Lamphone

Description: This week we address an accident that the Brave browser guys regret. We take a look at last week's Patch Tuesday and its several ramifications and consequences. We note a few odd new and unwelcome behaviors from this year's 2004 Win10 feature update and dip into yet another side-channel attack on Intel chips. But we also note that a long-awaited powerful antimalware technology is also about to ship from Intel. We look at the latest new SMB vulnerability named SMBleed, and conclude with an examination of the latest and more-practical-than-most techniques for covertly eavesdropping on a remote location - via a hanging light bulb.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-771.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-771-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Coming up we'll do trigonometry in assembly language. It's actually pretty cool, believe it or not. We'll also talk a little bit about the 129 vulnerabilities Microsoft patched on last Tuesday, and the light bulb that listens. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 771, recorded Tuesday, June 16th, 2020: Lamphone.

It's time for Security Now!, the show where we cover your security, your privacy, your safety on the interwebs with our safety officer, not wearing a red shirt today, that's good news, Mr. Steve Gibson.

Steve Gibson: No, those red shirt guys, they beam down to the planet, you know they don't have much longer.

Leo: That's it.

Steve: That's it.

Leo: You don't want to join that landing party. No way. No way.

Steve: That's right. That's right.

Leo: So what's up?

Steve: So we've got, for Episode 771, Lamphone, L-A-M-P-H-O-N-E. Our friends at the Ben-Gurion University of the Negev are at it again. They decided to find out whether a light bulb hanging down from the ceiling picks up enough vibration from the conversations in the room for them to listen in at a distance.

Leo: And I'm guessing, because we're talking about it...

Steve: Yeah.

Leo: ...that maybe it can.

Steve: It worked.

Leo: Wow.

Steve: So but of course, well, and actually originally that was not going to be the main story. There were a whole bunch of new vulnerabilities that have appeared. And I thought, huh. Let's just do newly vulnerable. But the Lamphone, you know, these guys always wrestle this stuff right down to the ground. It's one thing to say, oh, I wonder if a light bulb vibrates. But they've got the theta and phi spherical coordinates mapped out. And like they just - anyway, we're going to have a lot of fun with that.

But we also have an accident for which the CEO of the Brave browser apologizes with regret. We're going to take a look at last week's Patch Tuesday and its myriad ramifications and unintended consequences. Oh, boy. And of course they broke another record, that is, Microsoft did. We also note a few odd new and unwelcome behaviors from this year's 2004 Win10 feature update. Anyone care to print after that?

We're going to dip into yet another side channel attack on Intel chips, but also note the final arrival of a long-awaited powerful antimalware technology which Intel is about to start shipping. We're going to look at yet another new SMB vulnerability named SMBleed. And then we will conclude with an examination of this latest, what happens if you look closely at a light bulb research. So, yeah, I think we're all going to have fun. And a really interesting Picture of the Week for our more savvy, technically inclined listeners.

Leo: For the geeks amongst you.

Steve: Yes.

Leo: I would guess everybody who listens to this show is at least a little bit geeky.

Steve: I think so.

Leo: I'm thinking. Steve, I can't wait for you to explain what I'm seeing here, this Picture of the Week.

Steve: So you're talking about our Picture of the Week. The caption says: "In 1969 Apollo 11, the spaceflight that first landed humans on the Moon, used just 30 simple instructions to calculate the transcendental functions like sine and cosine essential for navigation." This made the rounds of social media a couple weeks ago, and I just grabbed it because I thought it was so clever.

Transcendental functions are historically very difficult and time consuming to calculate. I mean, even on modern machines. Of course now we've got the math coprocessor. And one way to just sort of sidestep the whole problem would have been just to use a big table. I mean, once upon a time, remember, what was it, the CRC that we had back in our college days.

Leo: Oh, yeah. You wouldn't calculate it. You'd do a lookup.

Steve: You would often just look up a sine or cosine in a book that had the tables printed because they were just - I was like, that's the way you found out the value.

Leo: That's how early Lempel-Ziv compression worked. It was just table-based because a lookup's table so fast, you know. You're not doing the work.

Steve: Well, yes. It's fast. But the problem is it's also big. And what we didn't have in 1969 for the Apollo 11 was memory. There just, you know, that still was rare. So what some brilliant mathematician, and I forgot her name, but it was a woman who was in charge of this aspect of the Apollo 11...

Leo: Was it Margaret Hamilton?

Steve: That sounds right.

Leo: She wrote a lot of the Apollo 11 code, yeah.

Steve: That sounds right. I think it was Margaret. She realized that you could, for the region, for the important region of a sine - and of course a sine and a cosign are just 90 degrees skewed, so if you have one, you automatically get the other. For this important range it was possible to use a relatively simple polynomial that tracks it through that range.

And for anyone who's interested, this is in the first page of today's show notes. It shows the actual function of sine as it sines itself along on a blue curve. And then this polynomial approximation, which comes out of nowhere, and then just in the nick of time settles right onto the sine wave and follows it virtually perfectly. And when I say "virtually," I mean, you know, we did get to the Moon. We didn't overshoot or undershoot.

Leo: Yeah.

Steve: You know? We landed. So it's just a beautiful piece of work. And then this also shows the polynomial which is, while it's not simple, it turns out it decomposes into a number of divisions and multiplications and some factorials which end up being pretty simple. So, and that ends up being expressible in just 30 instructions. So just a very nice piece of work.

Leo: So cool. So cool.

Steve: Really, really nice.

Leo: Yeah, just really neat. Wow.

Steve: Okay. So not so neat, although kind of a tempest in a teapot...

Leo: Yeah, I think not as bad as people made it out to be, yeah.

Steve: Yeah. Cryptonator 1337 - and of course 1337, as we know, is LEET sort of upside down, he or she was the first to note this behavior of the Brave browser, which they found objectionable, tweeting: "So when you are using the @brave browser and type in 'binance.us,' you end up getting redirected to 'binance.us/en?ref=35089877.' I see what you did there, mates."

So okay. So Binance, as in kind of a clever take on "finance," is a cryptocurrency exchange with whom Brave has a fiduciary relationship. They have Binance on their widgets on their new tab page, hoping that people will click on it, and that they'll get a little bit of referral money in order to finance the Brave browser project.

So Brendan Eich, the Brave CEO and cofounder, whose Twitter bio also notes that he cofounded Mozilla and Firefox and created JavaScript, so he's been around, he tweeted: "We made a mistake. We're correcting." He said: "Brave default autocompletes verbatim in address bar to add an affiliate code." He says: "We are a Binance affiliate. We refer users via the opt-in trading widget on the new tab page. But autocomplete should not add any code."

Okay. So then he goes on for another six tweets explaining himself and so forth. To which someone replied: "It's not a mistake. You did it on purpose." Which Brendan replied: "I think you used 'mistake' where you meant 'accident.'" He said: "I never said it was accidental." He says: "We were treating it like a search query, which all big browsers do, tag with an affiliate identifier to get paid by the search provider." He said: "But a valid domain name is not a search query. Fixing."

