TETRA:BURST

Description: It turns out that Advanced Persistent Threats have been leveraging satellite communications for many years. We start by looking at that. Then we'll find out what the next iOS release will be doing to further thwart device tracking. What new feature is Android 6+ releasing? What's the latest on the forthcoming seventh branch of the U.S. military? Why has Russia suddenly criminalized open source contribution? And what do we learn from VirusTotal's 2023 "malware-we've-seen" update? Then, after we share some of the terrific podcast-relevant feedback received from our amazing listeners following last week's second satellite insecurity podcast, we're going to examine one of the revelations to be detailed during next week's Black Hat hacking conference in Las Vegas.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-933.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-933-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk more about satellite communications. We've got some really expert listeners, some fascinating insights into that. We'll also talk about Russia. They've actually criminalized open source contribution. And then VirusTotal's 2023 malware-we've-seen update. Plus a look at a radio solution used by law enforcement all over the world that is woefully insecure. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 933, recorded Tuesday, August 1st, 2023: TETRA:BURST.

It's time for Security Now!, the show where we talk about your security, your privacy, your health and welfare online with this guy right here, Mr. Steve Gibson of the Gibson Research Corporation. Hi, Steve.

Steve Gibson: Hello, Leo. Great to be with you for this first day of August.

Leo: Yeah.

Steve: We're in the low 80s here, so we're in, like, paradise compared to the rest of the country. It's either thunderstorms and tornadoes, and my sister posted something on Facebook with like some huge lightning storm she was in in Colorado yesterday or last evening. And of course Arizona's breaking records at 110 degrees or higher for more days in a row than they've ever had. And here we are, it's a little humid, but otherwise it's great.

Leo: Really. That's nice. You don't get a lot of humidity, do you? You're in an arid clime.

Steve: No, it's odd. It's odd for us, yeah. Some El Nino thing happening.

Leo: Oh, yeah.

Steve: And I'll just note that we're two weeks away from finishing our 18th year of this podcast.

Leo: OMG. 18 years. Wow.

Steve: Yeah, I think it's August 18. Or, yeah, I think it's - I think it is August 18 is our 18th, is the end of year 18. We'll begin into 19. So wow. Very cool. So it turns out that Advanced Persistent Threats have been leveraging satellite communications for many years. So before we wrap up all of our staring-at-the-heavens discussion, we're going to look at that.

Leo: We should, and I hope you will, and I know you will, talk about what - it's a terrible name, Advanced Persistent Threat, for what it really is.

Steve: Yeah.

Leo: But I'm sure you'll explain that.

Steve: Yeah, yeah. Also we're going to find out what the next iOS release will be doing to further thwart device tracking. And I know you touched on that in your previous podcast on MacBreak Weekly, and also what new feature Android 6 is releasing. But you also cast some doubt on whether that was happening on MacBreak Weekly.

Leo: Yeah.

Steve: So I'm kind of curious to see whether we're talking about the same thing. Also we've got some news on the latest forthcoming seventh branch of the U.S. military. And we're going to wonder why Russia suddenly criminalized contributions to open source software, which is bizarre.

Leo: Can you believe that? Ugh. Because they don't control it; right?

Steve: Yeah, huh? And what do we learn from VirusTotal's 2023 "malware-we've-seen" update? Then we're going to share - we got an amazing amount of terrific podcast-related feedback from our astonishingly varied listeners. We've got more people who

know about satellite security, it turns out, who had some interesting stuff to add to our discussion last week, which I'm going to share.

And then we're going to examine one of the revelations to be detailed during next week's upcoming Black Hat hacking conference in Las Vegas, thus the title of today's podcast, TETRA:BURST. It turns out that when Europeans design a secure radio protocol that has four different encryption algorithms, which they allocate to different countries, you've got to say, "What? Why?" So everybody gets to use TEA1 except the European military gets to use TEA2. Don't you kind of want to wonder why?

Leo: You get an algorithm, and you get an algorithm, and you and you and you. Wow.

Steve: That's right. And they're all secret and unpublished.

Leo: Oh, that's not good.

Steve: And it turned out that only by leveraging some zero-days in a Motorola implementation of this encrypted handset were a bunch of guys in the Netherlands, researchers in the Netherlands, able to crack the Secure Enclave to for the first time ever get access to these proprietary encryption algorithms. And, oh my god, would you believe that they're not secure? So anyway, we're going to have fun today.

Leo: Why is it that people roll their own? I mean, it's not like the Enigma machine that it's security through obscurity; right?

Steve: Well, unfortunately, it's an attempt at that.

Leo: It is, yeah.

Steve: The only way I can give them, well, I would give them an out is to say that it's decades old. So this exists from the '90s. But, I mean, I'm giving away a lot of the podcast coming up.

Leo: All right, save it, save it.

Steve: But believe it or not, they've replaced them, having been caught, with new proprietary secret algorithms.

Leo: Now there's no excuse. If there's one thing we know about crypto, it's got to be open, which is exactly why Russia doesn't like open.

Steve: Yeah.

Leo: Now let's build the Picture of the Week, shall we, Steve?

Steve: Okay. This is just - it struck me as funny. I've got a killer one coming for next week, but this was repurposing an old photo with a wonderful caption that has got nothing to do with security. But I just got a kick out of it.

Leo: Okay. It's going to take me a second because I did not have it on this computer. I have two different computers, and I've got to make sure I get the one with the Picture of the Week. The good news is I haven't seen it yet.

Steve: That's correct.

Leo: So you will see my genuine reaction.

Steve: First reaction. It's just cute.

Leo: We'll do it together [laughter]. Okay. Okay, Mr. Spock. Never play with super glue.

Steve: Never play with super glue. This just shows Spock doing his Vulcan hand sign with the two fingers stuck...

Leo: It does look like they're glued. They might be stuck, yeah. Like he was trying to wave.

Steve: It's funny because my best friend at the time, a guy named Gary Rawlings, was my best man in my first wedding. And I said, "Rawlings, do not embarrass me. You're going to do the best man speech. Whatever you do, you know." Because, I mean, he knew where the bodies were buried, times, you know, to the power of infinity. And so he was very dangerous to have up there onstage. And he had like kind of a dry sense of humor where he could really go too far. So, I mean, I put the fear of god into him. And so he got up, he received the microphone, and he held his hand up. And he said, "Gibson told me I was forbidden for saying anything really that would embarrass him. So I'm just going to say 'Live Long and Prosper.'" Now, Gary could not do the Vulcan hand sign. So he had rubber bands around his...

Leo: Oh, no.

Steve: Around his fingers.

Leo: That's hysterical.

Steve: In order to make them do that. So anyway, I thanked him.

Leo: I could do it with one hand but not the other. I could do it with my non-dominant hand.

Steve: I'm clearly a double - I'm a double Vulcan. I can, yeah, I can...

Leo: I guess it takes practice.

Steve: ...animate them and do whatever they need to do, yeah.

Leo: Okay.

Steve: Okay. So before we wander away, as I said, at least for the time being, from the topic of satellite security - which turns out it's a rich field, I mean, there's been generated a huge amount of interest among our listeners. So I'm glad that we spend some time talking about this last couple weeks. I want to talk about another aspect of the use of satellites by bad guys, which again I wasn't aware of, but makes sense when you think about it, the deliberate routing of Internet connections through space. This is done as a means of thwarting the persistent efforts by law enforcement to track down, shut down, and sometimes take over the command-and-control servers and infrastructure which is being used by the major advanced persistent threat groups. Since it's another thing that we've never explicitly covered, I thought that now, while we're still looking skyward, would be a good time to add this to the growing list of things that we have covered.

So way back in September of 2015 - this is not news, this is eight years ago - Kaspersky published an informative research piece titled "Satellite Turla," T-U-R-L-A. Turla is the name of an APT group. And so their title was "Satellite Turla: APT Command and Control in the Sky."

Leo: What is an APT? Can you, I mean, I know it's resident, basically, a resident infection. Yeah?

Steve: I think the first time we encountered it on the podcast was when one was discovered at Sony Entertainment.

Leo: Right. They were wandering around for months inside the Sony systems.

Steve: Like a long time, yes.

Leo: Yeah, yeah.

