SECURITY NOW!

Transcript of Episode #647

The Dark Caracal

Description: This week's news continues to be dominated by the industry-shaking Meltdown and Spectre vulnerabilities. We will catch up with what's new there, then discuss the Net Neutrality violation detection apps that are starting to appear; a new app and browser plugin from the search privacy provider DuckDuckGo; a bit of welcome news from Apple's Tim Cook about their planned response to the iPhone battery-life and performance debacle; a bit of errata; and some feedback from our terrific listeners. Then we take a look into a state-level, state-sponsored, worldwide, decade-long cyberespionage campaign which the EFF and Lookout Security have dubbed "Dark Caracal."

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-647.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-647-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. And as always, lots of security news. We'll be talking about it all, including a security threat pinpointed by Lookout Security and the Electronic Frontier Foundation. Out of Lebanon, the Dark Caracal. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 647, recorded Tuesday, January 23rd, 2018: The Dark Caracal.

It's time for Security Now!, the show where we cover your privacy, your security, your deep-rooted insecurities with this guy right here, Steve Gibson.

Steve Gibson: And if you don't have them at the beginning of the podcast...

Leo: You will.

Steve: You will know what you've been missing.

Leo: Very good point. Excellent point. Hello, Steve.

Steve: Yo, Leo. Great to be with you once again for Episode 647 of this, what are we, we're in our 12th year and counting.

Leo: And this one has the most mysterious title I've ever seen.

Steve: The Dark Caracal.

Leo: What the hell?

Steve: C-A-R-A-C-A-L. Yes. I was hoping that we could be done this week talking about...

Leo: Oh, Intel? Oh.

Steve: Talking about Spectre and Meltdown vulnerabilities. But no.

Leo: It's worse than you thought.

Steve: Yes, it's continuing to generate news. So we will deal with the updates on Spectre and Meltdown, and also talk about Linus Torvalds' freakout.

Leo: Oh, he was mad.

Steve: Oh, he was really not happy. So we'll catch up with the news there. Then I think maybe because the whole industry is so preoccupied, there wasn't a ton of other interesting stuff. There's an interesting few apps which are appearing which are purporting to be Net Neutrality violation detection apps. So I want to talk about them a little bit. DuckDuckGo, the well-known privacy-focused search provider, has some new offerings that I want to cover. Tim Cook also recently said something that really made me happy on the iPhone performance battery front, that I just want to touch on.

We actually have some errata from a misstatement from last week, not that that should be a big surprise, but we haven't had any for a while, that I want to cover. Then we'll deal with some feedback from our listeners and then dip into the world of a state-level, state-sponsored, worldwide, decade-long, cyberespionage campaign which the EFF and Outlook Security have dubbed "Dark Caracal."

Leo: Oh, my.

Steve: This is sort of like the exact opposite of the news we've been covering recently where every one of our listeners has, like, comes away with action items. It's like, oh, crap, what button do I push? What do I download? I need to update this and blah blah blah. Here there's nothing you can do.

Leo: Okay, good. At least there's nothing I can do.

Steve: There's nothing you can do. But by sort of putting a point on this and looking at this - it was a 51-page report that I read last night in order to really read into this, with lots of details. I sort of wanted to synchronize us all with an aspect of what is clearly going on, not just in this case in a building focused in Beirut, Lebanon, but I'm sure every state, every nation-state has something like this - the U.S., China, Russia, and on and on and on.

And so here's an example of a big cyberespionage campaign that a nation has had underway for some time. And it's just for the, I mean, the title of the podcast is Security Now!. And it's worth us sort of all appreciating that, although there's nothing we can do about it, except be careful if you don't want to become a target, this is all happening. And it's sort of sobering. One line from the report really kind of caught my attention, and we will get to that toward the end of the podcast today. So I think something interesting for our listeners.

Leo: I know it's not in your notes, but there was a great article today in BuzzFeed News, "Inside the Fight for the Soul of Kaspersky Lab," that I highly recommend. Fascinating. And basically the conclusion was there was a struggle between Kaspersky and some members of management and others over how much influence the Russian government would have.

Steve: Should have or would have, yeah.

Leo: Should have, would have. Now, you've got to consider that they're relying on a story from a Russian publication called Meduza. But they also did their own research, and they concluded that the battle was won by the Russian security services. They even assert that Kaspersky has built into it a way to look at any computer that's got Kaspersky installed upon it, to look into their files. And, now, of course Eugene Kaspersky denies that heatedly, as does the company. But that's [crosstalk].

Steve: They have to. I mean...

Leo: I would if I were them.

Steve: Yeah.

Leo: Oh, yeah, we can do that, yeah.

Steve: Oh, is this Russian-backed spyware? No, no.

Leo: Yeah. No, no, no, no, no, of course not. But I think that was, I mean, if you were on the fence about Kaspersky...

Steve: That would tip you over.

Leo: I think this article would...

Steve: And not in the direction of renewing your subscription.

Leo: Yeah. Yeah. Basically this is the quote. According to a former senior manager who was involved with launching Kaspersky Labs and the product, or I'm sorry, KSN, which is the Kaspersky Security Suite, I think it was called, it was launched in 2012: "The product was referred to as 'cyberintelligence' inside the company. The system can be run manually from a remote location." Oh, I'm sorry, it's Kaspersky Security Network can be run manually from a remote location, he told Meduza, meaning an employee of the Kaspersky Lab can download any file from a computer on which KSN is installed without its owner's knowledge.

Steve: And you know, I mean, should that really surprise us that much?

Leo: No.

Steve: I've often commented here on the podcast that I'm amazed in the era of Linux, which is open source and beautifully designed and maintained and architected and functional, that non-U.S. companies are choosing to run Windows. I just...

Leo: Right. For the same reason. Right.

Steve: Yes. It just boggles my mind that they're like, oh, yes, we run Windows. It's like, what? Really?

Leo: But it's interesting because Kaspersky originally would run locally only; right? And when they launched this KSN in 2012, it was a network-based security solution; right? And they said moving to a cloud solution lets us analyze and neutralize new viruses and et cetera, et cetera.

Steve: Yup, yup.

Leo: And there are a number of other antiviruses that do this. But of course an always-on connection has - one of their sources says it's like an awesome kitchen knife that can be used for superbly slicing bread or stabbing people. Okay. You decide which use you want. Anyway, I thought I'd pass it on. I'll send you the link. I thought I'd pass that along because we've been talking about this back and forth. This one seems to put a nail in the...

Steve: Yeah, well, and I haven't gone into it because all we have is hearsay. I mean, you know...

Leo: Yeah, and that's still probably the case.

Steve: I want to talk about bits in machine status registers, which are clear and well defined, and I can say something.

Leo: You're an engineer. You like...

Steve: I could say something about that. But, you know, it's like, well, you know, yeah.

Leo: Our Image of the Week talks about philosophy versus engineering, as well. I like it.

Steve: It does.

Leo: Yeah. All right, Steve. Shall we do the picture? Yeah. I love this.

Steve: So our picture, it's a four-frame cartoon that a bunch of our listeners sent me because of course it's right on target for us. The first frame has - each frame has two characters talking. The first frame has the first guy looking at his laptop, saying, "But you can't do that. It goes against every basic security practice." And we can't see the screen. We're looking at them in the back of the screen, so we don't know what it is that this goes against. But the other guy says, "Yeah, I know." And then the second frame he continues, "But in philosophical terms, it makes sense." And then the first guy says, "You don't get paid for philosophy. You get paid to code."

And then the second guy in the third frame says, "But if it's just one character off, we should give a little encouragement; right?" And the first guy says, "No, no, no, no, no, no." And finally we see the screen that has generated this controversy. And it's your typical username and login. And the response to the password is, "The password is almost correct. You're close. Keep going." And then offscreen we see the little dialog box: "If the password is wrong, we show 'error' and that's it."