So anyway, essentially they decided, I mean, Brendan agreed that, if you enter something into the URL of a browser, you don't expect it to get changed with an affiliate tag added to it. So it no longer does. It wasn't nefarious. It wouldn't have altered the browser's cookie exchange in any way. So it wouldn't have had any effect on user tracking. Really the browser...

Leo: And it didn't say any information about the user to Binance, either, any more than they would normally get.

Steve: Correct.

Leo: Didn't do that.

Steve: Correct.

Leo: It just said, hey, we sent you. We'd like some money. Like a little, give me some bitcoin.

Steve: Yeah, exactly. And when you think about it, the browser's user-agent header could have identified the visitor coming into Binance as a Brave-sourced user. But I presume that Binance is only equipped to accept affiliate tags through the search query or through the query URL that arrives to them. So that's the way Brave did it. Brave is now a full Chromium-based implementation, which makes sense, since maintaining a secure browser, no matter how good you are, and keeping it state of the art is an impossible burden these days.

We would prefer, yes, not to have a monoculture in browsers. But a modern web browser has become like an operating system. It's just no longer something that can be built up from scratch. And if someone said, yeah, but I can do it, the proper response would be, why bother? We already have browsers, and they're fine. So this is just not a wheel that you can reinvent. But anyway, I sort of saw that in passing, and I thought, okay, well...

Leo: Yeah, there was a lot made of it. I feel like all browsers - look, Firefox doesn't do that, but Firefox lives - because these are free; right? It lives on the money it gets from Google referrals. And I just think that that's - I think Brendan, who knew that, having founded Mozilla, just kind of didn't think about it. So when he says "mistake," I don't think it was like a programming error. It was a mistake in judgment. And I think that that's fair.

Steve: Right. And that's what he meant. It wasn't an accident. It didn't happen by mistake or accidentally. They just thought, hey, that's one other way that somebody might go to Binance, so let's jump on its coattails and get some credit for that, too. And really, you know, they didn't have to show it. I mean, this got observed, but they could have done this surreptitiously, if they wanted to be sneaky, and they didn't bother to. So it's just like, okay, yeah, you're right, that's maybe pushing affiliate tagging a bit too far.

Leo: Yeah, I think they did the right thing. They took it out. And I do believe there was a switch. You could turn it off, as well, in the settings.

Steve: Yes, there is. Somewhere in the config options you can disable it. But again, we know that most people are not going to do that.

So Microsoft is, if nothing else, consistent. They are continuing their record-breaking streak, or as Sophos put it: "Whoosh. You hear that? It's the sound of Microsoft's security fire hose spraying out a river of CVE fixes."

Leo: Oh, man.

Steve: "That's right," they wrote, "Patch Tuesday was last week," the software giant's largest yet, releasing, yes, 129 fixes for CVEs. In other words, once again Microsoft has broken, well, set and broken its all-time record for the most patches released in one month. Do you remember, Leo, we used to talk about 11.

Leo: Yeah, yeah.

Steve: We were like, wow, there were 11. Okay.

Leo: Is this good or bad? I mean, they are getting found and fixed.

Steve: Yeah, but they're also, as we'll see shortly, they're also creating more problems than they're fixing. I mean, they really are.

Leo: Yeah. If they're introducing them, it's bad.

Steve: Yeah, really, they really are breaking things. So, okay. So just quickly, most of the 129 are rated just "important." Eleven of those 129 CVEs are CRITICAL, in all caps. And they're remote code execution vulnerabilities. It's pretty much the case now that if it's not remote code execution, it's like a yawn. Although, you know, privilege elevation can be, as we've seen, a problem, too. But anyway. So these are in Windows 10. They're all CVE-2020 somethings.

And so here we have 1286, a Windows shell remote code execution triggered by improper file path validation. There that is again. We have 1299, a remote code execution bug, again, that an attacker could exploit using a malicious .lnk, that is, a link file. And note that either we still haven't got link files working right, or we keep breaking them, since Windows has been having security problems with link files from Windows 95. And in this case, Microsoft warns us that if a malicious link file was placed onto a removable drive or a network share, clicking on the link file would run the attacker's malicious remote execution code contained in the file, or remotely provided, that they consider an RCE bug.

Then we have 1281, a vulnerability in Windows Object Linking and Embedded, OLE code, stemming from poor input validation. And it's exploitable via a malicious website, a file, or an email message. 1248, a memory object handling bug in GDI, the Windows Graphics Device Interface, which is deliverable via website, an instant message, or a document file. Those all affected Win10, of course, since Windows 7 is no longer being maintained, and many of those also affected the latest 2004 build of Windows 10 since, of course, most of the code never changes.

Not to be forgotten, IE had its own batch of critical vulnerability bungles. Both IE9 and 11 were susceptible to a remote code execution via bug, or actually three bugs - 1213,

1216, and 1260 - all memory handling errors affecting Visual Basic Script, VBScript. The original web browser - which I guess still isn't quite history yet, although as we talked about last week they're now working to replace it with Edge and then irrevocably replace it since you can't get rid of it once it lands.

The original Edge browser had a critical vulnerability, 1073, a memory-handling bug in its ChakraCore JavaScript engine. And then there's 1219, which affects both IE and Edge HTML, with more memory-handling issues. And finally, 1181, a bug in the SharePoint Server. It can be exploited by unsafe ASP.NET controls that don't filter properly. And apparently attackers who are able to upload a malicious page to the server, it's not clear how they would do that, perhaps through remote website authoring, could achieve pwnage. As a consequence, admins of SharePoint Enterprise Server 2016, Foundation 2010 SP2 and 2013 SP1, or SharePoint Server 2019 should all patch now, or should have.

And what's interesting is that those numbers are so low. This suggests that these were reported to Microsoft early in 2020. I mean, we're mid-June now, and I don't know where the CVEs are, but they're way above 1200. And so Microsoft has taken a while to actually get these out to the world. Oh, there's also 1300, a longstanding bug in Windows' handling of cabinet files. It affects most versions of Windows, including 7 through 2004, you know, Windows 10, and Windows Server. And believe it or not, those are just the 11 critical bugs. If I were to attempt to detail the other 118 important flaws, this entire podcast would have to be retitled "What Happened Last Tuesday." So I'm going to spare us that since we have plenty more to talk about. But in the meantime, Microsoft...

Leo: I do think the key here is most of those are old. So to my mind, if Microsoft's introducing new flaws, problematic. Right? But if they're fixing - the rate might be going up because they're doing a better job of finding and fixing old flaws, that's a good thing.

Steve: That would be a good thing. Which, thank you, Leo, for the segue, because we also have...

Leo: This is a new one.

Steve: ...Microsoft's disclosure of an oddball Win10 delight.

Leo: Weird. Yeah, unique to Windows 10.

Steve: Titled, I'm not kidding you, "USB printer port missing after disconnecting printer while Windows 10" - and this is versions 1903 or later - "is shut down." And it started - okay. And they stated this applies to Windows 10 1903, all editions; Windows 10 1909, all editions; and Windows 10, yes, 2004, all editions.