Steve: Yes. And so that may have been where this notion of, you know, so Advanced Persistent Threat. Advanced, obviously, means it's not some script kiddie doing, you know, up to nonsense. This is a serious, a serious organizational intrusion. Persistent meaning that, you know, again, it wasn't something that was executed and then died. It established a foothold in some sort of corporate asset, and from there it then was used

for surveillance over some long period of time. We've seen printers for example, you know, no one would think of a printer as being a computer, but of course they are. And their firmware is no more secure than anything else, unfortunately, these days. And so we've seen APTs that set up shop in printers where, as I said, no one thinks to look.

Leo: Right.

Steve: And then from there, you know, they're on the network, so they're able to go out and see what's going on. So all of these things need some means of phoning home in order to report the things that they have found and also, as we'll see, to create a means for allowing the bad guys back in over time. So Kaspersky in their write-up stopped short of explaining the detailed network packet flow. But they did provide enough for us to fill in the rest of the technology. So first, I've skipped over some of their warm-up introduction, which would be redundant for our audience, but I want to sort of create the background that they did create, and then we'll figure out how the packet flow works.

So Kaspersky said: "When you are an APT group, you need to deal with many different problems. One of them, and perhaps the biggest, is the constant seizure and takedown of domains and servers used for command-and-control. These servers are constantly appropriated by law enforcement or shut down by ISPs. Sometimes they can be used to trace the attackers back to their physical locations. Some of the most advanced threat actors or users of commercial hacking tools have found a solution to the takedown problem, the use of satellite-based Internet links." And again, this is in 2015. So this has only matured since then. "In the past," Kaspersky wrote, "we've seen three different actors using such links to mask their operations. The most interesting and unusual of them is the Turla group.

"Also known as Snake or Uroboros, names which come from its top class rootkit, the Turla cyberespionage group has been active for more than eight years." And that was more than eight years in 2015, and they're still a name that's around, so they've been at this for a while. Kaspersky said: "Several papers have been published about the group's operations, but until recently little information was available about the more unusual aspects of their operations, such as the first stages of infection through watering-hole attacks.

"What makes the Turla group special is not just the complexity of its tools, which include the Uroboros rootkit, aka 'Snake,' as well as mechanisms designed to bypass air gaps through multi-stage proxy networks inside LANs, but the exquisite satellite-based command-and-control mechanism used in the latter stages of the attack.

"In this blog, we hope to shed more light on the satellite-based command-and-control mechanisms that APT groups, including the Turla/Snake group, use to control their most important victims. As the use of these mechanisms becomes more popular, it's important for system admins to deploy the correct defense strategies to mitigate such attacks. For IOCs" - remember, indications of compromise - "see the appendix.

"Although relatively rare, since 2007 several elite APT groups have been using and abusing satellite links to manage their operations, most often their command-and-control infrastructure. Turla is one of them. Using this approach offers some advantages, such as making it hard to identify the operators behind the attack, but it also poses some risks to the attackers. On the one hand, it's valuable because the true location and hardware of the command-and-control server cannot be easily determined or physically seized. Satellite-based Internet receivers can be located anywhere within the area covered by a satellite, and this is generally quite large. The method used by the Turla group to hijack the downstream links is highly anonymous and does not require a valid satellite Internet

subscription. On the other hand, the disadvantage comes from the fact that satellite-based Internet is slow and can be unstable.

"In the beginning, it was unclear to us and other researchers whether some of the links observed were commercial Internet connections via satellite, purchased by the attackers, or if the attackers had breached the ISPs and performed man-in-the-middle attacks at the router level to hijack the stream. We have analyzed these mechanisms and come to the astonishing conclusion that the method used by the Turla group is incredibly simple and straightforward, as well as highly anonymous and very cheap to operate and manage.

"Purchasing satellite-based Internet links is one of the options APT groups can choose to secure their command-and-control traffic. However, full duplex satellite links can be very expensive." Now, this is in 2015. "A simple, duplex, 1Mb up/down satellite link may cost up to $7,000 per week. For longer term contracts this cost may decrease considerably, but the bandwidth still remains very expensive." And again, this is back in 2015, so things may have changed since.

"Another way of getting a command-and-control server into a satellite's IP range is to hijack the network traffic between the victim and the satellite operator, and to inject packets along the way. This requires either exploitation of the satellite provider itself, or of another ISP on the way," you know, in line. "These kinds of hijacking attacks have been observed in the past and were documented by Renesys, now part of Dyn, in a blog post dated in November of 2013, so two years before this one was written in September of 2015.

"According to Renesys: 'Various providers' BGP routes were hijacked, and as a result a portion of their Internet traffic was misdirected to flow through Belarusian and Icelandic ISPs.' They said: 'We have BGP routing data that show the second-by-second evolution of 21 Belarusian events in February and May of 2013, and 17 Icelandic events in July through August of 2013.'

"In a more recent blog post from 2015, these researchers point out that: 'For security analysts reviewing alert logs, it is important to appreciate that the IP addresses identified as the source of incidents can and are regularly spoofed. For example, an attack that appeared to come from a Comcast IP located in New Jersey may really have been from a hijacker located in Eastern Europe, briefly commandeering Comcast's IP space. It's interesting to note that all six cases discussed above were conducted from either Europe or Russia.'"

Okay, now, they write: "Obviously, such incredibly apparent and large-scale attacks have little chance of surviving for long periods of time, which is one of the key requirements for running an advanced persistent threat operation. It's therefore not feasible to perform the attack through man-in-the-middle traffic hijacking, unless the attackers have direct control over some high-traffic network points, such as backbone routers and fiber optics." And of course that's unusual, too. They said: "There are signs that such attacks are becoming more common, but there is a much simpler way to hijack traffic-based Internet traffic.

"Enter satellite link DVB-S hijacking." They said: "The hijacking of satellite DVB-S links has been described a few times in the past, and a presentation on hijacking satellite DVB links was delivered at Black Hat in 2010 by an S21Sec researcher. To hijack satellite DVB-S links, one needs the following: a satellite dish - the size depends on geographical position and satellite; a low-noise block downconverter," typically called an LNB, and that's generally part of the satellite that you get mounted on your roof if you're subscribing to Dish network or whatever. You also need "a dedicated DVB-S tuner" - which takes the form of a PCIe card these days - "and a PC, preferably running Linux."

They said: "While the dish and the LNB are more-or-less standard, the card is perhaps the most important component. Currently, the best DVB-S cards are made by a company called TBS Technologies. The TBS-6922SE is the best entry-level card for the task." And that can be had for about a hundred bucks. "The TBS card is particularly well-suited to this task because it has dedicated Linux kernel drivers and supports a function known as 'brute-force scan' which allows wide-frequency ranges to be tested for interesting signals. Of course, other PCI or PCIe cards might work as well, while in general the USB-based cards are relatively poor and should be avoided.

"Unlike full duplex satellite-based Internet, the downstream-only Internet links are used to accelerate Internet downloads and are very inexpensive and easy to deploy. They're also inherently insecure and use no encryption to obfuscate the traffic. This creates the possibility for abuse."

Okay. So Kaspersky's article, as I said, did not go into any more detail about how this works. They switched to providing tables of IP ranges that had been observed in the past and noted the satellite Internet service providers that were using those ranges. But fortunately, we have all the information we need to understand the advantage this gives to anyone who is attempting to hide their command-and-control server. The key is that these Internet communication satellites have extremely broad coverage areas, coupled with the fact that, just like the Internet, the IP packet traffic being carried is not, itself, encrypted. As we know, TCP and UDP are not encrypted protocols. They're just carriers of data that today is typically encrypted. That is, the data they're carrying is encrypted, but they themselves, the actual underlying protocol is not an encrypted protocol.

Okay. So imagine that some nasty advanced persistent threat malware has been surreptitiously placed into a high-value computer, and that more than anything the bad guys do not want their command-and-control infrastructure - which this malware will be reaching out to, to receive instructions and updates and things - to be discovered, commandeered, and shut down. Presumably this APT threat group has many such infestations which are all reusing the same infrastructure. So the loss of that command-and-control server would cripple the entire network that they had established.

Okay. So they have their APT malware periodically send a UDP packet to the IP of a previously chosen customer - probably a big stable customer - of a given satellite-based Internet provider. Having the malware send an outbound UDP packet has the effect of opening up return paths through any NAT routing and firewalls that would otherwise prevent unsolicited traffic from entering the enterprise's network and reaching the malware-laden machine. So you want the malware to initiate communications, which actually works in favor of this whole architecture.

