The Supply Chain

Description: This week we examine and explore an October Windows Surprise of a different sort. A security researcher massively weaponizes the existing MicroTik vulnerability and releases it as a proof of concept. Israel's National Cybersecurity Authority warns about a clever voicemail WhatsApp OTP bypass. What DID happen with that recent Google+ breach? Google tightens up its Chrome Extensions security policies. WiFi radio protocol designations finally switch to simple version numbering. Intel unwraps its 9th-generation Core processors. We've got head-spinning PDF updates from Adobe and Foxit. This isn't a competition, guys! And, finally, we take a look at the danger of Supply Chain Attacks, with a possible real-world example.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-684.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-684-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Great show planned for you. We're going talk about how hackers use your voicemail box to get into WhatsApp accounts. We'll talk about the problem with Windows 1809. Microsoft thinks it knows what went wrong. And, yes, the big debate over the Bloomberg Business Week story about Supermicro. Apple denies it. Bloomberg stands by it. Who's right? Steve weighs in. It's next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 684, recorded Tuesday, October 9th, 2018: The Supply Chain.

It's time for Security Now!, the show where we cover the latest security news and ways to protect yourself. It actually is all about understanding technology, too. And it's all thanks to this guy, who understands it better than anybody, Steve Gibson from GRC. Hi, Steve.

Steve Gibson: Hey, Leo. Great to be with you again, as always. So you asked if this was the SQRL episode, and I think what we'll do is probably have me up there live for that.

Leo: Good, good.

Steve: Because what I really want to do is I think get a couple of your gang together and do it sort of as an adversarial mode, where I explain what it is, and then you say, "But what about this?" "But what about this?"

Leo: But what about that, yeah, yeah. Because that's what people say when you talk about it.

Steve: Exactly. You and Mike Elgan and Jason probably would be a great board to say, okay, what about when this happens, and what about this. Because I think there's an answer for every possible "but what about" question. So anyway, I'm in the process of bringing up a third portal, essentially. I had to develop an API, a generic API, which I call the SQRL Service Provider API, because I want to bring it up for the SQRL forums, which are written or developed, well, they are XenForo forums. And while you can sort of add things to it, one of our listeners is a XenForo developer.

And so I thought, boy, you know, rather than having to learn PHP and then figure out all about what XenForo was doing, if I could give him a clean API to a SQRL Service Provider backend, then he could just write to that; and, of course, so could everybody else. So it's like it was a piece that I didn't appreciate having would really improve the adoption ease and rate.

So that's done. I've written that. I defined it and wrote it. And as far as I know it's finished, but now I need to test it. So I'm creating a third site because I already have the GRC demo site for SQRL. We'll have the XenForo forum. So there'll be a third one which explores, well, it's designed for me to test that backend API. But it'll also be a neat way for people to play with SQRL.

And the other thing it offers, and the thing that the API has, is something we call the "managed shared access" because one of the things that SQRL fights against is people giving their username and password away. You know, how many people are like, oh, you need to log in? Here's my username and password. So sharing username and passwords is a common practice because, well, because you can. But SQRL doesn't really support that. I mean, you have an identity, and you don't want to give somebody else your SQRL identity because that represents you for the world. So we've solved that problem, and this allows people to understand how that works and to play with it.

So anyway, people say, my god, Gibson, how long is this taking? Well, yes, it's taking a while because really, really, really solving this problem is something that nobody else has done. Everybody else has sort of played at solutions or solved part of the problem. But really solving it all the way is a little more challenging. But, you know, I'm getting there. So this is not the SQRL episode.

Leo: That's a long way of saying no.

Steve: This is the Supply Chain episode, our Episode 684 for October 9th, where we're going to talk about the rather, I mean, calling it earthshaking is probably not overstating it. It's also controversial because the Bloomberg story, which is what begat all of this, lacks evidence. And whenever that's the case, you can have the affected companies, like Apple and Amazon, screaming their denials out loud. And we're left with the, I hesitate to say in this day and age, the infamous "he said, she said." But, you know, its they said this, and the other people said that, and they're diametrically opposed.

But the point is, independent of whether it's true, and in fact just minutes ago I became aware of an additional Bloomberg story from this morning thanks to our listener Simon Zerafa, who when I went to tweet the link to the show notes, which you have probably grabbed from my tweet, Simon had just tweeted the link to this updated story, which has additional corroboration of the problem, different fact set than what Bloomberg reported. And because the idea that a major motherboard supplier, that is, Supermicro, was having

at least some of its subcontracted motherboards infected with a hardware implant, that's big news.

And we got the word "implant" from the Snowden documents originally because this is something that our own NSA is alleged to have done. In fact, one of the things that they do is intercept shipments and make a subtle alteration of the hardware that cannot be found. So independent of whether this particular instance is true, talking about the supply chain in general is certainly worthwhile. So that's how we're going to end the podcast.

We're also going to look at a different sort of October Surprise which dropped for a number of people who went to update Windows 10 with the October feature update and were unhappy with the consequences. Thus the October Surprise and our Picture of the Week, as you just noted. We have a security researcher who has massively weaponized an existing pair of MikroTik router vulnerabilities and released it as a proof of concept on GitHub. A very clever voicemail hack which bypasses WhatsApp's authentication, its texting authentication, which we'll talk about. And actually we talked about something like this before. So it'll be interesting to see that it's actually happened.

We have the question of what happened with that recent Google+ breach, which has a lot of people screaming and claiming that Google was keeping this a secret. We'll look at that. Also they've tightened up their Chrome extension security policies, which is good. And a welcome change to the Wi-Fi Alliance's ever-confusing radio protocol designations, those 802.11a, what, asdf or who knows what. Anyway, that's being all changed. We've got Intel unwrapping its 9th-generation Core processors, which we look at, of course, from the standpoint of what have they fixed of this year's treasure trove of Security Now! content with the Spectre and Meltdown problems.

We've got some head-spinning PDF updates from Adobe and Foxit, where we have to remind them, guys, this is not a competition to see who can have the most problems to fix all at once. And then we'll talk about the inherent danger of supply chain attacks. And, I mean, what a real, I mean, I would argue insurmountable problem they represent if we're going to - if the lowest bidder is going to be the people who build our products, and they're not our friends. So I think something great to talk about.

Leo: Lots to talk about.

Steve: Lots of great stuff to talk about. Unfortunately, people last week who manually checked for updates were treated...

Leo: Seekers.

Steve: Oh, seekers, yes.

Leo: That's what they're called, yes.

Steve: Well, seek and destroy.

Leo: In fact, I installed it on three machines with no problem. But obviously not everybody did.

Steve: Yes. And certainly not everybody has a problem. Unfortunately, those who did had a big problem. So what happened was that the original plan was that the Windows 10 October feature update was scheduled to be part of today's Patch Tuesday, at which point it would have been automatically found and downloaded and installed. But it was also put up on Windows Update a week before, where people who manually check for updates would find it and could install it. The thing that, you know, in retrospect, and I know this is a cheap shot at Microsoft, but their page describing it says: "With the release of the Windows 10 October 2018 Update" - and that'll be version 1809, right now we're at 1802 - Microsoft says, "we're empowering a new era of personal productivity."

Leo: Clear your desk.

Steve: Yes, exactly. You know, all those files that have been piling up in your documents folder, really, those are just getting in your way. Those are slowing you down. We're going to get rid of those for you. So what happened was, for those who don't know, some who updated, presumably in anticipation of being empowered by this new era of personal productivity, soon discovered that all of their personal documents had been irretrievably deleted from the computer. Seriously. And that was only one of a number of problems that this thing had.

Now, okay. We know that when - apparently it had something to do with Windows Cloud syncing, when syncing was not in use, although we haven't yet had a full articulation of why some people were hit and others were not. But a chunk of people were. And we know that, when files are deleted from a hard drive or, well, any mass medium, the entry in the directory is eliminated, and then the space that the file was occupying on the medium is just marked as available for reuse.

