keep using five any longer. Five-dot addresses started to appear out on the Internet, and that would be a problem. So then, in another related issue, there's this bizarre 25-dot space which we have talked about on this podcast in the past because this entire huge chunk of space, this is 16 million IPs, every IP beginning with 25. So from 25.0.0.0 to 25.255.255.25, technically 254, that entire block, huge block, was originally given, back when, oh, my god, we'll never run out of IPs, back in the old days to Great Britain's Ministry of Defence.

And get this. It just uses it like a private network. That is, it isn't routed publicly. It uses it just like a 10-dot network for its own internal communications. Meaning that in the same way that 5-dot used to be non-publicly routed because it had never been allocated, 25 is not publicly routed because the Great Britain Ministry of Defence owns it and doesn't have machines on the public Internet. They just use it for talking to each other privately. But it's not a 10-dot, which is really what they should be using.

So anyway, this is clever. This is Hamachi figuring out, or now it's LogMeIn because they're the people who purchased Hamachi, they figured out another Class A block of IP space that wasn't publicly routed, wasn't subject to being allocated because apparently the Ministry of Defence is not in any hurry to give up their 25-dot allocation of 16 million IPs, and it's sort of an interesting compromise. Now, apparently this switchover did cause people problems. People who were using Hamachi who weren't dynamically allocating the Hamachi IP, suddenly all of the 5s turned into 25s, and so people have scrambled around because their Hamachi links broke.

TOM: And they had to go manually change them?

Steve: Right. Also Hamachi does support already IPv6, and Hamachi has their own IPv6 allocation that can never be taken away from them. So all of this 5-to-25 switchover is just for the sake of maintaining compatibility with IPv4 addresses for those people who need that. But right now OSes have IPv6 stacks and have had those protocol stacks for a number of years. So you probably can just go ahead and switch over to using IPv6 and disable IPv4 completely. But this was an interesting hack. And as long as the Ministry of Defence holds onto their 25-dot allocation and does not route it publicly, I don't see any problem.

TOM: Does this mean, though, that you can't use Hamachi if you're inside the Ministry of Defence's network?

Steve: That's correct. And in fact, in their blog posting, LogMeIn said we do not have a single customer of Hamachi in Great Britain's Ministry of Defence. That would cause a problem. But not only isn't it a problem now, but as you said, that would cause a

problem; and so nobody in Great Britain's Ministry of Defence will be able to use Hamachi and a 25-dot address because it would foul up their network.

TOM: This makes me wonder if - CNET was, like, notoriously against allowing LogMeIn on their internal network. I had to get special permission. The IT guy had to whitelist the exact Ethernet port I was going to use when I went live and wanted to show LogMeIn. And I wonder if they weren't using 5-dot internally for something, now that I've been thinking about it. I'd never thought about that before. But it sounds like we need to go from MI6 to MIPv6.

Steve: There we go.

TOM: Question #7, Robin's kids in the U.K. must have Java. Why is that? My kid's laptop has just spectacularly failed MS Offline Defender with a dozen or more Java exploits. Java? Why is he allowing them Java, I hear you cry? Well, Minecraft. If they didn't have Minecraft, then life would be unbearable (for all of us). I have allowed MS Offline Defender to clean the PC, and I'm just rescanning it. Last time I restored a clean image, but there seems little point in doing this again as it will only get infected again. So my questions are, should I trust it if it comes up clean? What can I do to reduce the chances of it happening again? And should I remove Java and buy iPads for all three of my kids in the hope that this is immune to infection? And there's a Minecraft app for iPad. He says: Thanks for the invaluable resource and a tip of the hat to Leo for his impressive knowledge of the world and culture outside the USA, a major part of the reason that Security Now! is an international giant.

Steve: Okay. So if you have to have Java, and, for example, your kids need it for, like, one purpose, to run Minecraft, I did see in the mailbag this week somebody who was promoting the idea of using Sandboxie to sandbox Java. And the problem is I don't think that's probably secure enough. Sandboxie essentially filters the API, the application programming interface calls between applications and the operating system to prevent them from writing files to the OS and so forth. But the nature of Java exploits allows the user's own code to run in the app. And I just - to me it feels like maybe Sandboxie isn't enough isolation to really give you security.