So this UDP packet is sent out to a previously selected customer of a satellite ISP. So it will be received first by that ISP. So this is a block of the ISP's IP space. It comes to the ISP. But unlike other ISPs, the received packet is beamed upstream directly at a chosen communications satellite. This causes it to then be rebroadcast out across the entire coverage area of the satellite indiscriminately. Somewhere, down on the ground, is that subscriber of that Internet ISP.

But also somewhere else, anywhere else within that satellite's large coverage area, the malicious command-and-control server is silently lurking with its own satellite dish passively aimed up at the ISP's broadcasting satellite. It patiently listens for any UDP packets addressed to that IP. Since the subscriber will likely have their own NAT router or firewall that will simply ignore any unsolicited nonsense, as everything has to these days, and since that subscriber may have been pre-selected to make sure that that's true, their receipt of that incoming packet will be ignored; right? It's just a radio packet coming in on their satellite dish.

But it will be what the malicious command-and-control server base station has been waiting for. Upon receiving that UDP packet, the base station can reply by sending its own UDP packet via terrestrial ground Internet, since there's no need for it to be returned to space. Right? It's just an IP packet. So they can drop it on any IP connection and it'll find its way back to the original malware that initiated it. This allows the command and control system to send whatever commands it may wish back to the querying machine. And the traffic doesn't need to only be UDP. It's just sort of easier for this example. But nothing prevents the listening command-and-control base station from establishing a three-way handshake and bringing up an encrypted TCP connection.

The key to the hack is that it's the world's largest air gap. The outbound traffic is being sprayed over a huge geographic area, to be picked up by a totally passive satellite dish that there's no way of locating, and it could be anywhere. It could even be mobile, you know, within the range of the satellite. And the command-and-control system's IP address being used is someone else's, not theirs, and so it's an air gapped man-in-the-middle traffic interception attack that is going to work and prevent the command-and-control server from ever being discovered. So unfortunately you have to give the bad guys some credit for this hack. It's pretty slick.

Okay. So Apple just updated its developer program to further crack down on developers who are abusing some of its API features which are being used to collect data on user devices. And they're doing that as an underhanded means of tracking them online. Apple said that even if a user has given an app permission to track their activity, fingerprinting the underlying device is still not allowed, yet it is still going on.

So with the release of iOS 17 and macOS Sonoma this fall, developers who want to continue to have access to these features, which could and have been used to enable persistent device-level tracking, are going to have to provide a valid reason to Apple for having that right. Apps that don't provide a good reason will not be accepted on the App Store as soon as iOS 17 rolls out, and Apple begins to enforce this policy. And Leo, I'm astonished by the apparent value added by this tracking. I mean, tracking, it must be that it provides so much more benefit to advertisers above and beyond just putting their ad on a page where it makes sense for their ad to appear.

Leo: Well, I have strong opinions about this. Advertisers think it provides value. There's a lot of evidence that personalized ads don't in fact work better.

Steve: I agree.

Leo: Yeah. But if you're an advertiser, think about it. You would, I mean, there's a famous saying that I know that half of my ads work, I just don't know which half. They would love some idea that they're hitting an audience that's interested in buying, for instance. They haven't been able to, you know, on TV you really can't do that if you buy network television.

Steve: Right, right.

Leo: That's why it's mostly brands on network television. They know, well, we're enhancing the brand, Pepsi or Budweiser.

Steve: Right.

Leo: And so that million dollar ad on Super Bowl is worth it. But for podcasts and websites and a lot of the digital world...

Steve: We have some targeting.

Leo: You can target. You know, Facebook and Google live on this. And it makes them feel better. I don't know if, you know, there's a third category of advertising which is the advertising we do, which is called Direct Response advertising. That's why we always have a URL.

Steve: Right.

Leo: Or, you know, on late-night TV you'd see an 800 number or an offer code. That's another way of an advertiser kind of reassuring themselves that their advertising is working. They're all imperfect. And all the studies I've seen say that tracking is not a very effective way of, you know, that targeting your ads doesn't really make that much of a difference. But advertisers believe it. And maybe even if they know better, the agencies need something. They're grasping at straws.

Steve: Maybe what's happened is that this is all to support that sketchy data broker business.

Leo: No, no, no, I don't - I think that's a wonderful side business for the companies that sell it. But remember Google doesn't sell the data brokers space because they sell the data brokers.

Steve: So they're doing it for their own purposes, only internal.

Leo: Well, they're doing it because advertisers demand it. I mean, that's why we do it. We do very limited tracking. As you know, a podcast, it's impossible to know with RSS feed anything but the IP address of the visiting computer. And we don't do more than that. But we do use services, a variety of different services. Right now we're using a thing called Podsights that they're an independent third party. We send them the IP addresses of people who listen to Security Now!. And the advertiser sends them the list of IP addresses of people who visited their site. The third party goes, okay, 32% of the people who heard the ad visited your site. They don't give the information to the advertiser. We don't get the advertiser's information.

Steve: So there's no matching of IPs.

Leo: It's only done by the third-party in a private way. And, even that I resisted. But honestly, we would not be able to sell advertising because - and that's the thing. The advertisers, I don't think it's...

Steve: They're spoiled; right? They've gotten spoiled.

Leo: It's not even they're spoiled. They just - they have a faith, a firmly held belief that this information helps them. And they refuse to buy - they'll only buy ads where they can get that information. Frankly, we're lucky. We have a hard time selling ads against people like Google and Facebook, who will say, I can give you 25-year-old to 30-year-old men in Petaluma, California. Would you like that? Or I can give you people with income over $100,000 who live in the Northwest. Would you like that? We can't do that. You know, all we can say - so we are losing, frankly. We're losing out to Facebook and Google, which have about 88% of all the online ad sales, because they offer that kind of information. So they're going to keep doing it.

And, you know, you see Google doing all sorts of maneuvers to get us to trust them. So they've turned off cookies, and we were talking the other day about this new web integrity initiative that they're proposing, they're going to build into Chrome. That's just one more way of them knowing who's there. And advertisers insist on it. So they think they have to do it. Whether anybody believes in it working, I don't know. But they think they have to do it. We have to do what we have to do, or we would have zero advertisers. As it is, we lose a lot of ads because we can't give them, you know, people just go, well, I'm going to buy Facebook, or I'm going to buy Google.

Steve: Wow.

Leo: So some of our advertising, not so much on this show, but some of our advertising is now direct insertion where we use a company, LibSyn company called AdvertiseCast. We just started doing this. And we pause, put a little trigger in there, and they stick in an ad. And those, advertisers like them a little bit better because they can geographically target. Your IP address has a rough geographic location. So when one of our shows airs in Spain, for instance, a Spanish advertiser will buy that, knowing that, well, I'm only - because they don't want to buy U.S. listeners because they're not customers.

Steve: Right, right.

Leo: So they'll have an ad, and they say, well, we know these are the people who listen to the show in Spain. Here, you can have that. So it's another form of targeting. But, you know, advertisers, they demand it. And if you're an ad-supported media company, you have to find a way to balance your - we believe in our community, and especially your listeners, they don't want to be tracked.

Steve: No. And in this case, so we have Apple who's trying to thwart the surreptitious, underhanded device tracking. They haven't...

Leo: But they have all those ads themselves, this information themselves. So they have first - this is what we were talking about earlier, as you heard on MacBreak Weekly, is first-party tracking, like Facebook and Google and Apple do. And of course what they're really saying is we want this to ourselves. We don't want some app on your phone to have the information we have, the relationship with the customer we have. They're not saying we don't want advertising, we don't want to track you. They're saying, we don't want them to track you so we can. It's our advantage. So I'm a little - I'm cynical about this whole thing.

Steve: Yup. Well, and there's a different form of tracking that you also touched on, and that's a more deliberate form, and that's back into the deep dive that we took a couple months ago on AirTag tracking technology.

Leo: Oh, yes.

Steve: As we know, this AirTag tracking technology is Bluetooth based, so it's inherently crowd-sourced. So this of course relates to the Apple and Google agreement. It's in both parties' interests, Apple and Google, to have a single common standard which they share so that both Apple and Android handsets can provide the tracking location feedback for each other's ecosystems. And so what they announced when we talked about this a couple months ago was a joint specification. But it was really indistinguishable from what Apple had already been doing for several years with their AirTags. So what appeared to have actually happened was that Apple had opened their specification for Google, and Google was happy to take it because, you know, they already had an established ecosystem, and then people would be able to use their Android phones as track feedback devices, as well. So it's good for everybody.