So Microsoft explains: "If you connect a USB printer to Windows 10 version 1903 or later, then shut down Windows and disconnect or shut off the printer. When you next start Windows, the USB printer port will not be available in the list of printer ports. Windows will not be able to complete any task that requires that port." For example, printing.

So they said, under "Resolution: You can avoid the issue by connecting a powered-on USB printer before starting Windows." Because of course this is the 21st Century.

"Microsoft has confirmed," they wrote, "that this is a problem in the Microsoft products that are listed in the 'Applies to' section." And that was all Windows from 1903 on. "We are working to fix the issue in a future version of the operating system." What?

Okay. So according to reporting of this in the tech press, if you need to print something to your previously USB-connected printer, and you didn't have it turned on before you started Windows, no problem. Just shut down your computer, turn the printer on, and wait for it to finish initializing and settle down. Then you can fire up Windows, and the printer port would reappear, and you'll be able to print. With Windows. Because, you know, this is a state-of-the-art modern operating system.

And believe it or not, in a related but separate matter, last week's Patch Tuesday broke all printing, even to PDFs, for many users. Windows 10 users are reporting that they are unable to print to printers from several vendors after installing last week's cumulative Windows 10 patches for 1903, 1909, and 2004 OSes. Two specific patches causing troubles have been determined to be the cumulative updates, and these are KB4560960 or KB4557957. Although Microsoft hasn't yet gone official, a Microsoft Answers Independent Community Advisor has stated that Microsoft engineers are "already aware of this issue and working a patch to be deployable in the next update." So no printing for a month.

After updating their machines last Tuesday, users began flooding both Microsoft Answers forums and Reddit with reports of printing issues affecting various models of HP, Canon, Panasonic, Brother, and Ricoh devices. Typical postings included, for example, here's: "Unable to print after installing update KB4560960 and/or KB4561608. Uninstalling updates fixes the problem. This is happening to every Windows 10 computer in our organization as updates install."

Another one says, right after installing KB4560960 on multiple systems, users started reporting: "Windows cannot print due to a problem with the current printer setup," errors that went away after uninstalling the update. Someone else wrote: "Found this problem today where all clients at a customer site had the same problem. They have Ricoh, but a few other brands, too. Even the virtual PDF printers do not work anymore. Explorer.exe crashes completely when doing a test print."

A network tech posted: "HPs seem to be hit or miss with this issue. Ricoh, Canon, Brother, KM, Kyocera all seem to be experiencing problems. As everyone else is saying, backing out update KB4560960 and postponing updates seems to be our only salvation at this point." And then somebody else: "Hopefully Microsoft will produce a patch for this quickly. Call volume is picking up with everyone returning to work. This is going to make things awfully hectic." So, yes, new breakages.

Affected users have found that the printer's native driver can be replaced with PCL6 drivers, which reportedly work, or by uninstalling last week's cumulative updates to restore printing, and of course to also restore those 11 critical remote code execution bugs. You'll be fine. It's been determined that attempting to uninstall and reinstall the printer, or updating its drivers, does not help. PCL6 printer drivers do work, either vendor-specific PCL6 drivers or the universal Windows 10 PCL6 drivers for Canon, HP, Ricoh, Kyocera, and Brother. So, oh, brother.

Leo: Really amazing.

Steve: And Leo...

Leo: So that's an example of you don't want to see that. You don't want to see them introducing - especially stupid bugs. What do you think that's being caused by?

Steve: I have no idea. I mean, really.

Leo: It's so weird.

Steve: Like what could it be? And that's not all.

Leo: There's more.

Steve: Those running Windows 10 - yes, and wait, there's more. Those running Windows 10 who have moved this time only to 2004, who have systems based on SSDs, the 2004 feature update has broken Windows' awareness that it has ever previously defragmented the system drive. As a result, rather than only defragging occasionally - like once a month by design, which is known to improve the performance of Windows "volume shadow copy on write" system - now, if you have 2004, Windows 10 is defragging over and over and over, every time the system is started.

Now, it's not a huge problem since SSDs should have strong write endurance. But unfortunately it's not something you want to tax, and it's not what anyone wants. Microsoft has acknowledged the problem, but has not indicated when it will be resolved. The release notes for the insider preview build 19551 state: "Thank you for reporting that the Optimize Drives Control Panel was incorrectly showing that optimization hadn't run on some devices. We've fixed it in this build."

Now, that makes it unclear whether it doesn't show it or that it isn't acknowledging it. The screenshots I saw showed it running, suggesting that it is actually, if you open it up after you restart Windows, you'll see like, oh, yeah, we're redefragging your SSD, which we did for you yesterday, when you last turned Windows on, because we forgot we did that. So anyway, insider preview build 19551 fixes that. Unclear when other 2004 users are going to get it.

And in another oddity, Windows 10 2004 is for some reason also attempting to use the TRIM command against non-SSD drives. That of course fails and logs an error into the Windows error log because spinning drives don't support the TRIM command. Windows is not supposed to be trying to get them to TRIM themselves. Our long-time listeners will recall that SSDs have a TRIM command to allow the operating system to inform the drive of those regions that are not in use by the OS.

Normally drives, you know, traditional hard drives treat all sectors alike, and only the OS has any awareness of which regions are in use by its file system and which are free. Hard drives write data by simply overwriting what was there before. They just plow right over it. But SSDs are only able to set bits that have been previously reset by an erase cycle. And those are NAND-based SSDs. And erase cycles only erase large blocks of the SSD all at once. This means that to write a small region of a larger block, the previous contents of that larger block must first be read and held in a cache while the underlying physical SSD block is all reset. Then that cache data must be rewritten into the block.

But if the SSD has an awareness, thanks to the TRIM command, having been previously used, of which sectors are not needed, it can leave them reset, in the reset state, rather than rewriting them needlessly with junk data. And those that are reset and have not

been written to can later be directly written to because they've been left reset without needing that whole pre-erase cycle. So this TRIM command is a win. You definitely want it on SSDs. And of course our modern OSes all have that now. But it makes no sense for hard drives. So again, Windows 10 2004 introduced another new bug which is causing it to issue superfluous TRIM commands to spinning hard drives. Doesn't hurt anything, but it just sort of says, you know, Microsoft...

Oh, and there were also many reports that programs would no longer run at all after last Tuesday's updates. But I followed that down. It turned out to be a problem caused by an interaction, yes, with a recent update for Avast or AVG antimalware. There is, in the registry, there's the ability to support a debugger, that is, essentially it's a hook which allows, when you tell Windows to run one program, it actually - it goes to the registry and checks this little region. And it's possible for a user to have a debugger run instead, and then the debugger runs that other program, essentially a way of getting a debugger into programs in order to debug its startup.

And it makes me a little queasy that this is what Avast and AVG are doing, that is, I mean, they're deeply hooking the system such that, when they broke that feature, all the commands that they'd been hooking in this way, all the executables, stopped working. So they produced an update quickly, and this problem got fixed. Somehow it interacted with Tuesday's update. And we've seen instances before now where antiviral software, which is getting its hooks deeply into our OSes, are beginning to collide with what it is that Windows is doing.