Leo: Yes.

Steve: So of course we've even spoken of how, I mean, there are many ways to do this wrong. And one of the things that we sometimes see is a login that asks you for your username and your password and then says, well, you got the username right, but the password wrong. And it's like, wait, you know, we just gave away some information that wasn't necessary to give, very much like this, which inherently allows the security to be chipped away at. Properly designed, if you're going to ask for both at once, you say, "We're sorry, we couldn't find you," or like there's something wrong with your information. Again, don't give them any information that they don't already have. Just say [buzzer sound], exactly, just error. Not, well, you know, except for the fourth character...

Leo: So close.

Steve: ...you were really good. No. No. So, okay, Spectre and Meltdown. Now, first of all there's a lot of misinformation again, just because I think editors are rushing to get to their pages online. And so they're talking about Spectre and Meltdown together, when we know on this podcast they are related more in time than they are in nature. So Meltdown is the problem which is easy to resolve and may cause a performance hit or may not, depending upon how recent your chipset is. But significantly, its resolution, no event requires a firmware update for your chips from Intel, and AMD never had a problem because their chips were not vulnerable to this.

So the reason I bring this up is that Intel has just reversed its so-called "advice" to all of its OEMs, saying stop, stop, stop, stop patching anybody's firmware. We made a mistake in last week's firmware. And I don't remember if we touched on it last week, but there were - it's hard to nail them down. They're saying something as fuzzy as "May cause reboot issues," or "May make your system reboot more often." It's like, wait a minute. More often than never?

Leo: More often?

Steve: More often than what? In other words, this causes a catastrophic crash of your system. But they don't want to use those words. So yesterday the newsroom of Intel sent out, quote: "As we start the week, I," writes in first person this newsroom persona..."

Leo: Mr. Anonymous, yeah.

Steve: Ah, yeah. "I want to provide an update on the reboot issues we reported January 11." Okay, so that was 12 days ago, but 11 days ago yesterday. "We have now identified the root cause for Broadwell and Haswell platforms and made good progress" - Leo, they're making good progress.

Leo: Well, well.

Steve: "...in developing a solution to address it. Over the weekend" - oh, and you can just imagine there's a lot of late nights over there at Intel. "Over the weekend we began rolling out an early version of the updated solution to industry partners for testing." Oh, so everybody was awake. "And we will make a final release available once that testing has been completed." Okay, but not now. Not yet. "Based on this, we are updating our guidance for customers and partners. We recommend that OEMs, cloud service providers, system manufacturers, software vendors, and end users stop deployment..."

Leo: Oh, my god.

Steve: "...of current versions."

Leo: And those of you who've already installed it, oh, gee, we're sorry.

Steve: Oh, and actually that's happening. Dell has announced in following the guidance that they will be offering previous versions of the firmware immediately because, like, to just - because now Intel has scared everybody that reboots might happen any minute.

Leo: So this is, though, but it's Haswell and Broadwell systems. It's not Coffee Lake, Kaby Lake. It's not the more modern systems. It's older systems.

Steve: Oh, no, it's - no, no.

Leo: Oh, it is?

Steve: They haven't gotten to those yet.

Leo: Oh, great.

Steve: Oh, yeah. It's everything. So, okay. So they're saying stop deployment of current versions as they may introduce - get this, again - "higher than expected reboots." What? Well, okay. Like I'm not expecting any. So I guess a reboot which is unexpected, well, which you didn't Ctrl-Alt-Delete to make happen, or pull the plug out or something. Anyway, may reboot. And, they say, "other unpredictable system behavior." Now, by its very nature, unpredictability is something you are working to avoid in your processors. So, yeah. Then they say: "For the full list of platforms, see the Intel.com Security Center site." Which is very busy lately.

Okay. Meanwhile, the technical side - that was Mr. "I'm giving you an update from the coffee room." Here we have this title. This page is titled properly: "Speculative Execution and Indirect Branch Prediction Side Channel Analysis Method." So you know you're getting something a little more meaty than, ooh, more than expected reboots. Okay. So this was updated yesterday.

"We have now identified the root cause of the reboot issue impacting Broadwell and Haswell platforms and made good progress in developing a solution to address it." Still they're making progress. They haven't arrived yet. "Based on this, we are updating our guidance for customers and partners." So you can tell that Mr. Press Guy basically pulled his language from the technical guys. Same recommendations to everybody: "Stop deployment across the board of current versions on the below platforms," they write, "as they may introduce higher than expected reboots," blah blah blah.

"We also ask that our industry partners focus efforts on testing early versions of the updated solution for Broadwell and Haswell we started rolling out this weekend, so we can accelerate its release." In other words, we don't really know, apparently we're unable to decide that this is fixed. So we need everybody to try it and let us know if we're getting warmer. Did your unexpected reboots slow down? Whoa.

Then they write: "We expect to share more details on timing later this week." That is, this current week. "For those concerned about system stability" - oh, and I guess there are people who are not concerned. But "For those who are concerned about system stability, while we finalize the updated solutions, we are also working with our OEM partners on the option to utilize a previous version" - like one of the ones that worked - "of microcode that does not display these issues." Right, no one ever uses the word "bug" anymore. We don't have bugs, we have issues and more frequently than expected reboots. Okay.

But, they say, it no longer crashes, but it "removes the Variant 2 Spectre mitigations." In other words, of course, you're vulnerable again; but boy, you're able to stay online, which is handy. "This would be delivered via a BIOS update," which might be called a "downdate" in this case, "and would not impact mitigations for Variant 1" - which is not a problem, which did not require the firmware update - "and Variant 3 Meltdown," which never had any of these problems.

They say: "We believe it is important for OEMs and our customers to follow this guidance" - in other words, quickly run away back to a BIOS that worked, earlier firmware - "for all the specified platforms listed below, as they may demonstrate higher" - oh, here we are again - "than expected reboots and unpredictable system behavior. The progress we have made in identifying a root cause for Haswell and Broadwell" - get this, Leo - "will help us address issues on other platforms." Oh, we're going to learn why this didn't work for Haswell and Broadwell, and that'll help us fix the other ones. "Please be assured," they say, "we are working quickly to address these issues." So here it is.

"The guidance applies to at least some of the processors from Intel's last several generations of chips, with affected models in the Broadwell, Haswell, Coffee Lake, Kaby Lake, Skylake, and Ivy Bridge families."

Leo: But they said continue to patch the newer ones, which to me...

Steve: Yeah. "Certain lines..."

Leo: You know, Red Hat may have given them the warning on this because they pulled their update last week, saying we're no longer providing microcode to address Spectre Variant 2 due to instabilities that are introduced, causing customer systems not to boot.

Steve: Yeah. Minor problem.

Leo: Yeah.

Steve: "Certain lines," they write, "are affected more than others." I mean, this is all so weird and murky. "For example, only Ivy Bridge datacenter/workstation" - okay. "Datacenter/workstation processors are included." But chips from most recent consumer lines also appear to be impacted. So that sentence doesn't seem self-consistent.

"Intel says that it's identified the issue behind the unexpected reboots on Broadwell and Haswell processors and is working toward releasing an update that addresses the exploits without causing the issue. But the same problems have been seen on Ivy Bridge, Sandy Bridge, Skylake, and Kaby Lake processors, too. Intel says it's 'actively working on developing solutions' for those platforms, as well." So bottom line is nobody should - okay. So if you have not installed the firmware, don't.

Leo: God. I just got a pushed - Apple's pushing out an update right now.

Steve: Yeah.

Leo: But Apple did not get bit by this because they didn't do the microcode patches initially.

Steve: Right. I have my main laptop and two other - two Dells and a Lenovo. I've updated them all. I've been using them, and none of them have rebooted unexpectedly.