Last Thursday's news is that Google would soon be adding "Unknown Tracking Alerts" to Android. They said in their announcement, "Unknown tracker alerts, which we announced at I/O 2023, are beginning to roll out in Android 6.0+ users this month." And they also said, "Unknown tracker alerts currently work with Apple AirTags." And of course other third-party tags. They said: "We'll continue to work with tag manufacturers to expand this important protection to other tracking tags over time through our joint industry specification." Now, what you had seen was story that said that was going to be put on hold.

Leo: Yeah.

Steve: To the end of the year, which, I mean, and so I'm not sure if unknown tracking alerts, I mean, that's only - unknown tracking alerts is one aspect of the whole AirTag tracking. The other side being, you know, you own AirTags, and your device is telling you where they're located. So that's different than being aware of an AirTag that is traveling with you. So maybe we're talking about two different things, or maybe we're talking about everything being on hold till the end of the year. Now I'm not sure.

Leo: Yeah. So I was fooled, too, because the headline of the article I was reading said, "It's rolling out." And the last paragraph of the article is - so Google announced it at Google I/O in May. I mean, the problem is you can have AirTags following you around. Unless you've installed an app on your Android phone, it doesn't know about AirTags.

Steve: Right. Well, installed and it's running.

Leo: And it's running, and by the way doesn't work very well. So it was reasonable for Google and Apple to try to solve this problem by Google building it into Android and so forth. And so Google announced it in May that they were going to do this at

Google I/O. They had thought they were going to put it in Android. I think Android 14 is soon. But this article...

Steve: And I guess 6.0 is the kernel version.

Leo: Right. So this article, I'm reading along, and they're going to do it, they're going to do it. And then the last paragraph of the article is Google has announced it's putting this off until the end of the year because of the Apple Google Consortium. They want to work it out between the two of them. So I don't think it's in there now. I know I was very confused by this personally. So I don't - they've promised it. And we need it. But is it here? I don't know.

Steve: Yeah. And so what I picked up on said, you know, are beginning to roll out to Android 6.0+ users this month. So...

Leo: I saw that, too. And that was the same article that then said at the end of the article, well, except, no.

Steve: Except not.

Leo: I'm very confused by the whole thing. Let me see if I can find the article I read because I bet it was the same as the one you read. I bookmarked it in my thing here. And I think it was almost - it was as if they had written the whole article and then did a "never mind" on the whole thing. So, yeah, the headline of this article is, just as you said, "Android will now warn about unknown Bluetooth trackers like AirTag traveling with you." Sarah Perez writing for TechCrunch, July 27th. "Google today will begin to roll out a new safety feature, unknown traffic alerts." But then go down to the bottom, same article. "Today, however, Google says this update is on hold."

Steve: Wow.

Leo: Wait a minute. Also announced, Google said it would update its - okay. So I guess the alerts are there, but they are not updating the Find My network to work with third party. I guess that's just what you said.

Steve: Okay. Okay.

Leo: So if you read this carefully, which I didn't, apparently, the update that's on hold is updating Find My Device to work with third-party trackers. So they are going to - but then it says the decision was made to roll out these updates because Google is now working in partnership with Apple to finalize the joint unwanted tracker alert specification by year-end.

Steve: Wow. Really confusing.

Leo: I think we are getting the alerts. I think it was just a poorly written article, that we are getting alerts.

Steve: I think that is right. With Android 6.0 kernel, you'll begin to be told if something is traveling with you, the other stuff to come later.

Leo: Yeah.

Steve: And it does, it says, currently work with Apple AirTags. We will continue to work with tag manufacturers to expand this important protection. So I don't understand why it's not working with everybody because, you know [crosstalk] unified standard.

Leo: It all works the same; right? Yeah.

Steve: Yeah. But apparently not.

Leo: Also announced, Google said it would update its Find My Device network to help users locate other missing belongings, which can be located by third-party Bluetooth trackers. Now, Google doesn't sell a tracker. So anything Google works with is third party, including AirTags, Tile, Chipolo.

Steve: And I would say they don't "yet" sell one because, boy, I'm astonished by how popular Apple AirTags are.

Leo: Yeah. Oh, they're dominant.

Steve: They shared the numbers, it was millions of them were selling.

Leo: And I think what, really, there's lots of ways to track people.

Steve: And you know there's one fewer AirTag in use now, Leo, after you...

Leo: By the way, Burke gave me a much better hammer for next time I want to destroy something. I now have a mini sledgehammer. He didn't understand fully my plan.

Steve: Okay. So the National Defense Authorization Act, which successfully passed through the U.S. Senate last week, included a provision requiring the National Academy of Public Administration, whatever the hell that is, to conduct an assessment on the feasibility of establishing a new, formal, seventh branch of the U.S. military, which we've talked about several times, the U.S. Cyber Force. So this does appear to be happening.

Since many of our listeners have explained that wearing ridiculous camouflage clothing indoors is a bizarre requirement of the U.S. military - now, that's my word "bizarre," not

theirs - perhaps at least the Cyber Force's camo could have some cool cyber theme, like maybe like those green falling and fading symbols from The Matrix? Or maybe just do the whole thing as in ones and zeroes. That would be very cool; right? Like make camo out of ones and zeroes.

Leo: Well, remember it's supposed to be camo. I think why can't they just make something that makes you invisible? I mean, let's do it.

Steve: That'd be really good.

Leo: Yeah, yeah.

Steve: Make it stealth, a stealth camo. That'd be good.

Leo: Ooh, yeah.

Steve: Anyway, I do hope that someone gives this as much thought and serious consideration as is clearly needed because this U.S. Cyber Force, if they're going to have to wear some ridiculous outfit, let's make it techie and cool.

Leo: So a number of our listeners are saying, including in the Discord and somebody in the U.K. that they do have these alerts now on their Android phone. So it did roll out, yeah.

Steve: Oh, cool. Yay. Good. Thank you.

Leo: [Crosstalk] listeners.

Steve: Super use for feedback, yeah. Now, the other wrinkle is that both the Army and the Air Force, you know, obviously well-established branches of the military, have recently created their own new specialized cyber teams to support their traditional "kinetic teams," as we're calling them, you know, with cyber tasks related to intelligence gathering, electronic warfare, and sensors. And I think that makes sense, since those cyber teams which support the traditional kinetic forms of warfare are probably going to be highly targeted and specialized for their specific tasks; whereas the military's new seventh branch would be far more wide-ranging and not at all focused upon specific current Army and Air Force military operations.

So anyway, but through all this it is quite obvious that cyber - I know you love that term standing by itself, Leo - cyber has well and truly arrived, both on the front lines and soon in dimly lit dens filled with monitors and empty caffeinated beverage cans.

Leo: And pizza boxes.

Steve: So I want to know what they're going to be wearing. That's all I'm saying is, you know...

Leo: We should explain.

Steve: So for some reason this really matters to me.

Leo: We should explain that a couple of weeks ago Steve found a photo of the Cyber Defense Command, and they were all wearing BDUs, Battle Dress Uniforms, that were camouflaged. But obviously they're not in the jungle, so...

Steve: They're in a room. They're not even - you can't even observe them from satellite reconnaissance.

Leo: And we have to find some stealth uniforms for them.

Steve: Yeah.

Leo: Tempest uniforms.

Steve: Okay. So meanwhile, Russia continues to separate itself from the West. The Russian Parliament just passed three bills which, once signed into law by Putin, will ban Russian citizens from participating - I know.

Leo: This is crazy.

Steve: ...in the activities of foreign nonprofit organizations that have not specifically registered with the Russian government, and none have. Commentary about this over on opennet.ru notes that an unintended side effect will be that Russians using open source software would be prevented from contributing in any way to those projects, even from submitting bug reports. Now, as we know, today's open source software includes Linux, Firefox, most major database systems, and programming languages. Now, I read the entire piece after having Google translate it into English for me, and it only talked about the unintended consequences. I was unable to determine what the intended consequences of the three pending bills would be. Why would Russia think this was a good idea?

Leo: I know one of the reasons repressive regimes pass bills like this is for selective prosecution.

Steve: Okay.

Leo: So, you know, they need a way to get you.

Steve: If they want to stomp on somebody they have a law.

Leo: Oh, what's the company of Linux doing there? You're in trouble, big boy. That kind of thing. Yeah.

Steve: Yeah.

Leo: I mean, I can't imagine they want to stop all open source.

Steve: I mean, they're using it.