A great solution for Robin's kids and the Robin household would be to use a true virtual machine. Free virtual machine technology exists now, so that won't cost anything. And the idea would be all you have to do is train the kids to launch the VM; and, in that, nothing is installed except Java. I mean, you could put other things in there if you wanted to, but you probably don't want to have too huge a virtual hard drive established. So just put Java in the VM. Then the external system won't have it. You won't have to worry about it getting infected. And then you just have a contained environment, a virtual machine that shouldn't cost anything, that's got Java installed, and that's where you play Minecraft. And I think that's probably a robust enough solution.

TOM: Now, Chad Johnson, who hosts our Minecraft show, or his own Minecraft show, it's not even ours, and obviously was interested in the answer to this question, came running over and says, "Well, if I just disable Java in the browsers, is that not enough?"

Steve: I really think it probably is. And that's what I would - the way I've got - there are a couple things that I'm using Java for. And I have my browsers set up so that they will not run Java by default. So, yes, just disabling it in the browser, and we've discussed exactly this point before, I thought I would explore a little more of a containment solution. But, yeah, as long as your browser won't launch it. That's where the exploits are is from going to a page that invokes Java and takes advantage of unpatched problems. And clearly, if it had a dozen or more exploits in the case of Robin's kid's

laptop, then that had happened in the past.

TOM: So Box.net, Ubuntu, or Mint, and Minecraft for Linux, you're golden. All free.

Steve: Yeah, nice.

TOM: Nice. All right. Question #8, and leading us apparently into next week's cool topic, we have Craig Lewis in Wales of the United Kingdom asking: DTLS? What is it? Hi, Steve. I've been a long-time listener, from the very start. I've emailed before a few times. I've just run Windows Update, even though it isn't the second Tuesday yet, and there's an optional update. It adds support for something called "Datagram Transport Layer Security" in Windows 7. What exactly is it? Something to do with SSL? I wonder if this is something new and worth adding to Windows? Microsoft, as always, are very vague in their help and describing things. Maybe it warrants a topic? I have no idea. Thanks for your time, and all the best. Thanks to you and Leo for years of what is a great and well-valued show. Craig in Wales. P.S.: SpinRite is pure genius; and, yes, I have purchased it.

Steve: Well, that's a great question. And I think it's a great topic for a show. So next week we will talk about DTLS, which essentially is transport layer security for UDP datagrams rather than TCP connections. And it was designed specifically to be a very secure but more lightweight solution. It looks like it's going to take hold. So that'll be what we talk about next week.

TOM: What's an example of somebody who would use this?

Steve: Well, for example, DNS is famously very lightweight, but very insecure, inasmuch as, even though you may have SSL connections for your data, your browser's queries to a DNS server are in the clear. There is no encryption for DNS. And of course we've talked about DNSCurve, which adds encryption to DNS. OpenDNS has a secure DNS. But using something like Datagram Transport Layer Security, DTLS, would sort of give you the best of both worlds. You'd be able to use a low-level, lightweight standard where you don't have to establish a connection first, yet still get privacy protection of the data passing in both directions, and the other advantages that SSL brings, like authentication.

TOM: Right. Best of both worlds. I can't wait to hear more about it next week.

Steve: Next week.

TOM: Leo will be back then. Steve, it's been great doing these shows with you the past three weeks. Thanks for letting me fill in.

Steve: Has been. And it's been a pleasure for me, too. I'm glad we were able to continue cruising forward, even when His Nibs is off...

TOM: Cruising the ocean.

Steve: ...cruising in Australia.

TOM: Yeah, exactly. GRC.com, of course, the place to go. Any new things people can look for? Anything you want to point them to in particular this week?

Steve: Just our regular, you know, remember GRC.com/feedback allows you to send stuff like we were just reading. And otherwise I've got some stuff in the pipeline that I'll be talking about.

TOM: All right. We'll keep an ear out, keep an eye out, and keep listening. Security Now! of course on TWiT.tv/sn. It's live on Wednesdays. 11:00 a.m. Pacific is when they start to settle down in the chairs and talk some security. Like I said, Leo Laporte will be back next week. I'm Tom Merritt. See you later, folks.

Steve: Thanks, Tom.