Thus Norton's original claim to fame of "undelete," where he realized that, oh, you could remove the E5, which was the marker for this is an available entry, and guess what the first character was. Or the user could say, oh, I remember what that filename began with, using the rest of the name. And then, subject to the drive's fragmentation, which made guessing where the file was a little more tricky, you could unmark those sectors as free, that is, mark them as in use and oftentimes reclaim your file.

Well, that's exactly what happened here. When people were initially calling Microsoft saying, "Hey, all my documents are gone," the Level 1 support people said, "What? We don't know anything about that." Later, when Microsoft realized, oops, we have a problem here, and they started looking at it more closely, the advice changed to step away from your computer, stop using it immediately, and we'll get back to you.

Later down on the new era of productivity page, under the symptom - well, first of all they said "known issues updating to Windows 10 v1809" on that page. They added under "known issues" only one symptom: "We have paused the rollout of the Windows 10 October 2018 Update (version 1809) for all users as we investigate isolated reports of users missing some files after updating." Yeah. Like the ones that System Restore will not restore.

So under "workaround" they said: "If you have manually checked for updates and believe you have an issue with missing files after an update, please minimize your use of the affected device and contact us directly at 1-800-MICROSOFT, or find a local number in your area. Alternatively, use a different device to contact us. If you have manually downloaded the Windows 10 October 2018 Update installation media, please don't install it, and wait until new media is available. We will provide an update when we resume rolling out the Windows 10 October 2018 Update to customers."

And as this Picture of the Week shows, some poor guy named Robert Ziko posted on October 3rd, which would have been last Wednesday, under the title "Windows 10 October 2018 Update v1809 deleted all my files." He posted: "I have just updated my Windows using the October update (version 1809). It deleted all my files of 23 years in amount of 220GB. This is unbelievable. I have been using Microsoft products since 1995, and nothing like that ever happened to me."

And then he says: "Files were located at C:\Users\rober" - R-O-B-E-R, and of course his name is Robert Ziko - "\Documents." He says: "This location is still present, with no files. All of the files deleted. I am extremely upset. Not sure what to do. Please let me know."

So anyway, as we know, recovery is possible for those who were affected, subject to, you know, reuse of the space and also the degree of fragmentation. This is why I've still sort of thought, I mean, I don't defrag SSDs, but I do make images of my system. Like I have got a nightly imaging system, like a nightly imaging program now established, and with incremental images. And, I mean, I'm not taking anything - and of course that's what everyone says, right, in response to Robert is, well, you did make a backup; didn't you?

So anyway, this happened. Microsoft stopped it. It's gone away. It's not part of today's Patch Tuesday. In addition, there were reports that Task Manager was showing incorrect CPU usage levels. Microsoft Word, at least versions 2016 and 2019, would no longer save documents, no matter what anybody did - Save As, Save. But apparently that had something to do with add-ons, a conflict with Word add-ons such that, if you went into add-ons and disabled them, then Word would again save your documents. So if there is anybody who's listening to this who's been stuck and can't reboot or shut down or close because they've got unsaved documents, that apparently is a workaround. And then a number of games - Battlefield 1 and 4, Assassin's Creed, Hitman, Watch Dog, and Rainbow Siege - reportedly stopped working for some people after this. So Microsoft has some work to do. I'm sure they'll get it fixed. And hopefully people did have backups, and this did not represent a catastrophic loss for them.

For those who can, our takeaway is this is probably worth setting that deferral on feature updates to, like, 60 days. Unfortunately, Windows 10 Home does not give you that option, which is really infuriating, that all other versions of Windows offer it, but Microsoft said we're going to be stingy. We're not going to let Windows Home users defer the feature updates. But for those not using Home - Pro, Enterprise, and others - you can, under the Advanced Settings for Windows Update, Advanced Settings, you can elect to defer feature updates for up to 365 days, for up to a year. I would argue, you know, for old really cranky people, fine, set it for 365 days, and you'll never get a features update.

Sixty is certainly fine because, if a features update is going to cause a major meltdown, it'll be detected within two months, Microsoft will fix it, and your system, having patiently waited a couple months, will have been spared whatever pain that update may cause. So really I would say at this point definitely defer features updates for 60 days. I mean, you can always then decide you want to bite the bullet after a week of no problems. Or maybe set it to a week of deferral, I don't know. But clearly we're in a world now where this forward-rolling continuing eras of enhanced personal productivity being offered by Microsoft can sometimes backfire on people. So I would say not. Or if you have a plan to image your whole system and/or all of the important parts, then, yeah. Go ahead. Be a leader.

Okay. So this is really interesting. Unfortunately, MikroTik routers are back with us. An interesting and really good case study of leveraging multiple vulnerabilities unfortunately comes home to roost here. So we'll remember that the MikroTik routers which did not have a patch applied back in April ever since then have been vulnerable to an exploit which had widely become known because the patch itself was reverse-engineered, as we've often talked about can happen and is increasingly happening because people have

like a window of opportunity. Looks like it's bigger for MikroTik routers than places where systems are updating themselves autonomously.

But back in April MikroTik offered a patch. It got reverse-engineered, which allowed unauthenticated read-only access to any file on the router. Well, that's not good. What's also not good is the fact that the user database file that contains all of the authenticated username and passwords is stored in plaintext. No hashing of the passwords, no obfuscation, nothing. It's just so an unauthenticated user, as of April when this became publicly known, can obtain that file, which then allows them to look up the admin username and admin password in the clear.

And the issue then was what mischief could they get up to. After all, this was read-only access to any file. So it's not good, but it's also not the end of this router's life. However, it was possible to leverage that a little bit. There was a way to use the admin login credentials with that web interface to change a file. And in this case, as we reported, the error.html file was changed in the routers so that an HTTP error which went through the router's web proxy would return a page containing the Coinhive browser-based cryptomining software.

And remember that, at the time when we first talked about this, whereas a lot of these routers, as we've talked about recently, are $49 plastic boxes for homes and small offices, there are also big iron MikroTik routers for thousands of dollars rack mounting, which serve enterprises. And in those cases, if that web proxy got infected, then every user who hit an error page behind that router in an enterprise would, while that page was onscreen, be mining Coinhive. We don't know how much money that they were making. I don't think it was netting them a lot.

Okay. So that was the history. Near the end of August - so, what, six weeks ago, I think it was the 23rd of August - MikroTik security blog posted the following notice. They said: "Security issues discovered by Tenable," which is a well-known, reputable firm that we've talked about before. MikroTik wrote: "MikroTik was contacted by Tenable, Inc., who had discovered several issues in RouterOS web server. The issues only affected authenticated users," okay, so that's different. The previous one was an unauthenticated, nominally read-only attack. This one authenticated users, meaning, says MikroTik, "to exploit them there must be a known username and password on the device." Well, you can already see where this is going. Whoops.

"Your data, access to the system, and configuration are not under risk," writes MikroTik. "All the below issues" - four of them - "only allow the authenticated user," and they say "even a read-only user, to cause the www service to crash. Tenable has assigned CVE numbers to these issues." There are four of them. They're 1156 through 1159, of course in 2018. And their titles are "An authenticated user can trigger a stack buffer overflow." Which we know is not good. "File upload memory exhaustion. An authenticated user can cause the www binary to consume all memory." Or "Recursive JSON parsing stack exhaustion, which would allow an authenticated user to cause crash of the www service." And, finally, "Www memory corruption. If connections are initiated and not properly cleaned up, then a heap corruption occurs in www."

Okay. And then MikroTik finishes their security note, saying: "All of the above issues are fixed in the following RouterOS releases: 6.42.7, 6.40.9, and 6.43." So in all of those four potentially serious problems, MikroTik was careful to remind us and reiterate that only an authenticated user could cause these problems. That was good. However, in their disclosure they failed to be explicit about the fact that one of those bugs allowed an authenticated attacker to run any program of their choosing by making a web request to the router. But again, that attacker would need to be authenticated.