Leo: Yeah. The Russian official operating system is a Linux-based operating system.

Steve: Yes. And there are lots of really good Russian teams that are doing good work.

Leo: Well, maybe that's who they're stomping on.

Steve: You want to check the source code, but still.

Leo: Yeah. Maybe that's what it's all about. You know, as kind of retaliation for the sanctions or something.

Steve: We don't want you to have any of our stuff; right. So VirusTotal is out with their look at 2023 to date. It's always interesting since they've got a good snapshot, since everybody is submitting stuff to them, you know, whenever I, as I've mentioned before, when I download some old archive from some sketchy-looking site, I immediately hand it to VirusTotal to see what it thinks, just because, you know, it's better to be safe than sorry. So they get a really good snapshot of this. So they have some main takeaways from their most recent update.

First of all, email attachments, to no one's surprise, continue to be the most popular way to spread malware. However, traditional file types - Excel, RTF (Rich Text Format files), CAB, and compressed formats - are becoming less popular. Although the use of PDFs slowly decreased for the last few months, starting in June of 2023, the biggest peak in PDF usage was observed during 2023 compared to the previous two years. So PDFs are still a big deal, maybe a little summer slump for some reason. However, the big changes are in OneNote. OneNote and JavaScript, both distributed through HTML, are the most rapidly growing formats for malicious attachments in 2023, with OneNote emerging this year 2023 as a reliable alternative for attackers to the traditional use of macros in other Office products.

Malicious OneNote files usually embed an additional malicious file. So OneNote is just sort of serving as a recognizable container that seems benign. And I guess leave it to Microsoft, their various security permissions allow OneNote to be opened when you click on something on a web page, so yeah, let's have OneNote bring it in. So OneNote files usually embed an additional malicious file - a VBA, HTML and JavaScript, PowerShell, or

some combination of those. And as happens with malicious Office attachments, the attempt is then made to convince the user to allow its execution.

Payloads vary from one malware family to another, but many of them access external URLs to then download a DLL file which is camouflaged as a .PNG, you know, which is an old trick used to bypass simple firewall rules or just to appear less suspicious to anybody who knows to look.

The most usual kill chain, as it was noted and stated, where OneNote format is involved is three steps. The victim receives an email with a OneNote attachment. The mail body encourages the victim to click on a button to see a hidden or distorted image or document. Second, this button executes a script - VB script, a powershell, or whatever. And that will launch a payload, either embedded into the same script or downloaded from an external resource. And then finally the external payload might be yet another OneNote file, an image file renamed as a ".bat" file, a DLL file that's loaded into memory or even a Windows executable.

So we have inherently dangerous capabilities mixed with social engineering attacks. And only one mistake made by one curious or inattentive employee within a major organization is all that's required to invite the malware in to set up shop and, who knows, contact a satellite Internet provider in order to say, hey, I made it in, what do you want me to do?

Following behind OneNote, ISO image files for malware are now a flexible alternative for both widespread and targeted attacks. And their distribution as heavily compressed attachments makes them difficult to scan by some security solutions. So says VirusTotal. ISO files are being disguised as legitimate installation packages for a variety of software, including Windows, Telegram, AnyDesk, and Crypto Notepad, among others.

Virus Total said that - they said: "Our data shows that there was an increase in the number of malicious files attached to emails between March and April of 2023. In terms of suspicious attachments, for the past two years we have observed spikes in the number of suspicious PDF files linked to malicious campaigns. These files can be used for a variety of purposes, such as exploiting vulnerabilities or phishing, which is what happens most of the time."

And they said during 2023 so far they saw a significant increase in the use of JavaScript distributed alongside HTML, used in sophisticated phishing attacks which were designed to steal victims' credentials. Excel, RTF, CAB, and compressed formats, as I mentioned before, and Word, interestingly, seem to be declining in popularity along with the others as malicious attachments compared to OneNote and JavaScript. So that's the wrap-up on what's been happening so far in 2023. And we should have already taken a break, Leo, but let's do it now.

Leo: I'm ready.

Steve: And I'm going to share some amazing feedback from our listeners.

Leo: Okay, Steve. On with the show.

Steve: It's funny, too, when you're talking about password managers, I just - I can't imagine life without one.

Leo: It's so much easier once you're used to it; right?

Steve: Well, and I think that - yes, that. And maybe 10 years ago, 20 years ago, well, I mean, people had four or five online accounts, you know what I mean, because there wasn't that much to do. There wasn't that much going on online. Now our lives are online. And, you know, I mean, all of our utilities we have accounts for, and all of our various services we have accounts, and if you want to grab a car and drive somewhere, I mean, I just, you know, all of the airlines you have accounts, I mean, everything. And so if they're going to all have their own password, you just have no choice.

Leo: You have to. Even if it were memorable you'd have to. But I can't tell you how many people I know in my own personal family even, who know better, but you know, it's just, well, [mumbling]. Patrick Delahanty, our engineer, says his dad, who is a U.S. attorney, has a little black book of passwords. The problem is, yeah, you can do that. But then you have to generate unique passwords each time. It's just easier to use a password manager and let it do the heavy lifting.

Steve: Yeah.

Leo: I think easier than putting it in a notebook. You don't have to write it down. You don't have to write it down. You don't have to remember. You don't have to look it up. It just does it. Anyway.

Steve: Okay. So some feedback. Jeff Parrish, he said: "Thank you for another great episode. I am IT for a healthcare facility, and this episode" - referring to last week - "made me review the HTML of our EHR provider. I have now contacted them about the Google Analytics tracking they have on their site after we are logged in." So that was cool and useful to at least one of our listeners.

Actually, another one, Robert C. Covington, he said: "Longtime listener. I oversee cybersecurity for a large children's hospital system. Your podcast transcripts are frequently on my screen during team meetings."

Leo: Yay. Wow, that's awesome.

Steve: Yeah. He said: "Regarding website tracking and the recent OCR notice referenced in Episode 932" - last week - "there is a side consequence I've not heard mentioned. Cyber insurance companies are now declining to cover any legal actions arising out of website tracking and collection of personal health information. This is sending many healthcare orgs scrambling to get tracking tools off their websites. Keep up the excellent work. Robert Covington." He said, oh, he said: "P.S.: You fell into the classic trap on 932. It is HIPAA, not HIPPA." So thank you for the correction, Robert. And very interesting. That will certainly remove tracking from healthcare if they know that they're not going to get any insurance coverage from their providers if they do that, and anyone gets called out for having personal health information disclosed. If there's trackers on there, sorry, your insurance won't cover that. Wow.

Jon Dagle said: "Hi, Steve. Thanks for the shout-out on the 25th July episode." That was last week. He said: "I am the 'neat guy' who you saw on TWiS (This Week in Space) talking about orbital debris." And he said, he was joking, he said: "I'm fairly sure you

weren't referring to Geof, and I'm really sure you weren't referring to Rod, ha ha." So he said: "Thank you for your kind mention." He said: "I've been a Security Now! listener since Episode 1, proud SpinRite owner," he says, "and somewhere I have a certificate for a TWIT Brick." He said: "Pretty sure I've not missed a single episode, at least not a whole one. At the beginning I was in the U.S. Air Force. I stuck around for hobbyist purposes and with a plan to go into cybersecurity, but I made a detour into space policy.

"Orbital debris is a clear and present concern, if not actual danger." He said: "The space advocacy organization where I work considers this one of a handful of high priorities. While there are a number of public sources for tracking objects in orbit, they don't all agree. According to orbit.ing-now.com, a relatively approachable source," using that as a reference. And he said: "A high-level summary is available here," and I have a link to it, it's a long URL, nanoavionics.com/blog/how-many-satellites-are-in-space, and that's all hyphenated. And I think I clicked that, and it showed a picture of the Earth. And if that's an accurate depiction, it is a little sobering. And I think it may be accurate because it actually shows some of Elon's satellite trains.

Leo: 4,500 SpaceX satellites, which is half of all satellites are SpaceX.

Steve: Yes, yes.

Leo: This is not - not SpaceX, Starlink, which is from SpaceX.

Steve: Starlink from SpaceX, yes.

Leo: I thought Starlink, oh, I was so happy when I, you know, oh, we're going to low-cost Internet coverage to every corner of the globe. First of all, it's not low cost. It's very expensive. And second, he's going to put 42,000 satellites up.

Steve: Yes.

Leo: This is only one tenth.

Steve: It's going to be Starjunk instead of Starlink.