So my take is, much as Win10 2004 has some interesting useful new features, it really does feel as though perhaps holding off a bit and waiting for things to settle down might be prudent. I have an Intel NUC running Windows 10, and whatever it is that is new about 2004, Microsoft is saying, no, we're not ready yet on your hardware. And I'm just saying, good. I'm not ready yet, either.

Leo: Yeah. It's holding off installing it on places where they think there might be an issue. So who knows.

Steve: Yeah. And in fact, Leo, that would be everywhere.

Leo: Everywhere, yeah, all of them.

Steve: Yeah.

Leo: Yeah.

Steve: If you have a printer, don't install 2004.

Leo: Don't. Geez, Louise. Oy oy oy.

Steve: If your system has a hard drive or SSD, no, don't install 2004.

Leo: Yeah, no, unh-unh. Yeah, bad idea.

Steve: Floppy-based Windows, Win10 on floppies, not a problem. We have that nailed.

Leo: Where can I get Windows 10 on floppies? Oh, man.

Steve: So another SMB problem. We've got a new information leakage vulnerability which was introduced in Windows 10 1903, and it's present in all releases since. With Windows 10 1903, Microsoft's very troubled SMB, you know, Server Message Blocks, a.k.a. file and printer sharing, has been around also forever. But they can't leave it alone because it was updated, or upgraded, to support optional compression, which is present in SMB v3.1.1.

And the trouble is this allows a new class of information probing attacks to succeed, thanks to the foothold this provides an unauthenticated attacker. It's loosely related to the SMB Ghost vulnerability that so alarmed Microsoft earlier this year when they warned everyone that it could be turned into a wormable exploit. And as we know, that never happened, but it was successfully used against selected targets.

Microsoft Security Advisory published last week stated that "an attacker who successfully exploited the vulnerability could obtain information to further compromise the user's system." And they also said: "To exploit the vulnerability against a server, an unauthenticated attacker could send a specially crafted packet to a targeted SMBv3 server." In other words, an open file and printer sharing port. "To exploit the vulnerability against a client," they said, "an unauthenticated attacker would need to configure a malicious SMBv3 server and convince a user to connect to it." That seems much more far-fetched to me.

So the big threat is against servers, as usual. So the problem will mostly affect those devices, Windows 10 servers; and the core servers are affected, as well, since they proudly support SMBv3.1.1. And at this point anybody who has an SMB protocol service open to the public has few excuses when something uses that port to crawl into their network. It's not only an open port. It's an open invitation if it's Server Message Blocks, SMB.

The ZecOps guys who are the ones who disclosed and detailed their findings have been taking some heat for also simultaneously releasing a pair of proof-of-concept demos. One was an uninitialized kernel memory read proof of concept that creates a local file containing the target computer's kernel memory. Whoops. Like we know what's in the kernel: all of the keys, all of the cryptographic keys which are currently in use. They also posted a pre-Auth remote code execution proof of concept combining SMBleed with SMBGhost that opens a reverse shell with full system access. What could possibly go wrong?

So when asked why they didn't wait for Windows users to first patch their systems, because this all just happened, and chose to publish their proof of concepts when SMBleed was disclosed, the ZecOps guys said that the security vulnerability was not critical on its own, that is, all by itself. They argued that "After the patch was made available, the vulnerability was so easy to spot and reproduce," and "Only when combined in combination with another primitive, such as SMBGhost, would SMBleed be critical." And since SMBGhost was patched three months ago, they felt that people should be safe.

So the takeaway for our listeners, if your company still has SMB exposed to the public, really, really invest in a VPN solution having multifactor authentication. Since this attack works against unauthenticated attackers, not even multifactor SMB authentication would

protect your servers. And as I said, changing to some other port than 443 - or, wait, no, 445 - is no longer enough. You just can't move it somewhere because, for example, Shodan now looks at all the ports and figures out what service is running there. So, yeah, you're easy to find. Just changing to an obscure port, no. You may think that's tricky. No, not any longer.

Now, I'm sure that no one listening to this podcast would fall for this. But just for the record, this fraudulent website extortion scam has been in the tech press news recently because those behind it are, if nothing else, persistent. And because the email was apparently written by someone with at least somewhat passable English, which is a little bit unusual for some of this spam email, it's come to people's attention.

The email reads: "Subject: Your Site Has Been Hacked." And, you know, as a person who has a site, I thought, okay, what? It turns out that people with no websites get this, too. So the guys aren't even being that clever. But the subject of the email comes in, "Your Site Has Been Hacked." It's like, oh, okay. Then in all caps: "PLEASE FORWARD THIS EMAIL TO SOMEONE IN YOUR COMPANY WHO IS ALLOWED TO MAKE IMPORTANT DECISIONS." So I guess that's meant to be, yeah, we know you may not actually be in charge of the site, but your company's site has been hacked. So they say: "We have hacked your website," and then they say "URL redacted," which is, okay, a little [crosstalk].

Leo: Because we don't know it.

Steve: We don't actually, yeah, know the URL of your site.

Leo: Just, you know, fill it in yourself.

Steve: That's right. You know what your site is; right? Yeah. We're not going to bother telling you. And, they say, "...and extracted your databases."

Leo: Ooh.

Steve: And they say: "How did this happen? Our team has found a vulnerability within your site that we were able to exploit. After finding the vulnerability, we were able to get your database credentials and extract your entire database and move the information to an offshore server." Of course we're not sure which shore that is, but it's offshore.

"What does this mean?" they say. "We will systematically go through a series of steps of totally damaging your reputation. First your database will be leaked or sold to the highest bidder, which they will use with whatever their intentions are. Next, if there are emails found, they will be emailed that their information has been has been sold or leaked, and your site" - and again, brackets, website URL - "was at fault, thusly damaging your reputation and having angry customers/associates with whatever angry customers/associates do."

Leo: What the what?

Steve: "Lastly, any links that you have indexed in the search engines will be de-indexed based off of black hat techniques that we used in the past to de-index our targets."

Leo: Oh, yeah, those black hat techniques, whoa, boy.

Steve: They've got experience. Yeah. "How do you stop this?"

Leo: I don't know.

Steve: "We are willing to refrain from destroying your site's reputation for a small fee."

Leo: Oh, that's fair.

Steve: "The current fee is," and then "[ransom amount] USD in bitcoins." And they say: "If you decide not to pay, we will start the attack at the indicated date and uphold it until you do. There's no countermeasure to this. You will only end up wasting more money trying to find a solution. We will completely destroy your reputation amongst Google and your customers."

Leo: Oh, boy.

Steve: Yeah. So the ransom demands generally range between 1,500 and $3,000 equivalent in bitcoin. And of course this is all nonsense. However, I have received this email several times myself at the email address registered actually for my one remaining domain at Network Solutions. It's just not worth moving it. Unlike Hover, where I have moved everything that's movable, Network Solutions does not offer registered email blinding along with their service. And I refuse to pay Network Solutions any extra because it's like every year $5 or something. It's like, no. And the domains that are still there are just - they're prepaid, long-duration domains I bought, you know, they're just sticky.