Okay. However, in that previous bug, what it allows, thanks to allowing the passwords file to be obtained by an unauthenticated attacker, that allows them now to become authenticated. So in other words, by leveraging both of these problems, neither one by themselves being devastating, but used in conjunction they provide - the first one, which is otherwise much less severe, provides enough information for the second one, which would otherwise be much less severe, to essentially allow an attacker to cause the router to run any externally provided file that they choose as root.

So unfortunately Tenable, and while still obeying reasonable disclosure guidelines because, after all, a patch was immediately forthcoming and was made available, they posted a proof of concept on GitHub with full source, showing how to build the source and how to remotely attack any doubly unpatched MikroTik router. In other words, MikroTik was relying upon the fact that no unauthenticated user could become authenticated since April in their representation of how bad this newer problem is. But we now know that the Internet is full of many hundreds of thousands of MikroTik routers which are still today vulnerable to the problem from April, meaning that they are now also vulnerable to this proof of concept, which is widely available, and exploitation of it has erupted on the Internet.

So at this point what we have is the news of its exploitation, and without specifics. We know that Mirai has used it. We know that it's been used by this cryptominer injection. And we know that, as we were discussing last week, the VPN filter that Cisco Talos Group spent the summer analyzing, is also another user of the previous exploit. Or maybe this one. It wasn't clear whether the attackers may have been leveraging this latest one, as well. But we now have a situation where, publicly known by leveraging both of these problems disclosed this year which persistently exist in hundreds of thousands of these MikroTik routers, it is possible to simply cause the web server to fetch a remote file and run it as root and gain additional strength.

What I liked about this was this was such a clean and clear example of the concept that we've often discussed, where multiple vulnerabilities by themselves in isolation may not be nearly as devastating as chaining them in order to create something more strong. And I would also say, again, we know anyone can make a mistake. Mistakes happen. MikroTik is responsible, in my opinion, in patching them and responding to them immediately. Props for that. But their design policies they are responsible for.

And the policy which we've discussed, as we've looked at this more closely, which is their policy, of requiring a sophisticated understanding of remote access security from an unsophisticated purchaser of a $49 retail plastic box for the home or small office is absolutely wrong. A lot of people, our listeners, have defended MikroTik, saying, hey, it's got all these bells and whistles. You can lock it down tight. And as we found, yes, but it doesn't come locked down tight. You don't have to turn these things that you want on. They're on unless you turn them off. And that is absolutely wrong.

And as a consequence, I would argue that that they have earned and deserve the reputation damage that they are currently suffering within the security community. Most users are not listening to this podcast, and they'll never know. Unfortunately, people who purchased those MikroTik routers and haven't updated them since April are inviting all kinds of problems in their networks. And now we know, as a consequence of what Cisco found, that these things are setting up shop. They are, what was it, six different ways of making sure they're not removed, and they're turning around and starting to look into the networks to see what's there behind them. Which is the worst-case nightmare for something getting onto your enterprise network.

Okay. Here's a clever attack, which is clever because, again, it's sort of multifaceted. It uses sort of the unintended consequences of the way an overly complex system works. We've talked about it in the past. A security expert warned about it in the past, like a

year ago. And it's happened. So a wave of reports about hijacked WhatsApp accounts in Israel induced the Israeli government's cybersecurity agency to send a nationwide security alert last Tuesday. Huh. We got our nationwide security alert on Wednesday. Anyway. Of course, that was just a test.

The alert, authored by the Israeli National Cyber Security Authority, warns of an effective method of hijacking WhatsApp accounts using the mobile provider's voicemail systems. And as I said, it's clever and tricky. But it's been known of in theory for some time. It was first noted a year ago by an Israeli web developer at Oath, and we talked about this at the time back then. The underlying realization was that WhatsApp users who have voicemail accounts for their phone numbers are at risk if they don't change their voicemail PIN default, which in most cases starts off at 0000 or 1234. And apparently, based on the provider, it could be obtained from the phone number. You can even - you don't even have to guess what the default is. You can find out what the default is.

So the opportunity for an account takeover occurs when an attacker tries to add a legitimate user's phone number to a new WhatsApp installation on the attacker's phone. In that event, the WhatsApp service sends a one-time code via SMS to the phone number being joined, the victim phone number being joined to the attacker's account. So this would typically alert a user that something fishy was up if they get, like, a confirmation of account connection from someone they don't know, and they're not expecting it; right?

But an attacker could avoid their target being aware of this by timing their attack at night or when for whatever reason they were sure the user was away from their phone. After several failed attempts to validate the one-time code sent via SMS, that is to say the attacker simply retries the join request several times, and the WhatsApp service notices this, finally it switches to prompting the user to perform a voice verification requiring no prompting at the user's end.

So the WhatsApp service calls the user's phone number and speaks the one-time verification code verbally. If the attacker has timed all of this properly so the user doesn't answer their phone, voicemail picks up, and the one-time - and everyone can already see where this is going now. The one-time verbal authentication code goes to voicemail. The attacker, knowing of course the victim's phone number, or knowing the voicemail pick-up, because remember that also most mobile telco providers allow remote access to a customer's voicemail account. So they use that service to pick up voicemail from the non-mobile phone; enter the default voicemail PIN, assuming that the user hasn't strengthened it or changed it from 0000 or 1234; recover the spoken one-time code; enter it inside their version of the WhatsApp app; and link the real user's phone number to the hacker's device, effectively hijacking the account from the legitimate owner.

Once the attackers gain access to the WhatsApp account, they can then enable two-step verification to the user's account to prevent the legitimate owner from retaking control of their WhatsApp account because they won't know what the six-digit one-time password is which only the attacker knows. And while this may seem like an awful lot of rigmarole to go through, the reporting of this upsurge has noted that it requires minimal technical skills, no specialized equipment, and according to Israeli authorities it's been widely used in recent weeks, leading to many reports of hijacked accounts.

In their alert, the Israeli authorities recommend that users use either a strong password for their mobile voicemail account - yeah, I mean, by all means. And you can imagine someone might think, eh, why bother changing it? It just doesn't seem like a security issue. But when you couple that with failed SMS fallback to voicemail one-time password, suddenly now your voicemail account can contain sensitive information.

Leo: I think most of us just have like a four-digit PIN on our voicemail.

Steve: Yeah.

Leo: That's pretty common.

Steve: Yeah, exactly. And the Israeli authorities also said take the initiative of adding one-time password protection to your WhatsApp account so that you've got, you know, like using a third-party app rather than SMS in order to authenticate. And then the bad guys will have no way of getting it. So again, sort of obvious in retrospect. But wow, an interesting sort of hybrid combo of things.

So we ask the question rhetorically: Did Google hide a "major" Google+ security flaw that exposed users' personal information? The headlines about this were breathless.

Leo: Yeah.

Steve: "Google hid major Google+ security flaw that exposed users' personal information." "Google+ shutting down after bug leaks info of 500k accounts." "Google to shut down Google+ after failing to disclose user data breach." "Google is shutting down Google+ after a secret potential data leak." "Google+ to shut down after cover-up of data-exposing bug."

Okay. So what happened? Given what Google's own security research auditors found, which is who found this back in March, and given all of the deliberately privacy constrained information Google had at their disposal, which is not much because they were deleting it on purpose to protect users, I think they've acted responsibly and appropriately.

The controversy here is that the Wall Street Journal caught wind of this discovery of theirs back in March, which Google fixed and didn't say anything about at the time because Google, despite looking for any evidence that anybody else knew of this bug, was unable to find it. Since their logs were incomplete, they couldn't be absolutely sure it had never been used, but they had no reason to believe it had been.

So yesterday the VP of Engineering posted under Project Strobe. The posting is: "Protecting your data, improving our third-party APIs, and sunsetting consumer Google+." And of course we've talked about Google+. I know you guys on...

Leo: I used to love it, yeah.

Steve: Yeah. But kind of just was one of those things that never got off the ground.

Leo: No. It was their, you know, they really put all the wood behind the arrows there. They even gave people bonuses one year based on Google's social success because they were trying to stop Facebook. But it didn't take off. And then once it was neglected, as any social network does, it became spam and porno.