Leo: I mean, this is horrific.

Steve: So Jon said: "There are about 7,700-8,400 active human-made satellites in orbit around our planet. The vast majority, 90%, are in low-Earth orbit." And so that's less than 1,000 kilometers up. "About one third of these have been added in the past few years, mainly by SpaceX Starlink. About 7% of the total are in geostationary orbit, where these are LEO satellites we were just talking about, the GEO are geostationary," he said, "with the remainder in medium-Earth orbit, very few of those." He said: "Almost 2,300 'inactive' satellites," meaning they're up there but they died or they're dead or their battery ran down or something. And he said: "Thanks for the shout-out and," he said, "the 'brush with greatness.'" Jonathan, Washington, DC. He's the Policy Chair for the

National Space Society. And, yeah, you had that on the picture, that beautiful picture of the Earth. And I think, I can't see it there, but it showed like Elon's Starlink trains.

Leo: Yeah. There's lines, yeah.

Steve: Yes.

Leo: The New York Times this Sunday had an article about the concern, the geopolitical concern that Elon, who is, let us say, seemingly slightly erratic, controls this Starlink system, and the Ukraine military relies on it for military communications. And they're concerned. They asked in May, the Times reported, they asked the federal government what's the deal with this Elon, and the government basically went, uh, we don't know.

Steve: Well, and we're also in bed with him, right, because now we're contracting with him to launch our major space payloads.

Leo: Yeah. I bet they're regretting that a little bit right now. I mean, he just seems quite erratic. I'll find this New York Times picture because it's actually animated, and it's quite good. It's really, this looks like similar data because, yeah, you can see these Starlink trains in it.

Steve: Yeah. So we have another listener, John Sutherland, whose Twitter handle is @JohnOrion, which I got a kick out of. He said: "I wanted to offer a bit of knowledge I had about U.S. military satellites. I was active duty and what is now Space Force for 11 years, and I'm currently a contractor still supporting space. I 'flew' SAT-COM for four years, then taught for seven years."

Leo: Wow, that's cool. What an audience. We have amazing people in the audience.

Steve: We have amazing listeners, Leo.

Leo: Blows me away.

Steve: It is great. He said: "I taught both classified and unclassified classes, so I am very familiar with where the line is for what's classified. I can go right up to that line. Having just finished the second part of Satellite Insecurity" -meaning last week's podcast - "I can share that, luckily, most of the problems you talked about are not as true for U.S. DoD satellites."

Leo: Ah, I bet not, yeah.

Steve: Whew. Thank goodness.

Leo: Yeah, we're well-protected, I bet, yeah.

Steve: Yeah. He said: "The preconceptions that attackers would not have the equipment was never the case. China and Russia have always had similar ground station capabilities as we have. The oldest satellites I've worked with were developed in the late '80s, and they were highly encrypted and rolled keys constantly. For communication satellites the data is just routed so encryption is as good as it could be on Earth and not subject to the satellites' age. Controlling the satellites, i.e., moving them, changing configuration, is done with separate antennas that are monitored, and any communication with them is watched in real-time. If someone did break this encryption, it would quickly be learned.

"As for physical attacks" - and this gets a little interesting with his choice of words. "As for physical attacks, the arms of attacking satellites is only a start. When we table-topped attacks and planned responses, TTP (tactics, techniques, and procedures), we looked at jamming, ASATs" - which he'll explain in a second - "mechanical arms, and lasers, jamming being the most common and ones we have actually seen happen. Most jammers are big ground-based semi-trucks or ships that just try to overpower the uplink." So they're just blasting the same satellite target, hoping that it won't be able to receive the actual signal.

He said: "We have many mitigations to this, and I taught a class on RF Attack and Defense as part of operators' advanced training. ASATs," as he uses the term, "as you talked about with blowing up satellites from the ground are extremely unlikely at this point. We are much more concerned with small satellites with explosives, the idea being that an adversary would place and leave something small on a foreign satellite that could be triggered on demand at any time in the future." Whoa. So they're mining satellites without the satellites' knowledge. They creep up, stick something sticky on the side that's a bomb with a radio, and then leave. And that can then be detonated in the future. So, I mean, what a mess, Leo. Can you imagine, like everybody's satellites have all these bombs stuck to them from other hostile nations?

Leo: Good lord.

Steve: It just...

Leo: I really want to ask these - by the way, here's The New York Times animation. This is 10 minutes of Starlink satellites.

Steve: Wait, in the future.

Leo: No, no. This is as of July, launched as of July 10th. This is current.

Steve: Now. Look at all of them.

Leo: Yes. Yes.

Steve: Oh, my lord.

Leo: Yes. I'm wondering if we're having second thoughts about letting Elon launch all of these. This is crazy.

Steve: Wow.

Leo: This is half of the entire satellite load.

Steve: And they're only in a train before they've distributed themselves?

Leo: Yeah, they deploy from the train. But if you look, it looks like there's groups of two and three in some places. It's a really interesting - there's definitely method to the madness. Yeah, those trains you see are not yet deployed. They launch that way, and then they slowly deploy. Isn't that wild?

Steve: Wow.

Leo: I want our satellite experts, though, to tell me if I should worry about the Kessler Syndrome, Kessler Effect or not.

Steve: Right.

Leo: So if you blow up a satellite, and then debris from the satellite then blows up five more satellites, and debris from those satellites blows up 25 more satellites, and on and on and on, could you occlude the night sky? Or worse?

Steve: This has been demonstrated with dominos.

Leo: Yeah. I'm starting to worry.

Steve: And it's not good.

Leo: I know there are missions, that we're running missions and I think China is running missions to snarf up satellites like the Moonraker thing we were talking about. But, boy, I just, I mean, what happens? I mean, I guess when they've reached the end of life they just go through the atmosphere and burn up.

Steve: All I can say is we should hold onto our DVD collections because we do not want to become too dependent on space.

Leo: On the Internet.

Steve: Yeah, on space-based Internet.

Leo: Yeah, crazy.

Steve: So he finished, saying: "I cannot talk to the mechanical arms as the line beyond which I cannot talk is around this. But it's safe to say this has been looked at and is in some level of development by both sides." He said: "Lasers are not a threat to all types of satellites, but China and Russia have used lasers to blind sensors of low-flying 'spy' satellites. This is hard to guard against, but we do equip satellites with shutters now. And for satellites lacking shutters, we only need to spin them around."

And he finished: "There's more that cannot be talked about, but with your level of technical knowledge and a little imagination you could get close to guessing what's going on. I can tell you I've never been surprised when I got a security briefing." So, very cool.

Leo: Jiminy Christmas. Thank you, We have wonderful listeners. We thank you all. Just really fascinating.

Steve: And another one, Mikael Falkvidd. Mikael is on the board of OWASP in Gothenburg, Sweden, and he's the guy who invited me to present SQRL to their group. Well, it turns out that Mikael knows more than a little about satellite software. He said: "Regarding authenticated TeleCommands to satellites." Now, we talked about this last week; right? The idea being that TeleCommands are ways you tell satellites to do things. But what the guys who reversed the firmware found was there was a surprising lack of authentication.

So Michael said: "What satellite programmers are most afraid of is bit flips caused by single-event upset" - what is termed an SEU, a single-event upset.

Leo: You mean cosmic rays?

Steve: Yes. He says: "...which happen due to radiation in space." He said: "Imagine that an SEU flips a bit in the key used to authenticate the TeleCommand."

Leo: Right.

Steve: "Authentication would fail. And guessing which bit or bits flipped could take some time."

Leo: That's why you have ECC. I mean, we have ways...

Steve: He says: "There are of course mitigations, for example, using error correction codes or storing the key in multiple places. But complexity is the enemy of reliability, and resources (compute, flash, ram) onboard satellites have been very scarce historically. And people want reliable satellites, so they are hesitant to introduce new features. 'Flight-proven' [he has in quotes] is the mantra, so the old ways live on. The risk of losing the satellite because of an SEU (single-event upset) has been deemed higher than the risk that the satellite is hacked. Not an excuse today, but that's how the industry is."

And then he finished, saying: "I have written software for two satellites." And he said - yeah, like you said, Leo, our listeners are amazing.

Leo: Wow.

Steve: "SEUs are also one of the reasons TeleCommands exist to write to any memory location. NASA used this feature to restore a bit flip on Voyager 2 in 2010."

Leo: Vger.

Steve: Vger, 33 years after its launch.

Leo: Wow.