Leo: I haven't had any reboots, either.

Steve: So I guess the advice would be, I mean, it's better to be secure than not. Although remember the biggest problem, Meltdown, that is the only one that we've ever seen even a proof of concept for, isn't a problem. That got solved just with a Windows Update. No firmware update needed. Spectre is a known problem, but much more difficult to exploit. And for none of this have we ever seen an in-the-wild exploit. And it's not remote anyway. It's something gets into your computer and is able to peek across process boundaries. And it's not clear that Windows is good about maintaining those anyway for an individual. I mean, debuggers function by deliberately reaching into another processor and sticking its hands in there so that it can see what's going on and help a developer to fix broken code.

So the big concern had always been in the cloud, where you inherently had a single executing block of hardware that could have many different users sharing the hardware. And so it was cross-user boundary was the concern. But on your own machine, you're your own user. Unless you get some, I mean, theoretically you could get some malware in there that could get up to some monkey business. But no one has seen that yet.

So our advice would be, if you have not updated your firmware, don't bother. Why create problems for yourself? It's harder to fix your BIOS if you have a nonbooting machine than it is - because nowadays you're able to update your BIOS, typically from within the OS. You just download the app and say yes; and yes, I'm sure; and yes, I'm really sure. And then it says, okay, be sure not to kick out the power cord for the next minute while we reflash your BIOS, and you'll be all fixed.

So if you end up with an unbootable system, undoing that requires booting from a thumb drive and adjusting your BIOS to boot if you can and so forth. Potentially, you could brick yourself, so that's not good. So if everything's fine, and you have updated your firmware, like I have and you have, Leo, and nothing has rebooted, then we're probably okay.

Leo: Yeah. I just checked on my Lenovo T470s, and there's no updates.

Steve: Right.

Leo: So I don't know if Lenovo is rolling back those updates at some point or what they're doing.

Steve: Well, this has just happened. So it was like a long weekend at Intel, and now we have this advice that hit yesterday. And I just saw announcements. People are tweeting me that they're getting news from Dell saying go back to a previous BIOS. So I think what will ultimately happen is we will eventually get new firmware which is solid. And everyone should either update again or update for the first time at that point.

But given the fact that what we have now is flaky, if you are concerned, you could certainly go back to a previous BIOS. I know that generally when you go to the firmware and BIOS pages, you're able to say, oh, look, here's the earlier one, and the one before that, and the one before that. You could choose the one before last week and have a firmware which you know is going to be solid because it comes along with the BIOS. My feeling is, if there isn't a problem, essentially, if everything's fine, stay where you are. If systems have become unstable, back out.

And in all cases we'll wait till the dust settles, till everybody agrees that this problem has been solved. And then the good news is any manufacturer who just gave us a BIOS and firmware update two weeks ago will certainly still be around to fix the one that they gave us two weeks ago with something that is known to be more solid moving forward. And boy, it'd be fun, again from an engineering standpoint, to get some clarity into what is all going on. It's just difficult to understand from an engineering standpoint. What have they done that causes more reboots than usual? Wow.

Anyway, so that's where we are. I would say follow the advice of whoever your supplier is. For example, if Lenovo tells us you absolutely really must undo what we did two weeks ago, well, then I'd be tempted to. But they may be saying, as they generally do with BIOS updates, if there's no problem, don't mess with it. If everything's fine, just move along. So, wow.

Now, against this backdrop, Linus Torvalds has been having a temper tantrum meltdown.

Leo: Oh, boy, is he pissed.

Steve: In public. I mean, dropping the "F bomb," and calling this a bunch of S-H-I-T, and cannot believe what Intel is doing. And in pursuing this, I just hours ago, this morning, I found the document I had been looking for but hadn't tracked down yet, which is the ink is still wet on the PDF. It's titled - it's a 15-page PDF from Intel titled: "Speculative Execution Side-Channel Mitigations." This is the technical details of what Intel is proposing for the microcode fixes.

And just because I was putting the podcast together I had no chance to dig into it. But, boy, I'm going to have fun tonight because I'm really curious, having built the InSpectre to bring some clarity to this. And having taken the approach I took, now I'm really interested to - but I was never - I was always running from third-party comments. I found some ESX VMware documentation that I based some of the stuff on. And I'm so happy to get finally the original root document that is the center of this controversy.

So, I mean, and the problem is Linus isn't, in anything I have read that he has written, being clear. He's just ranting. And it'd be nice if he were to catch his breath and calm down and explain what it is that seems to be the problem. Near as I can tell, he's unhappy that Intel intends to call a "solution" the simple turning off of a couple flavors of branch prediction. That's what, you know, the whole Spectre thing is that it's possible to leverage this speculative execution of branches and branch prediction in a way that leaks information.

And so Intel's solution for now, and they've said for the time being, is to allow operating systems to turn it off. So his tantrum is that they're disabling all branch prediction, with the concomitant performance hit that that will bring, as their solution. And he's wanting them to fix it right.

Now, the reason I think he's wrong, that is, the reason I don't have sympathy for his position, I mean, yes, it would be wonderful if we could have that. But microcode is not like application code. By that I mean that microcode is tightly bound to the hardware. So it's not like there's a whole processor written in software that Intel can just say, oh, well, we've decided we're going to change everything. No. It's that what the industry realized decades ago was that implementing an instruction set as complicated as Intel's, in pure hardware, that is, in gates, in just simple logic gates that do all the work, was becoming impossible.

This is the difference with the ARM architecture. Because the ARM instruction set was so simple, was so clean, they've been able to give us a high-performance system that doesn't rely - I don't know if it relies at all on microcode. But if it does, not nearly to the degree that Intel's does. So essentially, Intel's instruction set got so complicated that you could no longer actually do it in hardware. You had to back off.

And so what they did was they moved some of the complexity into sequential functions which are driven by this microcode. But the point is you can't - for example, their branch table prediction system doubtless is a hybrid of hardware that does most of the work, and some microcode that just kind of points in different directions and says, oh, now raise this line, or let this data flow through here. Okay, now stop it, now run a compare, and that kind of thing. It's not actually doing the work, it's directing the hardware which is purpose-specific to do the work.

I have no doubt that Intel has learned some incredibly important lessons from this and that, downstream, a future generation of Intel processor will solve this problem. But they can't solve it with a patch to the microcode. And that's the problem. So what Linus is asking for is not feasible. What they've done is the only thing they could do with the existing chip hardware, which is say, "Oh, crap."

Leo: Doh.

Steve: And basically they've given us the "oh, crap" bit.

Leo: Yeah, we got that part set.

Steve: Yeah. And so here's the "oh, crap" bit. And if you turn it on, well, sorry, things are going to run slower. But that's the only thing we can do. We can let things run fast, and you hold your breath; or you can turn on the security bit, and it's going to go slower. But this class of exploits can't happen.

So Linus is not happy. But what he wants, you can't get there from here because microcode isn't like operating system code, where you could say, oh, we're going to change everything. On the same chip we're going to give you Windows 98. Now we're going to give you NT and Windows 10. No. It's the microcode is just basically the minimum required to simplify the hardware. But you're stuck with the hardware that you've got, which is doing all of the real work. Microcode is just like a traffic cop, raising and lowering, allowing data to kind of flow where it's supposed to and simplifying the process. So in fact I'm reminded of Wozniak's floppy disk controller because that was the genius that Woz designed where he gave us this six-chip floppy controller that was microcoded.

Leo: Ah, interesting.