But you know, Leo, security experts are nothing if not curious. And since what's been called "Breachstortion" - this campaign has been named that by the tech press, the "Breachstortion" campaign.

Leo: That's not a good word.

Steve: No. It lists the extortionist's pay-to bitcoin wallet addresses. And since it's possible to monitor the bitcoin blockchain for payments made to wallets, the web security firm WebARX decided to see how successful this campaign might be. By scouring social media postings for references made to this campaign, and there were many, WebARX identified the multiple bitcoin wallets being used to collect whatever ransom payments might be made.

In their coverage, WebARX wrote: "Unfortunately, looking at the different bitcoin wallets linked to these attacks, there have been at least five people who have fallen to the scam and paid ransom. One of the wallets linked to the attack has received close to $2,000

worth of payments in bitcoin. Another wallet used in this scam which has not yet received payments has been reported for abuse 81 times." So whoever is paying these are apparently not the brightest bulbs in the box since the threat and payment demand contains no email, no website contact details, is apparently just a blind spam. Which means that the scammers tell the recipient not to bother replying to the email at all. And there's no website where they're able to trace the payment to see whether they've received the money.

They explain to their intended victim that "Bitcoin is anonymous, and no one will find out that you have complied." So presumably that's meant to put the target's mind at rest by convincing them that the act of payment will not itself draw attention to the of-course-it's-fake breach, even though it means that you're relying entirely on the crooks then to keep track of which payments were made to protect which website's data. Which of course they're not. Unless they're sending, like, separate bitcoin addresses to each email address, which they're not. Or tracking them by using custom per-demand amounts, which seems highly unlikely. In other words, they have no idea if they receive money, obviously...

Leo: Thanks for the donation.

Steve: ...who they got it from. Yeah, it's just a donation to their spamming campaign. So this is just a blind spam gamble, hoping that it will pay off. But as one of P.T. Barnum's more well-known quotes observed: "There's a sucker born every minute." On the other hand, remember that SophosLabs similarly reported on the apparent success of those porn-scamming solicitations. You know, the one that claims to have activated your webcam and to have recorded what you were doing while you were watching online porn.

It's a bit unnerving the first time one of those arrives. I've had a few friends who have received that spam and asked me whether it's possible. Apparently they were a little worried. Sophos did some similar digging into the blockchain and discovered that that porn-spamming campaign was generating on the order of $100,000 a month for those guys.

Leo: Yeah. I'm not surprised, yeah.

Steve: So, yeah. Enough people, like, say, oh...

Leo: Why take a chance? Maybe it really did happen. I'm going to give them money.

Steve: Maybe it did happen because, you know...

Leo: I'll tell you, if they made it easier to pay, I bet they'd make more money.

Steve: I know.

Leo: It's complicated.

Steve: I bet you're absolutely right. The fact that it's bitcoin is a huge hurdle.

Leo: In fact, the email, because I've got it many, many times...

Steve: Oh, yeah, yeah.

Leo: ...goes through a bunch of steps that you have to - so if you want to know more - we talked about it a while ago. If you want to know more about bitcoin, read this article. Now, here's how you create a bitcoin. Blah blah blah blah blah. It's too much trouble. Most people are just going to give up.

Steve: Yeah, just watch some porn instead.

Leo: Go watch some porn.

Steve: So what is real is that unfortunately there are some authentic database ransom attacks. Insecure SQL, you know, SQL servers exist for online merchants, and there have been multiple accounts of their databases being copied and ransom notes left behind demanding payment or else. And being authentic, those attacks have been far more successful in collecting ransom. Researchers have tracked a total of 5.8 bitcoin, currently worth around $54,000 based on current bitcoin valuation, having been sent to the attacker's address by over a hundred victims into just two attacker wallets.

In the past, similar database ransom attacks have targeted MongoDB and also MySQL servers. So again, please, please, please never leave SQL servers exposed to the public Internet. Not even on, as with SMB, nonstandard ports. As I have said before, changing a port is no longer a sufficient protection. You know, if you have to do it, then use IP address filtering. That is, if someone remote has to connect to your system, they will almost certainly have a relatively static IP. So put up an IP filter and only allow incoming packets from that IP address. It's not perfect, but it's really, really strong protection because only someone at that IP is able to see that there is a port open. Better to use a VPN, obviously. But if you can't, certainly you ought to have a firewall at your boundary. So restrict access to anybody not at that IP.

And we have another side-channel attack on Intel chips. As we know all too well, the past two years have been a watershed period for security researchers poking into the seemingly unlimited supply of new vulnerabilities which have arisen as a consequence of our processor architectures, predominantly Intel's, because they cleverly attempted to squeeze every last microcycle of possible computational power from our chips. For that we thank them.

We don't thank them for the fact that, unfortunately, as seemingly benign as, for example, a cache is - which is used to decouple the demands of the fast processor from the comparatively lethargic reluctance of DRAM to give up its data - even a cache can, as we've learned, have its contents probed by other processes sharing the same hardware. And in today's cloud world, where you do have many different users sharing the same hardware, now you've got a problem. So throughout 2018, as we know, Intel's engineers learned to just what degree their clever engineering, meant to optimize the chip's performance, could be used by malicious code to exfiltrate data across virtually every security boundary that they had deliberately erected. They were all porous.

So today no one would be surprised to learn that yet another fault has been discovered and leveraged by researchers to penetrate Intel's SGX. That's their Software Guard Extension, which is their version of a Secure Enclave. Unfortunately, not as secure as they were promising. Last week two separate academic teams discovered two new distinctive exploits that do exactly that. They pierce Intel's Software Guard Extensions, which is designed to be by far the most secure region that Intel processors create.

And of course these new attacks aren't the first to soften the SGX security wall. We've talked about several before. Back in 2018 a different team of researchers used the attack, now known as Meltdown. And still another team broke SGX in a different way earlier this year. So things are not looking good for Intel security at the moment. And as we know, we talked about it at the time, Intel mitigated the earlier SGX vulnerabilities by introducing updates to their microcode. But they were insufficient.

So Intel has again, last week, released yet again new updates which should be available to end users before long. As we know, Intel's microcode can be patched at boot time, and either the motherboard's startup code or our OS's boot is able to apply those patches on the fly. So eventually we'll get them. I imagine Linux will have them quickly. Microsoft will push them into current versions of Windows. No longer supported versions of Windows will never get them. And maybe motherboard BIOSes will get them for those motherboards which are still being maintained. So yet another problem.

The most recent attacks are known as SGAxe and CrossTalk. Both use new and different side-channel attacks to infer what's going on within this walled-off region of Intel's processors. I've got more stuff in the show notes, but everyone gets the ideas now. That's all bad news.