Steve: Yup, yup. And very much like MySpace, which is just kind of like, eh.

Leo: Yeah.

Steve: Okay. So from their announcement, I pulled a few relevant bits. So they wrote: "At the beginning of this year, we started an effort called Project Strobe, a root-and-branch" - and I'm thinking, what? Is that like soup to nuts? I guess so. But that's a root-and-branch...

Leo: It's a corporate term of art. That's what that is.

Steve: "...review of third-party developer access to Google account and Android device data and of our philosophy around apps' data access." So that's all good. "This project," they wrote, "looked at the operation of our privacy controls, platforms where users were not engaging with our APIs because of concerns around data privacy, areas where developers may have been granted overly broad access, and other areas in which our policies should be tightened." So like a very what-we-want-from-Google-focused, self-imposed audit of their stuff.

So they continue: "Our review showed that our Google+ APIs, and the associated controls for consumers, are challenging to develop and maintain." I don't really - I didn't dig in enough to know what that means. That sounds like corporate speak for we didn't do it right or it's hard to do it right. I don't know.

Leo: APIs are hard.

Steve: Yeah. Tell me. Anyway, "Underlining this, as part of our Project Strobe audit, we discovered a bug in one of the Google+ People APIs." Okay. So now here's what we know and they wrote: "Users can grant access to their profile data, and the public profile information of their friends, to Google+ apps via the API. The bug meant that apps also had access to profile fields that were shared with the user, but not marked 'public.'" So these were meant to be kept private. So they were user-provided fields that the user did not intend to share publicly.

Leo: I'll give you an example what those are because as soon as this came out I went to my Google+ profile, which I set up ages ago. And there's mostly public stuff - your links, you know, stuff. And there were two things in there that were private. One was my personal phone number, and one was my address. I have no idea why I put my personal phone number and address in there because they weren't publicly viewable, but maybe there was something - I think there was...

Steve: May have been for account recovery at the time.

Leo: Could be. Or there were layers of sharing. I think there were. I can't remember. So worst case, I mean, nobody's putting a credit card number in there. Worst case...

Steve: Yes, and actually you can't.

Leo: Right.

Steve: Google explains: "This data is limited to static, optional Google+ profile fields," and they said, "including name, email address, occupation, gender, and age." Then in their posting they gave a full link. And it is worth noting that there is some additional stuff that's a little more worrisome. The full list contains a lot more. There is birthday, profile photo, relationship status, current and previous employers with dates, places previously lived, and skill sets. But again, still...

Leo: Yeah. But why would you - right. I don't know why you put those in if they're super secret.

Steve: Right.

Leo: I'm not sure. Maybe there was some incentive to do that.

Steve: So Google says it does not include any other data you may have posted or connected to Google+ or any other service, like Google+ posts, messages, Google account data, phone numbers - well, except in your case.

Leo: Well, I put them in. Yeah, I put them in. I had one public number and one private number in there.

Steve: Right. Or G Suite content. Then they say: "We discovered and immediately patched this bug in March of 2018. We believe it occurred after launch" - and I should mention that's 2015, so it had been present for years - "as a result of the API's interaction with a subsequent Google+ code change."

They said: "We made Google+ with privacy in mind and therefore keep this API's log data for only two weeks." Which in this case turns out to have been a double-edged sword. They say: "That means we cannot confirm which users may have been impacted by this bug. However, we ran a detailed analysis over the two weeks prior to patching the bug" - meaning they had those two weeks of previous log data - "and from that analysis, the profiles of up to 500,000 Google+ accounts were potentially affected." And that's a key word. "Our analysis showed that up to 438 applications may have used this [so-called people] API."

And then they conclude: "We found no evidence that any developer was aware of this bug or abusing the API, and we found no evidence that any profile data was misused." So I think they have behaved responsibly. Unfortunately, look what happens when an uncritical reading of the details gets loose. And so given that - and apparently the Wall Street Journal reported that they got a hold of internal memoranda which suggested that Google was concerned that Sundar would be yanked in front of a congressional committee a la Facebook and the abuse of Facebook's private information, and all kind of lumped into one mess, if back in March they had said we found a bug that we have no evidence was misused, but it could have been, and it could have affected up to half a

million Google+ users' private data. Again, an uncritical understanding of this even now led to those headlines.

So I think Google did the right thing. It would be nice to have had all of the logs of that. But then that could have been a problem, even a potentially bigger problem. So I don't think - as I said, it's a double-edged sword. I come down on the side of thinking that Google, given the situation, did the right thing. But I also recognize it's a little controversial.

Leo: I do wonder what their GDPR liability is because, you know, you have 72 hours to reveal a breach. But this isn't a breach. This is a bug that could have led to a breach, but as far as they can tell did not. So I guess they're not liable.

Steve: And was GDPR at the beginning of the year, or does it happen...

Leo: No, it wasn't in effect in March.

Steve: Right, right.

Leo: So they had no - yeah, I think, yeah. And I don't - it's not like a credit card or Social Security number breach. It's, I mean, certainly stuff people might have wanted to keep private and reasonably could have expected to be private.

Steve: Well, yeah. And if they marked it private...

Leo: It should be private.

Steve: Yes. Google is acknowledging that an API failure allowed this data that was not meant to be released to be released.

Leo: Right. Whenever I see an article about Google in the Wall Street Journal, I always go to it a little skeptically because the Journal hates Google because they compete with the Journal for advertising revenue, and has always in its editorial, which is odd, been willing to castigate Google, sometimes inappropriately. So I don't know if this is one of those. But we've observed this before with the Wall Street Journal.

Steve: Well, yes. And I would argue, too, that the technical press, all of those headlines that I cited were from the technical press. I mean, you know...

Leo: Yeah, they should know better.

Steve: Exactly. They should not just go for clickbait, but say...

Leo: Drives traffic. That's right.

Steve: You know, here's the story. So in this case I don't think Google did the wrong thing.

Leo: I'm glad to hear you say that. Good, yeah. It's sad they're closing Google+, but they also said in that post that you were referring to that 90% of visitors spent less than five seconds on the site.

Steve: Five seconds. It's like, oh, crap, where did I go? Hit back. Hit the back button.

Leo: Get out, get out, back out. It's a shame. Now, everybody who - like Mike Elgan and Trey Ratcliff. Trey actually wrote a long piece about what Google should have done with Google+. Those people, Trey Ratcliff was huge there. It really helped his career. And I think we're sad to see it go. It had such potential.

Steve: Well, and remember all the work that Gina Trapani did on that other thing that they, you know, I mean, Google does this. There's a bone yard.

Leo: Wave. She wrote a whole book about Google Wave.

Steve: Wave. Wrote a book. Yup.

Leo: Aw.

Steve: And then got washed up on the beach.

Leo: Washed ashore, yeah, yeah.

Steve: Okay. So Google has announced some additional information for their next release of Chrome, which will be 70, seven zero. We already know, because we've been talking about this a lot, all of that controversy in the display. They're not going to display files in the same way in the URL. They're going to simplify the URL by not displaying the www-dot at the beginning of those in 70. M-dot gets to stay for now. And we've also talked about - last summer we covered this - the problem with drive-by extension installation. It had been the case that websites could install Chrome extensions, which was a nice concept, I mean, the idea that you could have a website that would have like a site-specific extension that it could offer a Chrome user on the spot in order to enhance their experience with that site.

On the other hand, we know that installing anything on the spot is really fraught. And in fact standard wisdom now is never, never click something that a website you're visiting says you need for any reason whatsoever. Even if they're telling the truth, it's just not worth the risk that they might not be. So back on June 12th, Google announced their timeline for their gradual rollback of inline Chrome extension installation. They said on the 12th of June, they said: "Starting today, inline installation will be unavailable to all

newly published extensions," meaning that, if you weren't already in, you're not getting in.