Steve: He'd used microcode by taking some of the outputs of the ROM through a latch, and then fed them back into the address inputs of the ROM. And so by strobing this latch, he basically created a little microcode program in ROM that was, like, essentially, a little hardware processor that just did floppiness. And it was astonishingly elegant and incredibly inexpensive. And where everybody else had these - remember, the floppy controllers at the time were these monstrous huge cards where you couldn't get any more chips on the thing. He had three. Or, I mean, he had six, six little chips, just sort of wandering around lost on this small little floppy controller. And he's like, yeah, well, that's the way you should do it. It's like, with incredible modesty.

But anyway, so what I think, where we're going to be is that there's no question that Intel's future architecture was just scrapped. I mean, their classic back to the drawing board because they can't sacrifice performance for security in the long term. But it's not going to be - and I don't even know what their design pipeline is. It's like years long, I think. They've got multiple big teams running on stuff that we'll see next year and the year after that and the year after that. A lot of work is probably already committed. I don't know what this is going to do to their whole lifecycle system. But it's doubtless impacted it.

Leo: It doesn't affect Retpoline, right, the Google fix?

Steve: Correct.

Leo: Okay.

Steve: And we do have a solution for that from Google, which...

Leo: They say has no performance hit.

Steve: Correct, exactly, because you're able to - it's a very, like a cycle or two at the end of a subroutine which already burned up thousands and thousands of cycles being in the subroutine. So like not measurable.

Leo: Hey, here's some good news. Remember last week I ran InSpectre using WINE on my 2012 iMac, and it wasn't a very good result. I think it was No, No, No, across the board. I mentioned that Apple pushed out a fix today. And look at this. All of a sudden Yes, Yes, Good.

Steve: Oh, cool.

Leo: Oh, no. Yes, Yes, is not good. Yes, Yes is vulnerable to Spectre. Never mind. It's just as bad as it was before. I got confused by the yes. Yes is bad. No is good.

Steve: Yes. And in fact several people have said, "Steve, you know, you need to rethink the indicators because..."

Leo: [Crosstalk] patched for Meltdown, no. Vulnerable to Spectre, no. I mean, patched [crosstalk], something like that.

Steve: Yeah. I need to reverse that because...

Leo: Yeah, "Yes" sounds good, yeah.

Steve: It is confusing, and I'm relying on color.

Leo: Darn it.

Steve: And "Yes" sounds, oh, yes.

Leo: Darn it. I was all excited. So I may never get it fixed, this is such an old machine. And you did, and we should update this, we talked on Sunday, or Saturday on The New Screen Savers, say that your correspondents have said that WINE does give you a full and useful result on Macintosh. It doesn't hide.

Steve: Inside of a virtualized environment like VMware, you're not really seeing the chip. So you cannot rely on what InSpectre shows in a true VM. But WINE is not that. WINE is just an API layer. And my InSpectre app is reading the data from the chip's registers in order to determine whether you have the firmware update...

Leo: And those aren't hidden.

Steve: Right, those are not hidden.

Leo: So that's the microcode patch, which it turns out it's actually maybe "yes" is right because you don't want them [laughter]. So, you know, maybe I'm in good shape after all.

Steve: No one wants their system to be rebooting ever, let alone more than usual. So, yeah. Woody on Windows, who's got a great column over at Computerworld, put out a - he's been tracking all of this confusion. And, boy, it is confusing. One of the things that users of 32-bit Windows have been annoyed by is that Microsoft hasn't been keeping them updated. That is, their systems seem to be late in getting any of these patches. So I wanted to note that there is now, for Windows 10 version 1709, which was the Fall Creators Update of Windows 10, for x86-based system, that is, for 32-bit, there is an update, but you've got to go get it.

Woody wrote: "Cumulative update KB4073291 brings the Meltdown/Spectre Windows patches to 32-bit machines." And then he says: "What? You thought 32-bit machines already had Meltdown/Spectre patches? Silly mortal. Microsoft's Security Advisory [blah blah blah] has the dirty details in the fine print, point 7."

And I think we talked about that before, where they asked themselves the question, "I have an x86 architecture, and the PowerShell Verification output indicates that I am not fully protected from these speculative execution side-channel vulnerabilities. Will Microsoft provide complete protections in the future?" They answer their own question: "Addressing a hardware vulnerability with a software update presents significant challenges and mitigations for older operating systems that require extensive architectural changes. The existing 32-bit update packages listed in this advisory fully address" - and then there's two of the problems - "5753 and 5715, but do not provide protections for the 5754 at this time."

They say: "Microsoft is continuing to work with affected chip manufacturers and investigate the best way to provide mitigations for x86 customers, which may be provided in a future update."

So then Woody says: "It appears as if this is the first 32-bit version of Windows" - meaning the Fall Creators Update, that 1709 version - "that has a patch for the Meltdown vulnerability. Surprise." And then he says: "Like most of the patches I talked about yesterday, this one is available only through the Update Catalog. It won't be pushed to your machine." So I've got the catalog update link in the show notes. Don't know if there's an easier way to find it. Maybe you could Google, what is it, "KB4073291 x86" or 32 or something, in order to find it.

So I thought that would be of interest to our listeners because you're not going to get it automatically, but it's there. You just have to go get it yourself. And that adds mitigation for Meltdown, which again, of all of this, that is clearly the more important. It's the one for which people are developing attacks. Proof of concepts exist. They're able to read data from arbitrary locations in other processes and in the kernel. That's the one you want fixed. It is, however, if you have an older processor, the one that will also give you the biggest performance impact. And using InSpectre, you can flip it off and on, back and forth, to see if you notice a difference and then choose whether you want it to be fixed or not. But you don't have that choice unless you get the update from Microsoft. So it is available.

Leo: So you can just go to catalog.update.microsoft.com and put in that KB4073291.

Steve: Oh, good.

Leo: And when you search for it, you'll get the download directly.

Steve: Ah, good. So there are a couple apps which have just surfaced that are of sort of questionable value. The first piece of news I picked up on this, it was a couple days ago, and there was some controversy because Apple had blocked it in the App Store, and the people said, "Hey, wait, are you trying to prevent us from detecting Net Neutrality violations?" And Apple said, "What? No." Then they had some banter back and forth, and they apologized, and they said, you know, "We're busy."

Leo: We didn't understand what this was for, yeah.

Steve: Exactly. Exactly. So in what is arguably a clever URL, one of them, and I think it's a different one, I'm not sure, because there are a couple, TestYourInter.net is the URL. So I thought that was kind of clever. TestYourInter.net is the URL. And that will take you to an app called OONI, O-O-N-I, which is the somewhat awkward acronym for the Open Observatory of Network Interference. Now, the good news is it's open source. It's a nonprofit app for tracking Internet censorship. Their goal is for it to be sort of a crowd source-y sort of thing, where everybody downloads OONI for their iPhone - I'm not sure if it's cross-platform, if it's Android also. I think I remember...

Leo: Yeah, they have Android, too. Yeah, they have both, yeah.

Steve: Okay, good. I remember seeing two links. Yeah, there they are. And then this thing pulls data from a number of different sources and checks the bandwidth. In some of the coverage I saw, it was already generating false positive alerts because they were noting that, oh, I got 25Mb from Apple, and 14Mb from Google, but only 8Mb from YouTube. And it's like, yes, because streaming is different. And so I hope that this ends up not being a boondoggle because, I mean, the idea of being able to, like, the idea of having surveillance of ISPs' provider-based bandwidth, that's sound.

But it is the case, for example, that a streaming provider that is already providing an ungodly amount of bandwidth to the world, I mean, the proper operation is to feed out a stream of content as needed so that they're not sending you a huge blob where you look at the first minute or two and then say, no, I'm not interested anymore. So you end up discarding a lot of bandwidth that you received from them that you yourself used or downloaded, but then never actually consumed. So it's going to be important to, like, be careful with any of these apps about how we interpret what they show.