I do have some good news from Intel. Control-flow enforcement technology, which we're going to be talking about in the future, CET, control-flow enforcement technology is finally here. Intel just announced that its long-awaited hardware control-flow enforcement technology is ready and will be included first in their next Tiger Lake mobile CPUs. And they did also say that desktop and server platforms based on the Tiger Lake architecture would be getting them.

I talked about CET a while back. Recall that it's the technology that maintains a completely separate shadow hardware stack which, unlike the processor's traditional hybrid stack which mixes together subroutine return addresses with subroutine calling parameters and dynamically allocated data and buffers, the CET stack only contains the control-flow data, namely subroutine returns, and also indirect branch data.

Since this data is maintained by the hardware and is not visible in any way to software, unlike the software-visible stack, it's not subject to malicious manipulation. So when any return-from-subroutine instruction is encountered, this new Intel hardware will compare the return address it previously stored in its hardware to the return address that is being executed on the software stack. And if they're not in agreement, the executing thread will be killed.

This is a massive win for Intel processor security because in one fell swoop it eliminates a broad class of attacks. I mean, this is not just we're trying to protect our processors from information leakage. This is overt antimalware technology. And all of those surprisingly effective Return-Oriented Programming, those ROP attacks that the previous randomization of operating system, kernel, and program address space - remember ASLR and KASLR - they were attempting to mitigate these return-oriented programming attacks by randomizing where things were so that attackers who were blind to where things were loaded wouldn't be able to as easily craft return-oriented programming attacks which had otherwise proven so effective.

Since ASLR can only place software modules in a finite and relatively small number of places, like one of 256 places, we've seen that some attackers are fine with just guessing and crashing if they're wrong because some percentage of the time, like around 4% in the case of a one-in-256 chance, they will succeed. And when it's a numbers game, like in an attack that's just being sprayed out, that might be good enough for them. But in the presence of a hardware shadow stack, those attacks all fail, once and for all.

So I'm really excited about this, thanks to the fact that Intel has been working on this for the past four years, since 2016, when they first published their CET specification. And as we know, Intel has a relatively long processor development cycle. They've got a bunch of processors in the pipeline. And so it's taken four years for new architecture that incorporated CET in the microarchitecture to make it like to the point where it's ready to ship.

But the other good news is, since it took four years, and since Intel was able to give everybody a heads-up, software publishers have had time to get ready and add the required support to their code, waiting for the day when this next-generation active antimalware technology would finally ship. It's already present in the GNU standard C library, Glibc. And Microsoft has also added CET support to Windows Insiders, calling their support for this feature Hardware-Enforced Stack Protection.

So once we actually have the chips, the apps and operating systems will be able to enable their support and opt into the truly significant protection that CET provides. And as I noted, CET will first appear in Intel's Tiger Lake architecture mobile chips, apparently like now, and then it will also be available for desktop and server platforms.

So Leo, I think this is cool. This is like basically we've had no other hardware enforcement of these problems. We've had things like, remember, like stack canaries, where little tokens would be placed on a stack, and the subroutine return would check for the canary before it would accept it. And since bad guys, that was like something bad guys wouldn't be able to put on the stack necessarily, although, yeah, canaries have been bypassed, too. All the other things that have been done have been software-based solutions. Here we have a hardware-based solution that just ends this.

Now, of course, it does require that it be turned on. It requires that the hardware be able to support it. So unfortunately it's not applicable at all for any of the pre-Tiger Lake hardware. Thus it's going to take some time for this to filter into all of the hardware that we have, and then for everything just to be able to assume it's there, like we do now with things like math coprocessors that we didn't used to have. But still, this is nice. It means basically that address space layout randomization and kernel address space layout randomization, when CET is present and can be enabled, could be turned off if there was any cost or overhead for doing it because this is like the right way to end that kind of abuse, and anything else that goes wrong. So very, very cool.

Leo: Neat. All right, Steve. Now we will find out about...

Steve: As I said at the top of the show, I was originally going to title this week's podcast "Newly Vulnerable," since we talked about a bunch of new vulnerabilities.

Leo: Yeah.

Steve: And then put them all at the end. And I was going to place the Lamphone under "This Week in Wacky Surveillance Technology" category. But as they always do, the

researchers have a quite serious 15-page research paper, which always hooks me. And their work will be presented at Black Hat in August, the virtual Black Hat conference. So I decided that it won the top spot for this week's title and topic. And no one will be surprised to learn that this research hails from the Ben-Gurion University of the Negev, and the Weizmann Institute of Science, which produces a steady stream of new and imaginative data exfiltration schemes. Remember the singing capacitors of our power supplies. That was from those guys, too.

And what makes their work stand out is that they always wrestle every wacko scheme to the ground, applying solid science and physics to the challenge. In the show notes I have a schematic of their attack, which in the left-hand frame shows a couple of people talking, and that creating sound vibrations in the air, which apparently causes some microvibrations in the surface of a light bulb, like hanging from the - just a regular light bulb hanging from the ceiling. Then in their experiment up to I think it was 27 yards - maybe it was more than that. We'll get to it in a second. I mean, at some good distance away, they were able to use a telescope looking at the light bulb and a photodiode-based sensor to detect the microvibrations that the light bulb was subjected to just from regular conversation in the room.

Now, the light bulb, of course, is far from being a high-quality microphone. It's anything but. But they were able to work out the details. So this time I'm going to skip the abstract of their paper, which I normally share. Instead, I'm going to share their introduction, since this really does, it turns out, represent an advance in the state of the art of eavesdropping attacks, thanks to the backend computational transform that's needed to convert a light bulb's highly nonlinear frequency response into reconstructed speech.

If you think about a tuning fork, as we all know, you thwack it, and it vibrates at its natural resonance frequency. So in terms of speech, it's an extremely sharp band-pass, a very narrow band-pass filter. And a light bulb is not much better. So think about like a wine glass. You can ting the side; right? And if it's crystal, it'll vibrate, again at its natural resonance frequency. So rather than directly obtaining the speech waveform, what you get is more of a speech frequency modulated amplitude modulation. So they had to do some work in order to extract speech.

So anyway, they explain this. They said: "Introduction. Eavesdropping, the act of secretly or stealthily listening to a target or victim without his or her consent by analyzing the side effects of sound waves on nearby objects, for example, a bag of chips" - remember we talked about that years ago. You're looking at the vibration of a bag of chips, and we're able to extract the speech from it - "and devices, motion sensors, for example. It's considered a great threat to privacy."

They said: "In the past five years, various studies have demonstrated novel side-channel attacks that can be applied to eavesdrop via compromised devices planted in physical proximity of a target or victim. In these studies, data from devices that are not intended to serve as microphones - for example, motion sensors, speakers, vibration devices, and magnetic hard disk drives - are used by attackers to recover sound.

"Sound eavesdropping, based on the methods suggested in the abovementioned studies is very hard to detect because applications and programs that implement such methods do not require any risky permissions, such as obtaining data from a video camera or a microphone. As a result, such applications do not raise any suspicion from the user or operating system regarding their real use, eavesdropping. However, such methods require the eavesdropper to compromise a device located in proximity to the target victim in order to obtain data that can be used to recover sound and exfiltrate the raw processed data."