Steve: So Mikael also provided a link to a summary from JPL, the Jet Propulsion Laboratory in Pasadena, which documented events surrounding exactly this happening back in May of 2010. Somewhat astonishingly, Voyager 2 remains alive and functioning to this day, though something happened with it just last week, which I'll get to in a second. We last checked in on Voyager 2 nearly five years ago when, on November 5th of 2018, it became only the second spacecraft to ever exit our solar system's heliosphere.

And remember, Leo, we considered whether this event might break the simulation that Elon, among others, appear to be convinced we're all living within. But so far the simulation appears to be holding. We were wondering if Voyager 2 exited the heliosphere, was there a maximum radius at which the simulation would still be functioning and whether Voyager 2 might just spontaneously disappear because it got too far away.

Anyway, let's turn the calendar back 13 years to May 6th of 2010, when JPL wrote, they said: "Engineers have shifted NASA's Voyager 2 spacecraft into a mode that transmits only spacecraft health and status data while they diagnose an unexpected change in the pattern of returning data. Preliminary engineering data received on May 1st" - this would be May 1st of 2010 - "show the spacecraft is basically healthy, and that the source of the issue is the flight data system, which is responsible for formatting the data to send back to Earth. The change in the data return pattern has prevented mission managers from decoding science data.

"The first changes in the return of data packets from Voyager 2, which is near the edge of our solar system, appeared on April 22nd. Mission team members have been working to troubleshoot and resume the regular flow of science data. Because of a planned roll maneuver and moratorium on sending commands, engineers got their first chance to send commands to the spacecraft on April 30th. It takes nearly 13 hours for signals to reach the spacecraft, and nearly 13 hours for signals to come down to NASA's Deep Space Network on Earth.

"Voyager 2 launched on August 20th, 1977 [so, wow] about two weeks before its twin spacecraft, Voyager 1. The two spacecraft are the most distant human-made objects, out at the edge of the heliosphere, the bubble the sun creates around the solar system. Mission managers expect Voyager 1 to leave our solar system and enter interstellar

space in the next five years or so, with Voyager 2 on track to enter interstellar space shortly afterward. Voyager 1 is in good health and performing normally.

"Ed Stone, Voyager project scientist at the California Institute of Technology in Pasadena, said: 'Voyager 2's initial mission was a four-year journey to Saturn, but it is still returning data 33 years later. It has already given us remarkable views of Uranus and Neptune, planets we had never seen up close before. We will know soon what it will take for it to continue its epic journey of discovery.'" Meaning at the point where he's talking about this, something broke, and Voyager 2 is no longer sending data back, the science data that they wanted.

And he said: "The original goals of the two Voyager spacecraft were to explore Jupiter and Saturn. As part of a mission extension, Voyager 2 also flew to Uranus in 1986 and Neptune in 1989, taking advantage of a once-in-a-176-year alignment to take a grand tour of the outer planets." I just love this. It is just so cool, you know, real science.

"Among its many findings, Voyager 2 discovered Neptune's Great Dark Spot and 450-meter-per-second (1,000 mph) winds. It also detected geysers erupting from the pinkish-hued nitrogen ice that forms the polar cap of Neptune's moon Triton. Working in concert with Voyager 1, it also helped discover actively erupting volcanoes on Jupiter's moon Io, and waves and kinks in Saturn's icy rings created by tugs of nearby moons. Voyager 2 is about 13.8 billion kilometers, 8.6 billion miles, from Earth. Voyager 1 is about 16.9 billion kilometers, 10.5 billion miles from Earth. The Voyagers were built by JPL, which continues to operate both spacecraft. Caltech manages JPL for NASA."

Okay. So May 6th, 2010, and something is broken and has gone wrong with Voyager 2 such that the spacecraft's science data is no longer being properly formatted. Eleven days later on May 17th, 2010, we learn what went wrong: "Engineers at NASA's JPL said Monday, May 17th, that one flip of a bit in the memory of an onboard computer appears to have caused the change in the science data pattern returning from Voyager 2. A value in a single memory location was changed from a zero to a one. On May 12th" - so that was - yeah. So, "On May 12th, engineers received a full memory readout from the flight data system computer, which formats the data to send back to Earth. They isolated the one bit in the memory that had changed, and they recreated the effect on a clone computer at JPL. They found the effect agrees with the data coming down from the spacecraft. They are planning to reset the bit to its normal state on Wednesday, May 19th."

And then three days later on May 20th we have the report of the conclusion of this high-stakes drama: "Engineers have successfully corrected the memory on NASA's Voyager 2 spacecraft by resetting a computer bit that had flipped. Reset commands were beamed up to the spacecraft yesterday, Wednesday, May 19th; and engineering data received today confirm that the reset was successful. The Voyager team will continue monitoring the engineering data, and if the bit remains properly reset, commands to switch to the science data mode will be beamed up to Voyager 2 on Saturday, May 22nd. Receipt of science data would then resume on Sunday, May 23rd."

And all of that did happen on schedule. But I also noted that something else happened just last week. NASA's blog posting Friday, July 28th of this year, read: "A series of planned commands sent to NASA's Voyager 2 spacecraft" - right, still going strong - "on July 21st" - so toward the end of, like, just a couple weeks ago, toward the end of last month - "inadvertently caused the antenna to point 2 degrees away from Earth." Now, when you're billions of miles away, 2 degrees, baby, I mean, you might as well be looking in the other direction. So, "As a result, Voyager 2 is currently unable to receive commands [whoops] or transmit data [whoops] back to Earth.

"Voyager 2 is currently located almost 12.4 billion miles from Earth, and this change has interrupted communications [no kidding] between Voyager 2 and the ground antennas of the Deep Space Network. Data being sent by the spacecraft is no longer reaching the Deep Space Network, and the spacecraft is not receiving commands from ground controllers." Right. "Voyager 2, however, is programmed to reset its orientation multiple times each year to keep its antenna pointed at Earth. The next reset will occur on October 15th, which should enable communication to resume. The mission team expects Voyager 2 to remain on its planned trajectory during the quiet period. Voyager 1, which is almost 15 billion miles from Earth, continues to operate normally."

And, finally, a couple of interesting tidbits about the Voyager probes: Uplink communications to the Voyagers is via S-band at 16 bits/sec, while an X-band transmitter provides downlink telemetry at 160 bits/sec normally, and 1.4 kbps for playback of high-rate plasma wave data. Although I think that I saw that the plasma wave science equipment has been turned off due to power consumption. All data are transmitted from and received at the spacecraft via the 3.7 meter high-gain antenna. So that's the big high-gain dish, and obviously being a dish it's pointy, so you've got to point it in the right direction.

Electrical power is supplied by three Radioisotope Thermoelectric Generators (RTGs). The current power levels are about 249 watts for each spacecraft. As the electrical power decreases, power loads on the spacecraft must be turned off in order to avoid having demand exceed supply. Otherwise the voltage would drop. As loads are turned off, some spacecraft capabilities are eliminated.

So NASA maintains an extremely cool real-time Voyager status page which continuously shows the location of both spacecraft and other interesting tidbits such as which science modules are currently turned on and off, given the amount of available power. So I created a shortcut, grc.sc/voyager, because the page is so cool. We haven't looked at it since we last talked about the Voyager probes: grc.sc/voyager. Or you can just Google "Voyager Mission Status," and that will bring up as the first link that page. I mean, and it's updating as you watch it on the fly, how far both of these probes are, and also which science modules are turned on and off. So anyway, big thanks to our satellite-informed listeners for their information and feedback.

Leo: And we won't lose Voyager because it's going to reorient, so that's good news, yeah.

Steve: Right, right.

Leo: I do want to correct myself. It's not Vger. I was looking it up. I thought, well, which one was Vger, Voyager 1 or Voyager 2? Neither.

Steve: Oh.

Leo: Vger - is this a spoiler now? No, I won't tell you what I'm talking about. If you know, then you know. Vger was Voyager 6, which was, remember, this is a movie that came out in 1979, which was to be launched in 1999.

Steve: Ah.

Leo: And of course there is no Voyager 6.

Steve: It's a future Voyager that we haven't launched yet.

Leo: It's a future Voyager, yeah.

Steve: Of course.

Leo: Which explains how it got so smart. Because by 1999 AI was happening. It's funny how we thought all this stuff would be happening by now. Anyway, great story.

Steve: Oh, Leo, everyone wants to know where their flying cars are.