The idea I like, the idea of an observatory, which is how they're couching this, where lots of users download it, everybody runs it, and then this creates a distributed probing network that can send information back to the mothership, and they can understand what's going on. That seems useful. And I like the idea that ISPs will know that if they do get up to any hanky-panky, they're going to be - it's going to be spotted. It's not going to get slipped under the rug because we're all wanting to make sure that Net Neutrality, which they say, oh, we want to have the freedom, but never, never, never, never would we think of abusing that freedom. It's like, okay, well, trust and verify. Or at least verify.

DuckDuckGo is a popular search provider, and I just wanted to note to our listeners that they have a cross-platform offering, a browser plugin for Firefox, Safari, and Chrome, and also iOS and Android apps, where they're looking to extend their well-known privacy protecting search to also a general web browsing, blocking, tracking protection. They say smarter encryption. What they're actually doing there is they're doing something that we've seen before. Some URLs are not secure, that is, http://. Their plugin and/or their browser on the mobile platforms will attempt to make a secure connection to the same domain and, if it's available, will use that preferentially over the nonsecure.

So it's a small thing. Many sites now bounce people over from the nonsecured to the secured automatically. GRC has done that, for example, for a long time. But certainly it's nice just to have that as an additional benefit. So not a huge thing. But I thought that some of our listeners would find it interesting. And I know that DuckDuckGo has a following among people who like the idea of being able to search with some additional privacy guarantees.

And lastly, in our light news week, thanks to the fact that Spectre and Meltdown have just been so dominant, I did note an article in Gizmodo saying that Tim Cook was promising to let iPhone users turn off throttling, which to me makes the most sense. I think I will probably still, on my older iPhone 6 Plus, take Apple up on their battery offer, if they ever get batteries back in stock again. I guess they're - do you know how far backordered they are, Leo? Have you been following that?

Leo: No, I don't, offhand, yeah. You should do it anyway, I mean, it's so cheap.

Steve: Yeah, why not? Exactly. Really extend the life of that device. But the best solution is to give us the option. And so the good news is - Apple wasn't saying when. I'll be interested to know if we just got it with this iOS...

Leo: No, they said the next version. I think 11.3 will do it.

Steve: So not a security - so just today, a couple hours ago, we got 11.2.5, which adds some new features to HomePod and other stuff, and fixes a bunch of stuff, and also has a lot of security content. So that's a standard security update. The nice thing is that soon we should have an option to turn off throttling and take responsibility for Apple's concern that this is going to make our devices spontaneously reboot. I was always a little skeptical of that. And maybe if your battery is - the only thing I could ever figure is that the battery was so weak that, if the phone had a sudden surge of demand for speed, that is, because we know that it does dynamically change its power demands, it might just, like, hit the battery with a demand it cannot meet, which would cause the voltage to drop below operating threshold and would indeed make the phone reboot.

Leo: Yeah, that's exactly what Apple said the mechanism was, yeah.

Steve: Okay. And the problem, of course, was that those who had taken good care of their batteries, even though it was old, they just did that without telling us and then denied it for a long time. It's like, oh, what? I mean, our listeners got tired of me saying something is broken. This does not work the way it used to.

Leo: Yeah, but what do you mean, take good care of your battery? Did you not charge it?

Steve: I never used it.

Leo: You were never on battery.

Steve: Well, yeah. In fact, I have one of our closing-the-loop questions here is one of our listeners saying, "You use the term 'babying your battery.' What exactly do you mean by that?" So I do plan to address that.

Leo: We'll get to that. Good, good.

Steve: Errata. I misspoke, and a number of our listeners were confused by it, for which I apologize. Sean Spratt said: "@SGgrc On your podcast you said only post-Haswell has the INVPCID" - that's the Invalidate Process-Context ID support instruction. He says: "However, I'm reading elsewhere that INVPCID is included in Haswell. InSpectre says I have high-performance Meltdown protection. So, yes, I have INVPCID."

Similarly, Richard Tan: "Hi, hope you are well. I keep hearing on Security Now! that Haswell and down would not be patched with the performance fix, but Haswell seems to be the first set of processor that has INVPCID. So should it be that Haswell onwards should get the performance fix for Meltdown?" Okay, so yes. I misspoke, and I apologize. To get the high-performance Meltdown protection you need Haswell or later. Because we need two instructions. And I did initially say it correctly, and I must have just off-the-cuff misspoken later, like a week later. And so I confused everybody.

The first instruction was introduced way back on January 7th in 2010, so 18 years ago. Wait, no, sorry, eight years ago. That was in Intel's Westmere architecture. And that was the PCID instruction. Intel meant well with that instruction. Technically it allowed the instructions in the cache to be tagged with the process-context ID that called that data into cache so that it would then be able to be reused, but no other process with a different context ID would be able to see that. Intel, it turns out, though, hadn't asked OS developers if this would be useful. They just said, "Here you go." And so it turns out it was really awkward to actually use. It just ended up being something that was nice, but nobody used it because we didn't have to, and it was just very difficult to implement.

So Intel learned from their mistake; and 3.5 years later, in the summer of 2013, they fixed it with the Haswell microarchitecture which was also the fourth-generation core successor to Ivy Bridge. That's when they added the Invalidate PCID instruction, the

INVPCID, with Haswell. And that's, it turns out, the key that allowed all of the OSes, Linux and Windows, to in very short order fix this meltdown problem. So I apologize for the confusion. Haswell and later, that is, Haswell and more recent processors will all have this INVPCID instruction. So thanks for catching it and letting me know, and I'm glad to correct the record.

I got a nice note from someone who's not an English speaker. Ansiotropic is his name in Twitter. And he said...

Leo: Anisotropic.

Steve: Oh, Anisotropic. Oh, thank you. I always...

Leo: It's a video gaming term.

Steve: Ah, okay.

Leo: Don't worry about it. You don't need to know.

Steve: And you're right, I was completely - because there's no way that's ansiotropic. I just never pronounced it [crosstalk].

Leo: Anisotropic.

Steve: Anisotropic. As opposed to isotropic.

Leo: It's a substance having a physical property that has a different value when measured in different directions. Wood is stronger across the grain than it is along the grain.

Steve: Or fur.

Leo: But in gaming it's a kind of, in 3D computer graphics, it's a kind of filtering, anisotropic filtering. Fur is also, yes, yes, anisotropic.

Steve: Thank you for the correction.

Leo: I think we've all learned something here today.

Steve: Well, I have. And so did this guy. He sent me an earlier tweet that I found later. But he said: "Today's update." He said: "SpinRite 6 on Windows 10 64-bit." He said:

"Okay. After a grueling 17 hours at Level 3" - which given what it did doesn't sound like it was that bad - "on normal non-problematic hard drive, 1TB, I can safely say this product is awesome." He said: "From what I could notice, it fixed all oddities in, but also out of Windows." And then he said: "BIOS resolution. TY." So I assume that's thank you. And so sort of decoding that interesting tweet, there was something going on with his BIOS, and I guess some oddities in Windows. He did a Level 3 scan on that Windows 10 system on his terabyte drive, and everything got fixed. So Anisotropic, thank you for your tweet.

And some closing-the-loop bits. S. Wayne Martin said: "Hey, can you comment on how you determine performance when your InSpectre says performance is good? Seems that would require a benchmark to determine. Curious minds." And so, yes, what I'm doing is simply using the presumed performance hit which we expect from the Meltdown mitigation when your processor doesn't have the latest instructions, that is, Haswell and later, which allow the mitigation to be performed with low performance impact. So I'm not doing a benchmark.

And in fact, in the text that InSpectre also emits, aside from the one word "good" or "reduced" or "slower" words that it uses, it explains that your mileage may vary. You're able to flip the protection on and off using the Meltdown pushbutton and then reboot in order to see if you sense a difference in performance. But all I'm doing there is essentially I'm saying that the performance is good because I'm basing that on the relative low impact of the performance that's been reported for both the mitigated Spectre and the mitigated Meltdown vulnerabilities when you've got the processor support, with or without firmware update. So just sort of ballpark.