They said: "To prevent eavesdroppers from implementing any of the above-mentioned methods which rely on compromised devices, organizations deploy various mechanisms to secure their networks. For example, air-gapping the networks, prohibiting the use of vulnerable devices in sensitive areas, using firewalls and intrusion detection systems to prevent exfiltration. As a result, eavesdroppers typically utilize three well-known methods that don't rely on a compromised device.

"The first method exploits radio signals sent from a victim's room to recover sound. This is done using a network interface card that captures WiFi packets sent from a WiFi router, placed in physical proximity of a target or victim. While routers exist in most organizations today, the primary disadvantages of these methods is that they cannot be used to recover speech, or they rely on a previously collected dictionary to achieve their goal." In other words, it only works from previously known words.

And we talked about this, sort of something related, a long time ago, how reflected WiFi signals could be used to detect a person's motion or movements within a room. And apparently a group of researchers managed to detect the much smaller movements of the mouths of people speaking via WiFi signals, and then map them back into what they must be saying. Needless to say, it was strictly a proof of concept and would hardly be practical.

Anyway, our guys continue, saying: "The second method, a laser microphone, relies on a laser transceiver that is used to direct a laser beam into the victim's room through a window. The beam is reflected off of an object" - or maybe the window itself - "and returned to the laser transceiver, which converts the audio-modulated beam into an audio signal. In contrast to the previous limited radio-based methods, laser microphones can be used in real time to recover speech. However, the laser beam can be detected using a dedicated optical sensor.

"The third method, the Visual Microphone, exploits vibrations caused by sound from various materials - a bag of chips, a glass of water." I think we talked about balloons once. You could have, like, you know, inflated balloons in a room, and the balloon would vibrate, it being relatively elastic and pretty easy to pick up sounds. They said: "...in order to recover speech by using a video camera that supports a very high frame rate per second, high FPS of over 2200 Hz. In contrast to the laser microphone, the Visual Microphone is passive, so its implementation is much more difficult for organizations and victims to detect."

However, the main disadvantage of this method, according to the authors, is that the Visual Microphone cannot be applied in real time because it takes a few hours to recover a few seconds of speech, since processing the high resolution and high frequency 2200-frames-per-second video requires significant computational resources. In addition, the hardware required, a very high frame-per-second video camera, is expensive.

So, they said: "In this paper we introduce 'Lamphone,' a novel side-channel attack that can be applied by eavesdroppers to recover sound from a room that contains an exposed hanging light bulb. Lamphone recovers sound optically via an electro-optical sensor which is directed at a hanging bulb. Such bulbs vibrate due to air pressure fluctuations which occur naturally when sound waves hit the hanging bulb's surface. We explain how a bulb's response to sound, a millidegree vibration" - actually I think it ended up being microdegree vibration - "can be exploited to recover sound, and we establish a criterion for the sensitivity specifications required of a system capable of recovering sound from such small vibrations."

I was actually thinking I should have called this "Good Vibrations" or "Bad Vibrations" or something. But anyway, speaking of which, Leo, did you ever see the - this is completely

off topic. But I sent a note out to my friends. Did you ever see "Yesterday" when it was in the theatres? The movie?

Leo: Oh, yeah. I loved it, too.

Steve: Okay, because it's available on Netflix, and Lorrie and I rewatched it on Friday.

Leo: That's the one where the Beatles never existed, yeah. And [Jack] has a bike accident and somehow enters a different dimension where the Beatles never existed, but he remembers all the Beatles songs. So he becomes a massively successful songwriter.

Steve: Yeah. Very cool. So for what it's worth, completely just out of the blue.

Leo: Yeah, what does that have to do with this?

Steve: Nothing. Not at all. Except I was talking about "Good Vibrations." But that's...

Leo: But that's the Beach Boys.

Steve: That's the Beach Boys, not the Beatles.

Leo: I get your logic. I can - you enjoyed it, yeah, it's a good movie. I really liked it.

Steve: I really did. And speaking of which, since I'm off topic, it turns out, remember "Devs," the show that we loved?

Leo: Yeah?

Steve: The eight-episode series on Amazon?

Leo: Yeah.

Steve: The co-producer is a Security Now! podcast listener. The reason "Devs" was as accurate as it was is because of this podcast.

Leo: No. Because remember, I recommended it to you because in the first 15 minutes there's a debate over elliptic curve cryptography versus RSA. And I thought, gee, that's pretty good.

Steve: From the podcast.

Leo: Son of a gun.

Steve: And in fact, when some of the actors wanted some help in getting a sense for what techie talk sounds like...

Leo: No.

Steve: He said: "Listen to a couple episodes of Security Now!."

Leo: Oh, man. So when they bring that back - I don't know if they can. I guess they're kind of done.

Steve: The loop is complete. Oh, it was fabulous. If anyone, you know, it was on Amazon Prime, and absolutely top recommendations. So anyway, I thought that was very cool.

Leo: That's wonderful, yeah. Nice.

Steve: Anyway, so they determine the criterion for the sensitivity specifications of a system capable of recovering sound from such small vibrations. They said: "Then we evaluate a bulb's response to sound, identify factors that influence the recovered signal" - like I said, these guys go all the way - "and characterize the recovered signal's behavior. We then present an algorithm we developed in order to isolate the audio signal from an optical signal obtained by directing an electro-optical sensor at a hanging bulb. We evaluate Lamphone's performance on the tasks of recovering speech and songs in a realistic setup. We show that Lamphone can be used by eavesdroppers to recover human speech which can be accurately identified by Google Cloud Speech API."

Which I suppose means that it would be possible to set up an entirely automated listening station which triggers human involvement only when specific keywords were voiced because Google's Cloud Speech API could be doing speech recognition on the fly, turning it into text, looking for keywords, and then letting someone know.

Leo: Unbelievable.

Steve: And they also said singing, which can be accurately identified by Shazam and SoundHound...

Leo: [Laughing]

Steve: Oh, yeah, from a bridge located 25 meters, 27 yards away from the target office containing the hanging light bulb.

Leo: Wow.

Steve: They said: "We discuss potential improvements that can be made to Lamphone to optimize the results and extend Lamphone's effective sound recovery range. Finally, we discuss countermeasures that can be employed by organizations to make it more difficult for eavesdroppers to successfully use this attack vector."

So I was unsure how much time we would have, but we do have a little bit more, so I'll share a little bit, just so our listeners can get a sense for what they did. In Section 4 of the paper, bulbs as microphones. They said: "We measure the vibration of a hanging bulb, and we establish a criterion for the sensitivity specifications for a system capable of recovering sound from these vibrations.

"Experimental setup: We attach a gyroscope, the MPU-6050GY-5216" - for anyone who wants to duplicate - "to the bottom of a hanging E27 LED light bulb (12 watts). That bulb was not illuminated during this experiment. A Raspberry Pi 3 was used to sample the gyroscope at 800 Hz. We placed Logitech Z533 speakers very close to the hanging bulb, one centimeter away, and played various sine waves at 100, 150, 200, 250, 300, 350, and 400 Hz from the speakers at three volume levels - 70, 95, and 115 dB. We obtained measurements from the gyroscope while the sine waves were played.