"Extensions published," they wrote, "on June 12th, 2018 or later, that attempt to call the chrome.webstore.install function, will automatically redirect the user to the Chrome Web Store in a new tab to complete the installation." Which seems really, really very reasonable to me. Then, starting September 12th, so like three or four weeks ago, "inline installation will be disabled for existing extensions, and users will be automatically redirected to the Chrome Web Store to complete the installation."

So, okay. And developers have plenty of time, you know, what, June 12th to September 12th. So, what, is that six months? Or not quite, I guess, but five. So plenty of time to fix their website so that they're prepared for that redirect. And now, in early December 2018, the inline install API method itself will be removed from Chrome 71. Not 70, so you can still do it. But inline install API going away with the version after the one that we get in a couple weeks.

Okay. And now, last Monday, in their posting titled "Trustworthy Chrome Extensions, by Default," what they announced was an additional set of steps which will be appearing in this next release, so in a couple weeks. The thing that I'm - okay, there are several things that I'm excited about. One is they are giving users control over an extension's access to hosts explicitly. So in other words, it has been the case that if you put an extension in Chrome, it essentially became an equal participant with Chrome and could see into all of the sites you visit, just as Chrome can. So, I mean, so it really needed to be trustworthy.

What's changing is users can now change this to select specific sites where the extension will be active, require you to click the extension's icon to activate it, or allow it to be globally active. So we've never had those permissions. So underneath the extension, there's "This can read and change site data," and then you choose from three radio buttons: when you click the extension; on, and then it fills in whichever site you're on, like in this case in the screenshot material.io; or on all sites. So you can restrict an extension to one particular place, or shut it down completely and manually activate it. So that's nice. We get that in two weeks.

And they said: "Our aim is to improve user transparency and control over when extensions are able to access site data. In subsequent milestones," they said, "we'll continue to optimize the user experience toward this goal while improving usability." They're also going to change the extensions review process. This posting was aimed at developers, so there's some developer-aimed language. But they said: "Going forward, extensions that request powerful permissions will be subject to additional compliance review. We're also looking very closely at extensions that use" - and this makes me shudder - "remotely hosted code." I mean, the idea that an extension could be allowed to just go get ad hoc code from somewhere is horrifying.

And they said: "...with ongoing monitoring." That is, they will be monitoring this. Extension permissions should be as narrowly scoped as possible, and all code should be included directly in the extension package to minimize review time. So to me, this feels like another thing any sane developer will read the tea leaves here and realize that in the future that's going away completely. Your extension will not be allowed to go reach out and get whatever it wants from wherever it wants. So start reengineering that now so that you're not inconvenienced later because this really feels like it's going to go away.

And then this is really good. Under new code reliability requirements - and again, it's like hard to believe they've allowed this. "Starting today, Chrome Web Store will no longer allow extensions with obfuscated code. This includes code within the extension package as well as any external code or resource fetched from the web. This policy applies

immediately to all new extension submissions." So turn off code obfuscation for anything new you submit. They said: "Existing extensions with obfuscated code can continue to submit updates over the next 90 days, but will be removed from the Chrome Web Store in early January if not compliant."

They said: "Today, over 70% of malicious and policy violating extensions that we block from Chrome Web Store contain obfuscated code. At the same time, because obfuscation is mainly used to conceal code functionality, it adds a great deal of complexity to our review process." Yeah, no kidding. I mean, they've got to deobfuscate it and then figure out what the heck it's doing. They said: "This is no longer acceptable, given the aforementioned review process changes."

And they also said: "Additionally, since JavaScript code is always running locally on the user's machine, obfuscation is insufficient to protect proprietary code from a truly motivated reverse engineer. Obfuscation techniques also come with hefty performance costs such as slower execution and increased file and memory footprints." And they said: "Ordinary minification, on the other hand, typically speeds up code execution as it reduces code size and is much more straightforward to review. Thus, minification will still be allowed, including the following techniques: removal of whitespace, newlines, code comments, and block delimiters; shortening of variable and function names; collapsing the number of JavaScript files."

They said: "Extensions in the store with obfuscated code must be updated using the recommended minification techniques for Google developers and submitted before January 1st, 2019." After that, they're being removed, to which I say bravo. They should have never been allowed in the first place. I mean, I guess I can understand someone believing that they're protecting some proprietary secrets by encrypting a JavaScript download. But exactly like with DVDs, where the DVD player has to have the decryption key in order to decrypt it to show it to its owner, the web browser has to decrypt the JavaScript in order to run it. So it never made any sense. So goodbye to obfuscation, especially when 70% of malware is obfuscated. And, I mean, I want Google to be able to understand and audit the extensions that I'm getting from the Chrome Web Store. I mean, so, yeah.

Also, on the issue of account takeover, because as we know that was how there was that recent problem that we covered, they are going to require two-step verification for developers logging onto their Web Store accounts. They said: "In 2019, enrollment in 2-Step Verification will be required for Chrome Web Store developer accounts. Popular extensions," they write, "can attract attackers who want to steal it by hijacking developer accounts, and 2-Step Verification adds an extra layer of security by requiring a second authentication step via phone or physical security key. Google strongly recommends enrollment as soon as possible, and it will be required in 2019." So, yay. I'm glad that Google is continuing to move forward to strengthen things.

And Leo, oh, this was just too long in coming. Okay. So we ask rhetorically, what's the latest version of WiFi? Well, those of us who've been paying attention can issue the magic incantation, 802.11ac.

Leo: Or ax.

Steve: Or ax. Or, well, yeah, actually, that's coming.

Leo: Yup.

Steve: Or what about "n"? What about "g"? What about "a"? What about "b"? Or a combination of. So the Wi-Fi Alliance, the group that I would say I love to hate them, but no, I just hate them, they manage the implementation of WiFi, much to all of our disservice, by not making any of their specification production process public, as I've often complained about. We now have WPA3 coming out, but only people who are members, who are paying dues to be, and they're not cheap, have seen these documents. What they publish is just the Table of Contents, thank you very much. So we'll have to wait until it leaks, as it of course inevitably does because people have to implement it, like in Linux, which is open source.

Okay. So big change coming. The next WiFi standard, correctly, as you noted, Leo, 802.11, wait for it...

Leo: Ax.

Steve: A-X, yes, ax, the ax version, will instead use a simpler naming scheme, simply being called WiFi 6.

Leo: Yeah.

Steve: Yeah. And so we have backported. The reason it's 6 and not 1 is that 802.11ac is WiFi 5; "n" is 4; "g" is 3; "a" is 2; and "b" is one. So gone. So now, I mean, and this is good. For the typical consumer, we'll find out like the next - I guess at some point we'll get an update from Google with Android and Apple with iOS that our new devices support WiFi 6, if they don't already. Maybe they do. And you'll go, when you're shopping for your next WiFi router, it'll just say, you know, "Supports WiFi 1 through 6," probably. And it's like, oh, good.

So the Wi-Fi Alliance wrote: "To help users identify devices that provide the latest WiFi experience, Wi-Fi Alliance has introduced simplified generational names that may appear in device names and product descriptions. WiFi devices supporting the latest generation of connectivity are based on the 802.11ax standard and are known as WiFi 6 devices."

And then they've also produced a series of icons which they're hoping that our OSes will incorporate into the UI. And I've got a picture of it here in the show notes. And I hope people understand that that's a version number and not how many people are connected because...

Leo: It's not clear at all.

Steve: It's not at all clear, exactly. It's like, wait. There are six people on this WiFi? Maybe I should use the one that only has four people. I don't know. Yeah, maybe six is better or stronger, who knows. Anyway. So for what it's worth, we're entering WiFi 6, which is a - what was it WiFi 6 was supposed to do? It's mo' betta. Anyway, that's enough.

Leo: I can't remember, actually.

Steve: It's higher data rates, increased capacity, better performance even in dense environments such as stadiums or public venues, and increased power efficiency, making it a worthwhile upgrade for smart home and IoT devices.

Leo: Oh. Just as you said, it's mo' betta.

Steve: Yeah. Yeah. Mo' betta. We want 6. Get 6. We also want 9th-generation Core processors from Intel.