And, frankly, I'm wanting to put a little bit of heat under Microsoft because, as we've been reporting, they're not yet taking advantage of the Invalidate PCID and PCID instructions on their older OSes. Hopefully they're planning to do that. They're settling for that performance hit, which they need not do so on newer processors. And I'd like to just keep that on everybody's radar so that we can make sure we stay aware of the fact that there is no need, there shouldn't be a need to go to Windows 10 Fall Creators Update, which apparently is the only platform where they are offering this high-performance mitigation of Meltdown, which is annoying.

Alessandro Canepa, he says: "Regarding WebMon," and he gives the URL, "doesn't seem to monitor HTTPS websites. Do you know of any alternatives?" And that's referring to the question I answered last week about the utilities I use to monitor web pages for updates. And it hasn't been updated for I think six years, during which time, as we know, websites have moved from mostly HTTP to mostly HTTPS. And unfortunately this WebMon utility, which exists outside of the browser, it's just a freestanding tool, does not appear to have been updated in order to handle secure websites. So I just wanted to correct that and make a note of that.

The add-on Check4Change plugin for Firefox, and I'm not sure if it's available cross-browser, it doesn't care whether pages are secure or not because it's just using the browser's framework in order to repull the page however often you ask it to and check the marked region for any modifications. But it looks like this little freestanding Windows app is sort of seeing its end of life because, unless it gets updated to do HTTPS sites, it's not going to be very useful. And the question that I referred to earlier, Guillermo Garcia asks, from Security Now! 644, so a couple weeks ago, "How do you 'baby' your iPhone's battery?" And you asked the same question, Leo.

What I mean by that is I am acutely aware that lithium-ion cells, while they definitely do not like to be overcharged, they really dislike being discharged. We know that, unlike the nickel-cadmium and nickel-metal hydride technologies which preceded lithium-ion, those had a severe memory effect. And so, for example, for the longest time, in the early days of portable phones, we all learned that it was better to deep cycle those batteries, running them all the way down and then putting them on the charger. Everybody got trained, that's the way to handle batteries.

It turns out that's completely the worst possible thing you could do with today's lithium-polymer, LiPo, and lithium-ion batteries. Those chemistries are killed by deep cycle discharging. And I've had several friends who have said that phones and pads die on them. And I note that their batteries are always down in the red. And I've said, "Well, yes, because you're running the battery down." They go, "What do you mean?" I go, "That kills lithium-on." You absolutely never want to do that.

So anyway, so when I say I baby my phone's battery, my lifestyle allows me to only have the battery, or the phone, off of a charger for maybe a couple hours a day, when I'm out getting a meal, or I'm driving from point A to point B. I even have a cigarette lighter plug with a lightning cord, and I dutifully keep it plugged in. So I'm very conscious about preserving the life of that battery, which is why I contend that mine's probably in really good health, despite the fact that it is a few years old. Age itself is also known to upset lithium-polymer, lithium-ion. Basically the lithium chemistry of our battery. So even if you really kept it in good shape, it would, after five or six years, just age, there is a shelf-life problem, but it is exacerbated by deep cycling. So for what it's worth - yeah.

Leo: Don't chargers do micro cycling, though? I mean, don't they, like, let it charge up and let it go down, let it charge up and go down?

Steve: Yeah.

Leo: In little increments? Does that not add up, though, to as much as a full charge and discharge in the long run?

Steve: I don't think so. I can't speak definitively. I know that I was very impressed that my Lenovo, that the Carbon X1, after I'd had it, and of course it was plugged in, for like a week or two, I remember turning it on, and I got a notice. And it said: "We notice that you seem to be having your laptop live on the AC line. If that's going to be your normal habit, we recommend, for the sake of the battery, running it down to 50 percent. And we will hold it there." And I thought, wow, that's the nicest thing I've seen. Because, I mean, that's really what you want is if there's any way to maintain it.

Again, if I were to, like, grab it and run out the door, I would have limited life use, and down at the scary end because I really don't want to hit bottom with that battery. So in fact, when I am planning a trip, I flip it off of that mode. It then brings the battery up to full charge, and off I go. So it would be nice if we had a battery technology where none of this was the user's responsibility. And unfortunately, Apple has tried to give us that. Apple has just, "Oh, don't worry about the battery. There's nothing you need to worry about it. We've got complete management." Well, but unfortunately they really don't because if the user interacts with it a little bit, as I do, for example, with my laptop, then I'm able to get much better long-term battery life.

Leo: Yeah, yeah. I mean, there are - I think there are battery monitoring programs on iOS, there certainly are on Android and computers, that can tell you what your battery capacity is compared to its original capacity, things like that.

Steve: Yeah.

Leo: But I hope Apple does the right thing and gives us that information. That would be much, much more useful.

Steve: And in fact, I'm glad you mentioned that, Leo, because they said they're not only going to give us the ability to turn off throttling, but they are going to surface their internal battery health metrics to the UI. So we'll be able to see what it is that they think the state of our battery is. Basically, they just sort of wanted to pretend this was not a consumable thing. It didn't get old. You never had to worry about it. Besides, you'd be getting a new iPhone before this became an issue. Now they're saying, okay, we're going to show you what's going on inside. So that's very cool.

Leo: I'll pass along one other note that I just saw that Ursula K. Le Guin has passed away.

Steve: Oh, wow, a major sci-fi author.

Leo: Very, very significant.

Steve: I cut my teeth on her in my youth.

Leo: Yeah. I'm trying to think of - what was your favorite Le Guin? I'm trying to think of what my favorite is. She was 88 years old.

Steve: I was just going to ask how old she was.

Leo: Yeah, 88. Nebula Award winner. Kind of a bete noire for people like Cory Doctorow because she was very actively promoting the rights of writers against people like Google Books and was always very angry about what she saw as copyright violation with Google Books. But I'm sure there'll be some good obituaries we can read about. She wrote her first published sci-fi story for Astounding Science Fiction when she was 11 years old.

Steve: I was just going to say, I was wondering how old she was because, I mean, I was reading her forever.

Leo: Yeah, a long time, yeah.

Steve: Yeah. She has a website bearing her name.

Leo: I think a lot of it was fantasy kind of focused, of less interest. "The Lathe of Heaven," that was a PBS film based...

Steve: Yeah, and as we know, I'm much more of a hard sci-fi person.

Leo: Yeah, it's much more fantasy. I do remember "The Lathe of Heaven." I think it's probably the only one I ever read. Still, a very important person.

Steve: Thank you for the heads-up.

Leo: Sure.

Steve: So a listener, Darrel McQuienn, I guess, asks, he says: "With respect to precision time, it's quite important in a lot of applications. I do," he writes, "traffic control systems and astronomy, both of which require precision time." And I got a kick out of that because of course there's precision, and then there's precision. He's referring to my suggestion that one of the most obvious mitigations for all of this is to fuzz the application's access to what's called the RDTSC instruction, the Read Time Stamp Counter. Nobody, nobody except crazy entropy harvesters or somebody doing this level, like branch cache miss cache hit resolution timing needs that kind of precision. That timestamp counter runs at the system clock rate, which is, what, 3GHz often, or 2GHz? So that's half a trillionth, is it, half a nanosecond, 500 picosecond or better timing that you just - nothing needs that.

So my point was we're not talking about, like, being off by a second or two, or a millisecond or two, or even a microsecond or two. Even if we reduced or fuzzed timing down to the microsecond, the millionth of a second, that's enough to completely thwart the detection of whether data is cached or not. That is, it's only by determining whether something is one cycle or 30 cycles at the resolution of the timestamp counter, so 30 nanoseconds. Or what is that, .03 microseconds, that these decisions are being made.