"Results: Based on the measurements obtained from the gyroscope, we calculated the average peak-to-peak difference in degrees for theta and phi. The average peak-to-peak difference was computed by calculating the peak-to-peak difference between every 800 consecutive measurements, that were collected from one second of sampling, and averaging the results. The frequency response was a function of the average peak-to-peak difference, revealing three interesting insights.

"The average peak-to-peak difference for the angle of the bulb is, one, very small, 0.005" - okay, so it is milli - "to 0.06 degrees." So that's between 5 and 60 millidegrees. "Two, increases as the volume increases." Okay, not surprisingly. "And, three, changes as a function of frequency." Which you would expect, right, because it's not, I mean, it is far from being a neutral frequency response. You don't buy your light bulbs based on, oh, has a flat frequency response within 3 dB. No.

Leo: I might start.

Steve: "Based on the known formula of the spherical coordinate system, we calculated the 3D vector," they say, "x,y,z, that represents the peak-to-peak vibration on each of the three axes by taking the distance between the ceiling and the bottom of the hanging bulb into account. We calculated the Euclidean distance between this vector and the vector of the initial position. The results show that sound affected the hanging bulb, causing it to vibrate in 300 to 950 microns between the range of 100 and 400 Hz."

Then, under "Capturing the Optical Changes: We now explain how attackers can determine sensitivity of the equipment - an electro-optical sensor, a telescope, and an A to D converter - needed to recover sound based on a bulb's vibration. We've established the criterion for recovering sound. The attacker's system - consisting of an electro-optical sensor, a telescope, and an analog to digital converter - must be sensitive enough to capture the small optical differences that are the result of a hanging bulb that moves in 300 to 950 microns. In order to demonstrate how eavesdroppers can determine the sensitivity of the equipment they will need to satisfy the above criterion, we conduct another experiment.

"Experimental Setup: We directed a telescope at a hanging 12 watt E27 LED bulb." And as I said, I have a diagram at the top of the show notes. "We mounted an electro-optical

sensor, the Thorlabs PDA100A2, which is an amplified switchable gain light sensor that consists of a photodiode" - which you need for frequency response, a photo transistor is not fast enough, says he who did the light pen for the Apple II - "used to convert light to electrical voltage, to the telescope. The voltage was obtained from the electro-optical sensor using a 16-bit A to D converter, the NI-9223 card, and was processed in LabVIEW script that we wrote," they said. "The internal gain of the electro-optical sensor was set at 50 dB. We placed the telescope at various distances - 100, 200, 300, 420, 670, 830, and 950 centimeters - from the hanging bulb and measured the voltage that was obtained from the electro-optical sensor at each distance.

"The results of this experiment were used to compute the linear equation between each two consecutive points. Based on the linear equations, we calculated the expected voltage at 300 microns and 950 microns of bulb displacement. From this data we can determine which frequencies can be recovered from the obtained optical measurements. A 16-bit A to D with an input range of from minus 10 to positive 10 volts, for example the NI-9223 card used in our experiments, provides a sensitivity of 300 microvolts, provided by a 16-bit A to D converter, which is sufficient for recovering the entire spectrum from 100 to 400 Hz in the 200 to 300 centimeter distance range, because the smallest vibration of the bulb, 300 microns, from this range is expected to yield a difference of 300 microvolts. However, this setup cannot be used to recover the entire spectrum in the 670 to 830 centimeter range, so an A to D converter that provides a higher sensitivity is required."

And anyway, they go on like that, and you get the sense. It is now possible, they ended up verifying, after they'd calibrated everything and figured out what equipment they needed, they did do the 27-yard experiment with a light bulb and people talking in the room. It is now the case that anybody talking in a room which is visible to the outside can be overheard if there's something in the room that is vibrating because we now have the technology to look at it and recover the speech in the room from that. So congrats to the guys at the Ben-Gurion University and the Weizmann Institute of Science. And I wonder who else might be reading their paper with interest?

Leo: Oh, come on, Steve. The NSA's had this for years. You know they have.

Steve: That's true. That's true. This is not news for them.

Leo: I mean, you could even speculate this would be possible. So I'm sure they thought of this way back when. It's cool, though. It's cool research, and not surprising at all, yeah. I'm sure plate glass windows vibrate. It's just a question of whether you can measure such a small vibration; you know?

Steve: Yeah. With that, you know, we know that a laser interferometer will do that.

Leo: Sure.

Steve: So you can bounce a laser. But then you're going to see a red spot on the window.

Leo: Right, right.

Steve: And so what's cool about this is that they want to take a totally passive approach. Nothing in the room, no listening devices, just find something that's vibrating. And my goodness, you know, think about it.

Leo: Everything vibrates.

Steve: Yes, exactly.

Leo: Right.

Steve: So, I mean, and a light bulb is relatively rigid. So it's going to be resistant to vibration compared to, like, a leaf on a plant, a plant in a flowerpot or something. So they tackled the tough one.

Leo: I hope they got an "A," that's all I can say.

Steve: Somebody is now "Dr." who wasn't before.

Leo: Doctor, you ought to be the doctor. Steve Gibson, he is at GRC.com, the Gibson Research Corporation. That's where he makes SpinRite, the world's finest hard drive recovery and maintenance utility, currently working hard on 6.1.

Steve: I'll have an announcement next week. I'll have some code for our listeners to play with.

Leo: See, now, if you buy 6 right now, you'll get a free upgrade to 6.1, but you'll also be involved in the beta testing of it. So I think this would be...

Steve: Yes, sir.

Leo: Head over to GRC.com. While you're there, lots of free stuff to test out, including ShieldsUP! and, oh, it's just a beautiful nest of fun facts to fill your mind. And of course this show is there. He has 16Kb audio, 64Kb audio. He's got transcripts so you can read along as you listen, or use them for searching. That's really a great tool there because you can search for a phrase or a term or a name, and it'll jump right to that part of the podcast. All of that's at GRC.com.

We don't have anything so fancy, but we do have video, so you can watch Steve, if you really want. That's at TWiT.tv/sn. We do the show Wednesdays right after MacBreak - sorry, Tuesdays. We only changed this years ago, just it's not sinking in. Tuesdays, 1:30 Pacific, that's 4:30 Eastern time, 20:30 UTC.

If you want to listen live or watch live, go to TWiT.tv/live. There's streams there. Chat with the chat room while you're listening, irc.twit.tv. Get on-demand versions at the website, TWiT.tv/sn for Security Now!. And let's see, I guess the one more

thing to say is subscribe. That way you'll get it automatically. You can start building your collection of 771 Security Now! episodes. Very cool.

I love that "Devs" story. That blows me away. That is awesome. It makes perfect sense in hindsight, but that's just great. Steve, have a wonderful week. Stay safe. And I will see you next Tuesday on Security Now!.

Steve: Thanks, buddy.