Leo: Yes, we do.

Steve: Which they unwrapped yesterday to much fanfare. They of course did a whole bunch of them. I was focused on the desktop, where, boy, are these things getting fast.

Leo: Twenty-eight cores. Holy cow.

Steve: Yeah. Yeah.

Leo: Every one of them with speculative execution built right in.

Steve: Built right in, yup. Just speculate all you want. And also in the case - oh, there is, for example, the i9-9900, the K versions, which are the unlocked speedy guys. This 9900K can run up to, can burst to 5 GHz speed. And let's harken back to the 4.77 MHz. So what is that, that's...

Leo: It's faster.

Steve: A thousand times. It's fasta, mo' fasta.

Leo: Mo' fasta.

Steve: Yes. You can also cook your breakfast on it because these things...

Leo: 265W on the Xeon. That's a lot for a processor.

Steve: Yeah. Maybe, you know, if you had one of the heat sink with fins, you could just use it to dry your hair. Oh, my goodness.

Leo: They do make a bitcoin miner that doubles as a room heater. I think this would be perfect for that.

Steve: Yes, they do, actually, yes. Just blow some cold air over it, and it'll come out hot. So the good news is they have built, as one would expect, some additional mitigations are built in. There was a slide during yesterday's presentation read: "The new desktop processors include protections for the security vulnerabilities commonly referred to as Spectre, Meltdown, and L1TF," which of course we've covered extensively all year. "These protections include a combination of the hardware design changes we announced earlier this year, as well as software and microcode updates."

Okay. So in this case there was time to change the hardware; whereas of course until now we've had to settle with just software and microcode. So there are the five problems. The Spectre V2 branch target injection, that's not a hardware-based fix. It is microcode and software. But of course in this case it's already got the microcode. And all of the latest Windows and Linux and I assume Apple, I have not been really keeping my eye on them, but the OS knows about the new features which the microcode offered. So whereas the challenge for existing systems has been you would need to somehow get your microcode updated, which could sometimes be impossible, these of course come with the microcode current.

The second was the Meltdown V3, which was the Rogue Data Cache Load. That is fixed in hardware. So the new chips have it. It's not possible to have that in the old chips. The third one was the V3a, which was Rogue System Register Read, fixed in microcode and now built into those chips.

The fourth one is the Variant V4 Speculative Store Bypass. That, as was the case with the first one, it was Spectre V2, not subject to hardware fix. That's microcode and software. And we've got the microcode, and we've got the software. So everybody will be in good shape with that. And the last one was the L1TF, the L1 Terminal Fault, which was 100% hardware fixable and is fixed with this latest round of chips. So it's 9th Generation. They are i9 down through i3. So I guess they're all 9000 series; right? Because the i3 is a 9000T with four cores and four threads, so it's not hyperthreaded. And it runs at 3.2GB. And then it goes all the way up to i5s and i7s and then the i9 monster.

In this case, this one has eight cores and is hyperthreaded and nominally runs at 3.6GB and can be bursted up to 5. Yikes. Of course, it won't run Windows 7. So too bad for Windows 7 people. Thus Microsoft ultimately wins the battle of the Windows versions by not supporting the older OSes on the newer hardware. Which of course they announced some time ago.

Okay. So Adobe and Foxit need to understand that this is not a competition. Last Tuesday of this month, Adobe released security updates for Windows and Mac versions of Acrobat and Reader, including 47 critical vulnerabilities and 39 that were merely important. Forty-seven. Of the 47 critical vulnerabilities, 46 of those allowed for remote code execution; one allowed for escalation of privileges. Whereas the 39 important vulnerabilities were information disclosure. So, and remember that in the case of a PDF document which a reader is reading, there's never been, well, it's probably safe to say there's never been an interpreter as troubled and incredibly complex as a PDF interpreter. And we know how difficult it is to get an interpreter to work right. But they haven't yet.

And so 47 critical vulnerabilities, 46 of them remote code. And this means, if you've got your web browser configured to use Adobe Reader as its reader, which I don't think anyone should do anymore, you know, use the built-in PDF viewers. But if you did, and somebody sent you a deliberately crafted PDF, they could run code of their choosing on your computer. So it is a huge, huge attack surface for today's machines.

However, that was not the largest number of PDF problems. Not to be outdone, Foxit needs to be updated, too. So if you have left the Adobe train and switched over to Foxit -

I'm using Nitro as mine. When I switched to my - when I rebuilt this new system after my WinXP machine died, as we'll remember, a few months ago, I had Acrobat there, and it was a fully purchased, licensed one. But there was no way to recover the license, and Adobe wanted, you know, I just didn't want to go with them. So I ended up choosing Nitro PDF.

Leo: That's good. I think it's good.

Steve: I looked at 12. And many of them are nice, but they do not produce small PDFs. They produce big, lazy, blobby PDFs. And so, no, Nitro.

Leo: I should have because this morning when Google announced its new Pixel Slate, they announced: "And it comes with Adobe Acrobat." And I went [moaning]. Really? But it is a problem on Chrome OS to do PDFs, so maybe they had to do it.

Steve: Yeah, yeah. Anyway, Foxit, their update fixes 116 vulnerabilities.

Leo: Holy cow.

Steve: I know.

Leo: Is there something inherent in PDF rendering?

Steve: It's just really, really difficult.

Leo: It's an interpreter; right?

Steve: Yeah, it's an interpreter. And it's probably, I mean, something like Foxit is probably based on old Ghostscript code, which they grabbed, and then they took in-house. And Ghostscript was written just with good intentions, but once upon a time to sort of be a free and open source alternative for PDFs. But programmers who didn't have a commercial interest just said, okay, let's accept this. I mean, because a PDF is, as we know, it is a fully parameter-driven tokenized description of a page. It's a page description language. And it is really extra difficult to, like, check every possible special case.

I was just telling Lorrie the other day that I'm creating this third website web portal for SQRL. And a lot of it is just plumbing. It's like, if the user clicks the login button when they've left the username and password blank, because I also explain, I show how you can use it alongside of SQRL, I have to say, you have to fill in your username and password if you're going to log in. Or if they leave either one of them blank, but not the other, then I have to say, I mean, the point is, a lot of this is just brain-numbing but necessary plumbing that you have to do. And in my case here it's just UI.

But in the case of a PDF reader, which in this world now where anything you expose to any content externally can be malicious. And it just, I mean, it's really - we're in a

different world than we were decade ago, where it's like, oh, email macros, how cute. It's like, oh, look. It's a malicious email macro. How sweet.

Leo: Awww.

Steve: Yeah, isn't that adorable. You know? Now it's just like, oh, my god. Even our AV software creates an additional attack surface. So you put that in to protect yourself, but it's buggy because it's having to interpret what's coming in, too. Anyway, update Foxit. If you're a Foxit user, make sure you're running an update. It's probably a lesser large target than Adobe. But, for example, if a corporation were known to have standardized on Foxit, and somebody wanted to target that corporation, they could send one of their employees a PDF designed to remotely execute code.

We owe Cisco's Talos Group for these vulnerabilities: 18 of the 116 alone were discovered by Cisco's Talos Group. And all of the 18 vulnerabilities, as well as many of the others fixed by this update, are labeled "critical" because they could lead to code execution if you just visit a website, and you're using the Foxit Viewer on your browser. Or let alone if you receive a PDF in email, you make sure it's from someone you trust, you open the PDF, and wham. So be careful.

And I saw a note in my mailbag from Mark Taylor in - I didn't practice pronouncing this ahead of time, Leo - Wausaukee?

Leo: Wausaukee.

Steve: Wausaukee, okay. I felt like it needed to have another syllable in there.

Leo: It feels like that, yeah, yeah.

Steve: Yeah. Lot of vowels. Wausaukee, Wisconsin. Anyway, his subject was "SpinRite and encrypted drives." And I see this question a lot, so I just thought I would take a minute to respond to Mark and anybody else who's curious. He says: "Steve: All the normal accolades. Can SpinRite be used on encrypted drives without making the drive unusable because of moved sectors that were fixed? Thanks, Mark." And the answer is yes, SpinRite can. This version of SpinRite is completely happy if the drive is encrypted. It looks at it. It sees that it has no idea what this drive is. So it just considers it a blob of opaque storage, and it fixes the sectors.