So it'll be interesting to see. I mean, I'm not - it would be better for Intel to fix the microarchitecture, I mean, the way I'm sure they will be moving forward, in order to maybe make the branch table buffer be also tagged with a process-context in order to solve the problem. I mean, I'm sure there will be lots of clever solutions. But one possibility is just not being quite so sure about what time it is. But not quite so sure, even giving applications a microsecond would have to still be fine with traffic control and astronomy, but would completely make it impossible to see what was going on closely enough to execute these attacks. And that's one 2,000th less resolution than they have now. So it would completely blur it. It'd be interesting to see if that mitigation alone would solve the problem.

And, finally, Markzip said: "@SGgrc You say that DV certificates are free and easy, so

HTTP should/will die. But what about those of us on shared servers in hosting providers who are not offering Let's Encrypt? These large providers have no plans to do so," he writes, "because it cuts in to their charging for certificates." And I will suggest that the only thing that is still allowing them to do that is inertia. There are alternative hosting providers, and there will be more, who are using Let's Encrypt and, as a benefit for luring people away from hosting providers that aren't offering free certificates, are offering free certificates.

So again, this is causing a sea change, and it will not be long before those providers who are still charging for certs are looking around at the hosting provider comparison charts, and one of the columns is "Offers free certificates." And they've got a No in their column, where everybody around them is getting Yeses. So they'll be moving to free certificates. And again, further putting another nail in the coffin of HTTP. And, boy, it just - we're at a point where you just can't do HTTP anymore. It's becoming wrong. Okay.

Leo: Time for a kitty cat.

Steve: Isn't that a neat photo?

Leo: I love it.

Steve: Yeah. That was from the EFF's page, I think.

Leo: That's a caracal.

Steve: Dark Caracal, C-A-R-A-C-A-L. And the authors explained: "In keeping with traditional APT naming" - that's Advanced Persistent Threat - "we chose the name 'Caracal' because the feline is native to Lebanon, and because this group has remained hidden for so long. From the Wikipedia entry, Wikipedia says: 'The caracal is highly secretive and difficult to observe and is often confused with other breeds of cat.'"

Okay. So for anyone who's interested, in the show notes I have a link to the full 51-page, very, very detailed readout on this, which is posted at Lookout.com. This was jointly announced research conducted by the EFF, our friends at the Electronic Frontier Foundation, and Lookout. It describes with as much detail as you could ask, basically they captured all of these bits of malware and reverse-engineered them so that they're showing the command vocabulary and the response vocabulary that these different pieces of malware provide. They track down the specific building in Beirut, the General Directorate of General Security, GDGS, where this is all being managed. They found three or four personas, pseudonyms under which domains and hosting is provided and email is being routed through. And then from those pseudonyms they were able to find other instances of those on the web, and basically assembled this spider web of thinuous - thinuous.

Leo: I like it. It's a new word.

Steve: Thinuous. That's a thinuous connection - of tenuous connections in order to pull together the whole picture. And really, I mean, this is the way this kind of research goes, where you just pursue each lead. You gather the information. You record it. And then something else, you realize, ooh, wait a minute, that's using the same port as this. And then you go over to Shodan, and you do a scan of that port, and oh, that finds a bunch of other machines. And then you go look at them, and you say, hey, wait a minute, those IPs have the following provider in common, and that then tells you what provider is being used. And then you go look at the provider. And from that perspective, looking out, you can see more than you would have known to find before. And so that's the way this goes.

So this report pulls this all together, all of these disparate leads - no matter how thinuous they were when they were initially found - and builds a picture. So they find this Dark Caracal network which is, interestingly, reusing the same infrastructure - the command-and-control servers, IP addresses, hosting and database providers, domain registrars because they've got the whole bunch of spoofing domains and waterhole domains designed to trick people into not noticing, oh, wait a minute, where I'm getting my Flash player from is not actually Adobe, but something that looks sort of the same, unless I looked closely.

This infrastructure had been previously seen in another similar campaign which was targeting journalists, lawyers, and dissidents - so a little more on the political side - who were critical of the government of Kazakhstan. The Dark Caracal effort has been conducting a multiplatform Advanced Persistent Threat. And that's a term that I think maybe we first used after the Sony revelations, the idea that there was an APT that had gotten into Sony Pictures' network and set up shop and over a long period of time was performing surveillance from inside their network.

So this is an Advanced Persistent Threat surveillance operation, which they have verified now is targeting individuals and institutions across the world, not just focused on political targets. They have located and rummaged through hundreds of gigabytes of data identified as having been exfiltrated from thousands of victims spanning more than 21 countries, including North America, Europe, the Middle East, and Asia. And in their report they have a map that sort of gives a sort of a geographic sense of what the spread is. Lots in the Middle East and Asia. There's a mobile component of this Advanced Persistent Threat which is one of the first they have ever seen, which was executing global espionage at this scale.

Lookout, which is the partner of EFF here, focuses on mobile malware and mobile attack analysis. They found that, in looking at the connections and the data, that this Dark Caracal had successfully compromised the devices of military personnel, enterprises, medical professionals, activists, journalists, lawyers, and educational institutions; that it was targeting governments, their militaries, utility infrastructure - utilities, power, gas, electric and so forth - financial institutions, manufacturing companies, and defense contractors. So it's got its tentacles everywhere. They found exfiltrated data, including documents, call records, audio recordings, secure messaging client content, contact information, text messages, photos, and account data. So again, pretty much everything that was of interest to this kind of cyberespionage.

They also were able to piece together what the typical attack chain looked like, that is, from not having infiltrated a target, how does it proceed? And they found that the attack chain that they typically see is what's underway, relying primarily on social media and phishing. In some cases physical access is obtained to target systems, devices, and accounts. But that's far and away the rarity. Generally they're learning about their target. They're figuring out who they are, what interests them, where they spend their time, and just doing that from looking at what's available in social media. And then they began spearphishing, trying to lure them to a watering hole or to some phishing sites in order to begin the process of establishing a beachhead in their hardware.

Some of Dark Caracal's espionage technology appears to have been developed in-house. And in fact there is a cross-platform Java-based tool that they called CrossRAT, R-A-T as in Remote Access Trojan, CrossRAT as in cross-platform Remote Access Trojan, written in Java. It's the first time they've encountered that, appears to have been written by this organization. And then there's other technology which is either borrowed from or purchased from the dark web. A version of FinFisher that we've talked about in the past has also been used as part of this campaign.

The guys at Lookout first discovered the presence of something they named Pallas, P-A-L-L-A-S, which is their name for an implant used in multiple trojanized Android applications. They spotted that last summer, in May of 2017. And they have identified it in 11 different Android applications, including - and here's what's a little bit freaky - Signal, Threema, Primo, WhatsApp, and Plus Messenger, which is the Telegram messenger. What they do is they, I mean, this group is good. So they're taking those apps and trojanizing them, essentially installing Pallas into an application. It remains fully functional as what it does.

So, I mean, if you think you downloaded Signal, but you didn't get it from the Google Play Store, you got it from a spoofed version of that store. And so it operates, it functions, it's the Signal app in every way, except you're also carrying a trojan, which sets up communication with their command-and-control network, allows it to exfiltrate any data on your phone, enumerate the hardware devices, like how many cameras you've got, and turn on the microphone and surreptitiously record and send the data out of your phone. And that's the case for Signal, Threema, Primo, WhatsApp, and Telegram.