So, I mean, SpinRite never really cares what data is there. In years past, it got itself much more involved with the file allocation table and following chains and all that. A much enhanced version of that will be making a reappearance in 7.0. But all of the 6.0 versions, where we are with 6.0 now and the forthcoming .1, .2, .3, they will continue. They will be much faster and much more capable, but of only doing the same thing, which is fixing the drive that you have without moving anything around.

So SpinRite never changes the content of any data. If a sector is unreadable, then it can often, as we know, make it readable again or show the drive that it is dangerously close to becoming uncorrectable, in which case the drive will move the sector. But what happens is the data is moved into the new sector, and the old sector is taken out of service. So, and this is happening in the background even without SpinRite. SpinRite

augments that process. So definitely worth doing, and it works just fine on encrypted drives.

Let's talk about this mess with our supply chain.

Leo: Yeah, and I should mention that Microsoft has put out an update. They think they understand now what was going on with the Windows feature update, and they are pushing out a new version of it, or will be soon. They're pushing it out to Windows Insiders. They say: "We've fully investigated all reports of data loss, identified and fixed all known issues in the update, and conducted internal validation." And they're going to push it out to the Insiders before rolling it out more broadly.

"It appears" - I'm reading from The Verge - "that the bug that caused file deletion was related to Windows 10 users who had enabled a feature in Windows called Known Folder Redirection, to redirect folders like desktop documents, pictures, and screenshots from the default location." I've done that. In the past I've had my documents folder moved to D. And you can do that. Microsoft has a facility to do that. Microsoft introduced code in its latest update to delete the empty and duplicate known folders. So when you move the folder to D, it's moving all the contents. But if there was stuff left behind, perhaps intentionally, it would delete that, thinking it was an empty folder. But it appears it wasn't always empty. I can understand why they would do that because normally when you use redirection, that's because the data folder, the documents folder is now on D. It's not on C.

Steve: Sure.

Leo: There shouldn't be anything left. But I guess that's not always the case. Microsoft has developed fixes to address a variety of problems related to these folder moves, and these fixes are now being tested with Windows Insiders. They say, and this is about what I would expect, that they believe that the data loss happened at a rate of one 100th of 1%. So that's what, one 10,000th? One 1,000th? 100 times 100. So that's one 10,000th. One in 10,000. But if you have a million people install that, one in 10,000 is quite a few people. So that's what happens when you have a massive user base.

Steve: And Robert Ziko is not happy.

Leo: Yeah, he's the one 100th of 1%.

Steve: Oh, and I did forget to mention that, with the October feature update, there's another change that I just wanted to point out that might catch out our listeners. And that is that, for those who use the disk cleanup utility, and I'm a big fan of it because it just goes through and gets rid of a whole bunch of crap, they've added the downloads folder. And so pay attention. It's not checked by default. But if your habit is to turn all the checkmarks on, as has mine been, and then say, yeah, get rid of all this junk, if you are a person who knows that you've got downloaded things that you deliberately have left in the downloads folder, make sure that you pay attention if you turn that on because this update to disk cleanup adds a line item that you can turn on, and it will wipe out your downloads folder.

Leo: Yikes.

Steve: Yeah.

Leo: We thought they were going to take that out, so I'm glad to see that it's not gone because, you're right, I use it all the time.

Steve: I still like it, too. And they're replacing it with their smart something or other.

Leo: Of course.

Steve: Fixer-upper, cleaner, triple-scoop whoop-de-do. But I don't know.

So I will begin by sharing an update on this Bloomberg story, which was just posted a few hours ago. A major U.S. telecommunications company, and I'll explain why we don't know who in a minute, discovered manipulated hardware from Super Micro Computer, Inc., in its network and removed it in August, fresh evidence of tampering in China of critical technology components bound for the U.S., according to a security expert working for the telecommunications company.

The security expert, Yossi Appleboum, provided documents, analysis, and other evidence of the discovery following the publication of an investigative report in Bloomberg Business Week that detailed how China's intelligence services had ordered subcontractors to plant malicious chips in Supermicro server motherboards over a two-year period ending in 2015. Appleboum previously worked in the technology unit of the Israeli Army Intelligence Corps and is now co-Chief Executive officer of Sepio Systems in Gaithersburg, Maryland. His firm specializes in hardware security and was hired to scan several large datacenters belonging to the telecommunications company. Bloomberg is not identifying the company due to Appleboum's nondisclosure agreement with the client.

Unusual communications from a Supermicro server and a subsequent physical inspection revealed an implant built into the server's Ethernet connector, a component that's used to attach network cables to the computer, Appleboum said. The executive said he has seen similar manipulations of different vendors' computer hardware made by contractors in China, not just products from Supermicro. Quote: "Supermicro is a victim. So is everyone else," he said. Appleboum said his concern is that there are countless points in the supply chain in China where manipulations can be introduced. And deducing them can in many cases be impossible. That's the problem with the Chinese supply chain, he said.

So that's a perfect preamble because, as it happens, I didn't title this "Supermicro Attack," I titled our podcast "The Supply Chain" because this is a bigger problem than just this one particular "in the news at the moment" issue. And I don't know how it has a solution, Leo. I mean, we covered - what was the chip? There was another chip that had a backdoor in it. Oh, it was an Ethernet chip. Was it RealTek?

Leo: I remember that.

Steve: It was somebody's - yeah. I mean, so...

Leo: We've known this is a problem for a decade. Brian Krebs said he remembered 10 years ago when he was writing for the Washington Post he found a Chinese printer or a hardware update to a printer that would copy everything you printed and send it to a server controlled by the Chinese military.

Steve: Yes.

Leo: This stuff's been around forever.

Steve: So there is a fabulous link, the last link I've got on the last page of the show notes, to LightBlueTouchPaper.org, which was written on the 5th of October, so that's Friday, "Making Sense of the Supermicro Motherboard Attack." I commend this to our listeners to read. It's an engineering paper. And so the guy obtained the baseboard firmware from Supermicro and disassembled it to take a look at it, to better understand what's going on. It is the case that the world has switched from parallel interfaces to serial interfaces. All of this PCIE stuff, you know, the x4, the x8, the x12. Well, that x number is the number of paralleled serial interfaces. So you could have x1, which is a single connection, which nonetheless runs at crazy high speed. And in fact USB is a serialized protocol.

And it turns out that, because the interconnection density has become a problem, speed is not. So it's easier to send bits out faster from a single pin than it is to have 32 or 64 pins. That's just too expensive in this day and age. And they're also extremely noisy. Every time you have signals going up and down you are radiating energy. And you're also having to supply power to pump the capacitance of that signal up and down. So the more copper you have, the more electrons you need to pour in and suck out in order to raise the voltage and lower the voltage again. So it's just a mess.

So serial is the future. And as a consequence there is - it's called the SPI, the Serial Peripheral Interface, which is a well-defined single pin. Actually there's a clock line and a data line. So it's a two-wire interface. And people who have studied this particular instance with the Supermicro motherboards and have seen the pictures that Bloomberg posted of what looks like little - they describe it as a grain of rice and looking, I think they called it like a "signal conditioner." Well, I'm sure that's common speak for a capacitor. And it showed six connections.

And people have said, how can that possibly do anything? Well, it could be that the motherboard allowed for expansion by having an additional SPI flash ROM, and they can be chained in series. So if there was a spot on the motherboard that was unpopulated, all somebody would have to do is to move the official chip over into the second place and put the SPI chip in in the first place. And six pins, you need power and ground and clock and data - clock in, data in, clock out, data out. That's exactly six pins. So that could work in order to allow that chip to have its contents downloaded by the processor when it starts up. That baseband processor is an ARM9. It's an older version of an ARM9 processor running Linux.

So we just - we see a motherboard that doesn't have the Intel 265W processor and its heat cooling technology on it and everything else. We sort of look at it, and it just looks like a bunch of sockets. But there is a Linux OS on that motherboard from the manufacturer, which starts up and brings our system to life. And we've talked about how the baseband has Ethernet communications capability, how it's able to be on the network by itself, how it can communicate even without an OS on top and loaded. So it's fully

awake and aware and has drivers for Ethernet in addition to all the other peripherals that it needs to set up at boot time.

So from a standpoint of skeptics who doubt the technical veracity, it is very clear from an engineering standpoint that something like this is completely possible. I would also argue that, I mean, that this is huge. It is probably understood and is probably enjoying some containment within our intelligence services because we're talking about an economic global partner in China with whom we do a lot of trade, who is apparently, if we believe the facts claimed, where their intelligence services have engineered this technology to infiltrate the hardware of a U.S. major motherboard manufacturer.

Supermicro isn't as well known as Gigabyte and ASUS and some others. But it is a major player in the server farm business with at least 30 customers, among whom are Apple and Amazon. And so an infiltration of this sort, I mean, has epic scope and consequence and really does, I mean, I have no reason to disbelieve it. Apple and Amazon's denials probably, given the fact that Supermicro's stock dropped in half, down 47% in one day after the news...

Leo: Somebody believes it.

Steve: Yes. And then it dropped another chunk, I don't remember how much. There wasn't much left for it to go. But this most recent news of this morning's report of an Ethernet connector having an embedded SPI hardware implant knocked its stock down even further. So, yikes.

Leo: You know, I guess for this show and in general our audience, the takeaway, you can debate about whether this happened, whether Supermicro knew, why are Apple and Amazon and others so vigorously denying it, and yet Bloomberg stands by their story. All of that is a lot of heat, but there's no question that these kinds of supply chain modifications happen. I was talking to somebody who buys a lot of servers, like a lot of servers. And he said it's routine. It's unusual for it to happen at the factory, much more often happens in transit. It's usually done as a targeted attack. So you know that we've got these 15 servers going to Sony. Let's modify those.

Hardware modifications are not unusual, like this Ethernet port modification. They're very difficult to detect and often go undetected. And because the supply chain, well, even if the supply chain weren't in China, we don't have a secure way of transiting this stuff. So, and by the way, the other final point is the U.S. does this just as often as every other country.

Steve: As I said, we learned - the first time we used the word "implant" on this show...

Leo: It was ours.

Steve: ...was Edward Snowden because the NSA was implanting and doing exactly this. They were intercepting packages bound for targeted targets, and making a change.

Leo: Right, right. So there really isn't anything in this story besides the specific allegations about Supermicro and the fact that they were ending up in Elementals,

which were ending up in the Department of Defense and the Mormon Church and TWiT, by the way. Those specifics you could debate. But it doesn't essentially change the point, which is this is happening, and has been happening, and is extraordinarily difficult to stop. Right?

Steve: Yeah, yeah. And the only, I mean, so I wanted to make it very clear to our listeners that technically there is no problem with this. That little pea size, that grain of rice size thing could very likely be a serial peripheral interface flash ROM. That's not in question here. And the idea that the motherboard, they're typically 32K, the motherboard engineers, the designers, could have just left some pads available for expansion space so they wouldn't have to redesign the motherboard. They would just populate the six little pads with an additional grain of rice if they needed more than 32K of storage in the future.

Leo: Right.

Steve: And things like infected Ethernet connectors, that's not science fiction, unfortunately. We may wish it were, but it isn't. And I think, Leo, the problem with this kind of hardware, I would say implantation at scale, is that it seems like it was a spray. What we know from Bloomberg's reporting is that it wasn't Supermicro themselves, nor their Chinese plant, but probably the subcontractors they use for overflow capacity that took the opportunity. And so somebody somewhere in China took it upon themselves, we don't know from how high up the orders came, if we believe Bloomberg - and again, we have no evidence at this point, and I don't think we're going to. I mean, again, Apple cannot have it be known that their server farms were infected with malicious hardware from China. They can't. Nor can Amazon. So unfortunately, for the good of their stockholders and their reputation, they have to say no. I'll bet there are people who know otherwise.

Leo: If this really happened, there will be more. There will be motherboards out there that are modified in that way, and they will surface. There must be thousands. Alex Lindsay says he has...

Steve: I would imagine tens of thousands.

Leo: Yeah. Alex Lindsay has an Elemental with the old Supermicro motherboard from 2015. I mean, it must be hard to detect. We've got to send it to Chipworks or something for them to figure it out. But this will - I think we've not heard the end of it, I guess.

Steve: Yeah, well, and this morning we heard something, another credible source, who says he was asked by a major telecommunications firm to inspect their datacenter. He found data traffic from the baseband processor that was unauthorized, traced it back, and found a fake Ethernet connector on a Supermicro motherboard that had been affected.

Leo: How hard is that to do? Wouldn't that be the first avenue is to see if there were unexpected traffic? Or is that just very hard to find?

Steve: Yes, yes. And in fact that's how this was found. It was a ping that was phoning home. And somebody - but again, Sony, in your example, it was perfect. Where I was headed was the idea of spraying the world with this is dangerous because someone in the world is going to find it. Whereas modifying 15 servers that are headed for Sony is much less dangerous, just because of the law of numbers. And Sony isn't a mainstream, like, datacentering is not their business. They're just in business to do something else.

And so computers are just arriving, and some IT guy is slapping the server blades into a chassis and firing it up and loading the OSes. And so they would never be any the wiser. Nor would they be as likely to catch a little bit, a little ping leaving in order to phone home. Whereas it was traffic, pursuant to this story, that tipped off somebody that, wait a minute. What? And our listeners will remember all the trouble I had when I upgraded my server. I was having all kinds of weird behavior on an Intel motherboard until I moved the connector away from the primary NIC1 to NIC2, and all the problems went away, because something about the baseband processor, which is only on the primary NIC, was causing these problems.

Leo: It's a fascinating story for a variety of reasons. It really does resonate with the political environment. It is a "they said, they said" story. It's fake news, according to some. It's the god's truth according to others. It's really an interesting story. I'm sure we'll hear a lot more.

Steve: Anyway, for anyone who's interested, this BlueTouchPaper, LightBlueTouchPaper.org piece, "Making Sense of the Supermicro Motherboard Attack." If you're interested, it's written by a techie and was sent to me by a good friend and has definitely got a nice take on this.

Leo: Yes.

Steve: I mean, like from a, you know, is this reasonable? And the answer is yeah.

Leo: Yup. Oh, I think this is a story that's only just begun.

Steve: Yeah.

Leo: But this show has come to an end. So there you go. But we'll be back every Tuesday about 1:30 Pacific, 4:30 Eastern, 20:30 UTC. It's Security Now! time, and a jolly cheer goes up throughout the geek community as Steve takes to the air and explains it all. I mean, you read a story like that, and I know my only thought was, well, I can't wait till Tuesday. I can't wait.

Steve Gibson is at GRC.com. That's where you'll find all of his stuff - including his bread and butter, SpinRite, the world's best hard drive recovery and maintenance utility - plus a lot of free stuff, all the information you'd ever want to know about SQRL, Perfect Paper Passwords, the Healthy Sleep Formula, I mean, that goes on and on. I was just the other day looking up the Vitamin D stuff that we did, those Vitamin D podcasts, because I was trying to figure out what dosage to get. It's all at GRC.com. And this show is, as well. He has audio versions of Security Now!, and he's the only source for really good transcriptions written by Elaine Farris that make it easy for you to read along while you listen: GRC.com.

We have audio and video at TWiT.tv/sn. You can also subscribe if you'd like in your favorite podcast application. That way you'll get it every episode. You really want to have it promptly every week, and you want to keep a backlog, too. The archives are just as valuable on this show.

Steve, have a great week. Thank you for being here. We'll see you next time on Security Now!.

Steve: Thank you, my friend. Till then.