So, I mean, this is like state-level surveillance, cyberespionage. It also makes extensive use of a Windows malware known as Bandook, B-A-N-D-O-O-K RAT, the Bandook RAT, and employs a continually evolving and changing global network infrastructure. Many of the threads that they pulled on they found led to abandoned IPs or abandoned servers, which were at one point in use, but had been rotated away from. So this suggests to them that there's some money behind this; that doing that, maintaining a constantly evolving footprint is definitely more secure, but also significantly more expensive. You have to have a whole separate team that is evolving this entire structure continually.

They found that the command-and-control servers generally preferred the use of Windows and the XAMPP stack, which is the cross-platform stack, rather than the more traditional LAMP stack that is Linux-focused. And their research produced more than 90 what they called Indicators of Compromise: 11 Android malware IOCs, these Indicators of Compromise; 26 desktop malware IOCs. And so, for example, things like this Java-based malware for cross-platform compromise that affects Windows, Mac OS X, and Linux platforms; and then also 60 domain names, IP addresses, and WHOIS information that they all dredged up during this research.

They tracked down WiFi networks and SSIDs, the IP addresses these guys were using. It turns out that they were almost all using one particular hosting provider. I didn't have it in my notes here, but it is in the documentation, if anyone's curious. And it's apparently, it looked like it was an Asian provider that gave absolutely no concern whatsoever for what content was being hosted on their servers. As long as you paid the bills, they were open for business to you. And very difficult for law enforcement to have any access to.

There are, and they found and have screenshots of, fully mature watering hole servers which are offering these malicious versions of these well-known messenger apps that sort of tended to be their focus, as well as phishing domains which closely mimic

Facebook and Twitter. So you're following a link from a message that you receive. You think you are going to Facebook or Twitter. Unless you are really paying attention, that's not where you are. You then log in, and they essentially are then leveraging your trust in your belief at where you are for everything that you're presented and the links that you then get. And from that point on, they say, oh, you need to update your Flash player in order to view the latest content that we have. Click here, and we'll take you to Adobe. And in fact they take you to an Adobe clone site, where you download, obviously, whatever it is that is malicious and that is now inside your system.

So it relies primarily, that is, as I mentioned before, getting a foothold primarily relies on social engineering via posts in a Facebook group and WhatsApp messages in order to compromise target systems, devices, and accounts. And they wrote that, at a high level, the attackers have designed three different kinds of phishing messages, the goal of which is to eventually drive victims to a watering hole controlled by Dark Caracal. Interestingly, neither the desktop nor the mobile malware tooling use zero-day vulnerabilities. That isn't the way they get in. They're not leveraging them or assuming them. They're simply downloaded instead of the intended application, and then they rely on whatever permissions have been granted at the time of the installation in order to give them access to sensitive user data.

So anyway, it is a - oh, the one chilling thing that I said I mentioned, at one point they were enumerating what the Windows client malware was obtaining. In addition to full transcripts of all the messaging that's being done through the desktop messaging, they get screenshots of the desktop. And at one point they said this data included full screenshots taken at regular intervals and uploaded to AdobeAir[.]net. And the authors wrote: "By observing these images, it is disturbingly simple to watch a victim go about his daily life and follow that individual every step of the way."

They wrote: "Not only was Dark Caracal able to cast its net wide, it was also able to gain deep insight into each of its victims' lives. It did this," they wrote, "through a series of multiplatform surveillance campaigns that began with desktop attacks and pivoted to the mobile device. Stolen data was found to include personal messages and photos, as well as corporate and legal documentation. In some cases, screenshots from its Windows malware painted a picture of how a particular individual spent his evenings at home."

So anyway, this is obviously, as I said at the top of the show, this is not something that is affecting us individually at the level that Spectre and Meltdown do. But this is a look inside a mature, nation-state level, well-financed, cyberespionage campaign. This of course is not an outlier. We have to assume this is what all significant nation-states, and even insignificant ones, are probably doing to varying degrees and varying levels of success.

Unfortunately, this is the downside to this amazing global network and PCs and mobile devices and the free availability of software that does useful things. The flipside is it can all be turned against us. And as individuals, I don't think any of us are probably targets or victims. But it's worth being aware that, when targeted, and we've talked about this before, it is very, very difficult to never make a single mistake. All it takes is a single wrong click on something that you have been given every reason to believe is safe. And you can then have something installed on your system that will remain as an APT, an Advanced Persistent Threat, moving forward, and give somebody you don't intend to have given access to essentially every detail of your life.

Leo: So it's coming out of Lebanon. Does that mean that's - who made this, you think?

Steve: It's not clear.

Leo: Lebanese? I mean, I don't think this is a - is there a Lebanese secret police that has interests around the world? I mean...

Steve: Well, I would argue that any large country does now. I mean, we know that North Korea has been bragging about their cyberattack people. We know in the U.S. we have a whole cyber program that is spun up and running. We know that China is active. And we know that Russia...

Leo: It doesn't necessarily have to be coming from Lebanon to be, I mean, they pinpointed Lebanon, but that doesn't mean that's who's originating it.

Steve: Well, they pinpointed the intelligence services building in Beirut. So, yeah. I think in this case they were stopping just short of providing - it wouldn't be until you got a formal declaration of, yes, we're doing this, and here's our annual budget for it.

Leo: Right. Responsibility, yeah. But then in that case it sounds like something that was targeted perhaps at Lebanese citizens that just kind of got loose? Or maybe somebody got the code and is exploiting it?

Steve: No. They're, I mean, if you read this paper, I mean, it's all there. They have a global network targeting more than 21 countries where they're actively surveilling individuals across the globe - in military and government, in power utilities. I mean, this is the nature of the world we're in today.

Leo: Okay. It's interesting that GDGS somehow has some interest in what's going on in some distant country. Just surprises me a little bit.

Steve: Yeah. I just think that's the nature of spooks, national spooks. They're like, oh, yeah, we have our hooks into Pacific Gas & Electric in California, just in case that's of some use to us.

Leo: Yeah. Interesting. Okay. Well, there you have it, the Dark Caracal. Steve Gibson is at GRC.com. That's where you can get InSpectre, his free utility to check the status of your Meltdown and Spectre mitigations and the impact. And there's lots of information in the app, and it's free. In fact, he's got a ton of free stuff.

But while you're there, check out SpinRite, the one thing, the one thing he charges you for. Might be the most useful. It's the world's best hard drive maintenance and recovery utility. You'll find that there. Also copies of this show, audio and transcriptions, handcrafted by Elaine Farris: GRC.com. If you want to leave a question or a thought or a comment for Steve, it's on the Twitter at @SGgrc. And he accepts DMs from strangers. I've told him to stop it, but he persists. Nevertheless, he persisted. Tips and all sorts of stuff coming that way: @SGgrc.

We have audio and video of the show if you choose to see our smiling faces at TWiT.tv/sn. You can watch us do the show live every Tuesday, 1:30 Pacific, 4:30 Eastern, 21:30 UTC. You can be in the studio audience if you email us: tickets@twit.tv. We've got a couple of nice fellas here, Chris from New Waverly, Texas, and Mark from Oakland, California. Welcome. Good to have you both.

If you can't be here live, if you can't listen live, you can always get on-demand audio and video from your favorite podcatcher. Just subscribe to Security Now! and complete your set. You need all 647 episodes to really say - we get emails, I'm sure you do, too, every week from people saying, "I just completed listening to all the back episodes."

Steve: Yeah.

Leo: That's a major effort.

Steve: That is a labor of love, they call that.

Leo: Forty-three days, three hours, 20 minutes, and one second. And counting, according to [crosstalk].

Steve: Of nonstop audio.

Leo: Nonstop audio.

Steve: And counting. Yay.

Leo: Thank you, Steve. It's great to see you again, and we will see you all next week.

Steve: Thanks, buddy.

Leo: On Security Now!. Now you can say it.

Steve: Thanks, buddy.

Leo: Take care.