Leo: Yeah, yeah.

Steve: You know?

Leo: Yup.

Steve: Yeah. And now I know that would be a very bad idea, so, yeah.

Leo: Voyager 2's been out there for 45 years.

Steve: Unbelievable.

Leo: Isn't that amazing?

Steve: That is really - that is.

Leo: Yeah.

Steve: Wow. So Jon David Schober, he said: "Hey, Steve. On SN-932 I heard you talking about how you're keeping the rack of servers at Level 3, and not moving to the cloud. In case you wanted some interesting reading, here is a blog post from David Hansson, founder of 37signals and Basecamp and creator of Ruby on Rails."

Leo: DHH, David Heinemeier Hansson, yes.

Steve: Yes, David, DHH, yes. "He discusses how they regret moving their business to AWS, and how expensive everything was, and how much better life is being back on their own hardware."

So first of all, Jon, thanks very much for the pointer. Since this topic is quite near and dear to my heart, and since I think it might also be extremely interesting to a large number of our listeners, I want to share the blog post that Jon pointed to. As Jon said, this was written by David H. Hansson, and it was posted just last October 19th, 2022, titled "Why we're leaving the cloud."

David wrote: "Basecamp has had one foot in the cloud for well over a decade, and HEY (H-E-Y) has been running there exclusively since it was launched two years ago. We've run extensively in both Amazon's cloud and Google's cloud. We've run on bare metal virtual machines. We've run on Kubernetes. We've seen all the cloud has to offer, and tried most of it. It's finally time to conclude: Renting computers is (mostly)" - he has in parens - "a bad deal for medium-sized companies like ours with stable growth. The savings promised in reduced complexity never materialized. So we're making our plans to leave."

He continues: "The cloud excels at two ends of the spectrum, where only one end was ever relevant to us. The first end is when your application is so small and low traffic that you really do save on complexity by starting with fully managed services. This is the shining path that Heroku forged, and one that has since been paved by Render and others. It remains a fabulous way to get started when you have no customers, and it'll carry you quite far even once you start having some." He says, parens: "(Then you'll later be faced with a Good Problem once the bills grow into the stratosphere as usage picks up, but that's a reasonable trade-off.)"

He says: "The second" - meaning the second use case - "is when your load is highly irregular. When you have wild swings or towering peaks in usage. When the baseline is a sliver of your largest needs. Or when you have no idea whether you need ten servers or a hundred. There's nothing like the cloud when that happens, like we learned when we launched HEY, and suddenly 300,000 users signed up to try our service in three weeks instead of our forecast of 30,000 in six months.

"But neither of those conditions apply to us today." And I would say neither of them apply to me, GRC, and actually probably to TWiT. He says: "They never did for Basecamp. Yet by continuing to operate in the cloud, we're paying an at times almost absurd premium for the possibility that it could. It's like paying a quarter of your house's value for earthquake insurance when you don't live anywhere near a fault line. Yeah, sure, if somehow a quake two states over opens the earth so wide it cracks your foundation, you might be happy to have it, but it doesn't feel proportional; does it?

"Let's take HEY as an example. We're paying over half a million dollars per year for database (RDS, relational database) and search (Elastic Search) services from Amazon. Yes, when you're processing email for many tens of thousands of customers, there's a lot of data to analyze and store, but this still strikes me as rather absurd. Do you know how many insanely beefy servers you could purchase on a budget of half a million dollars per year?

"Now the argument always goes: Sure, but you have to manage these machines. The cloud is so much simpler. The savings will all be there in labor costs. Except, no," he says. "Anyone who thinks running a major service like HEY or Basecamp in the cloud is simple has clearly never tried. Some things are simple; others are more complex. But on the whole, I've yet to hear of organizations at our scale being able to materially shrink their operations team just because they moved to the cloud.

"It was a wonderful marketing coup, though. Sold with analogies like, 'Well, you don't run your own power plant either; do you?' Or 'Are new infrastructure services really your core competency?' Then lathered up with a thick coat of new-new-new paint, and The Cloud [he has in caps] has beamed so brightly only the luddites would consider running their own servers in its shadow.

"Meanwhile, Amazon in particular is printing profits, renting out servers at obscene margins. AWS profit margin is almost 30%," and he says, "$18.5 billion in profits on $62.2 billion in revenue, despite huge investments in future capacity and new services. This margin is bound to soar now that 'the firm said it plans to extend the useful life of its servers from four years to five, and its networking equipment from five years to six in the future.' Which is fine. Of course it's expensive to rent your computers from someone else. But it's never presented in those terms. The cloud is sold as computing on demand, which sounds futuristic and cool, and very much not like something as mundane as 'renting computers,' even though that's mostly what it is.

"But this isn't just about cost. It's also about what kind of Internet we want to operate in the future. It strikes me as downright tragic that this decentralized wonder of the world is now largely operating on computers owned only by a handful of mega corporations. If one of the primary AWS regions goes down, seemingly half the Internet is offline along with it. This is not what DARPA designed.

"Thus I consider it a duty that we at 37signals do our part to swim against the stream. We have a business model that's incredibly compatible with owning hardware and writing it off over many years. Growth trajectories are mostly predictable. Expert staff who might as well employ their talents operating our own machines as those belonging to Amazon or Google. And I think there are plenty of other companies in similar boats.

"But before we can more broadly set sail back toward lower-cost and decentralized shores, we need to turn the rudder of our collective conversation away from the cloud-serving marketing nonsense about running your own power plant. Up until very recently, everyone ran their own servers, and much of the progress in tooling that enabled the cloud is available for your own machines, as well. Don't let the entrenched cloud interests dazzle you into believing that running your own setup is too complicated. Everyone and their dog did it to get the Internet off the ground in the first place, and it's only gotten easier since. It's time to part the clouds and let the Internet sunshine through." So anyway.

Leo: He's kind of a - he's a crackpot, but okay.

Steve: Yeah.

Leo: There's a lot of reasons you'd want a cloud. For AI training, for instance, you're not going to go out and buy a thousand cards from NVIDIA and a bunch of servers and stuff just for the training, and then what? And then just let them sit in the basement?

Steve: Well, you just gave a perfect use case for the cloud, and I've heard that suggested. You would use the cloud to train the model [crosstalk] run the model.

Leo: And then run it locally. Yeah, lots of people do that.

Steve: Yes.

Leo: I mean, I think [crosstalk] is very common.

Steve: But Leo, you...

Leo: You say you're not in the cloud, but Level 3 isn't on-prem. Aren't you in the cloud?

Steve: Well, everybody has some Tier 1 service provider. I mean, so you have an IP.

Leo: But your servers are in your house.

Steve: No, no, my servers are a short drive away in a datacenter.

Leo: But that's not the cloud because you own the hardware?

Steve: Correct.

Leo: Okay. All right. I mean, my website is right down the hall. It literally is on-prem here.

Steve: Except you've talked about how expensive Mastodon is. Mastodon for me would be free.

Leo: Mastodon's running in the cloud, yeah.

Steve: No matter how big it gets. So that's a...

Leo: Well, there's an example. I wanted to run Mastodon in the cloud because I didn't want to maintain it and run it off the servers here, and because we don't have enough local bandwidth to run it. I mean, obviously 37signals can afford to buy many, many gigabits of bandwidth; right? I mean, come on. He's kind of a crackpot. He's a well-known crackpot. But you know, since he put that out, he has succeeded. They are all off the cloud now, yeah.

Steve: Cool.

Leo: We'll see. I'd like to see what his bills are for running it locally. The problem really is that he doesn't see those bills because it comes in the form of rent and electricity and air conditioning and things that he doesn't consider cloud costs.

Steve: Well, I pay about a grand a month for all of GRC and all of my servers and all of my bandwidth.

Leo: But you're kind of a cloud because you're running in a network operations center. You're not running on-prem.

Steve: Well, no, he's talking about renting machines.

Leo: Oh, he's talking about the same thing, a colo? Yeah. Yeah. A colo is going backwards a little bit, I think. But okay, fine. It's whatever. There's a lot, I mean, there are a lot of businesses who will dispute that.

Steve: And that of course was the whole point of his blog post was that is sane to go backwards.

Leo: Okay.

Steve: That, you know, the promise of the cloud did not materialize. Anyway, I wanted to share it with our listeners.

Leo: No, it's good, yeah.

Steve: It is my position. It is what I'm doing. And I have fixed costs. I could run Mastodon servers till the cows came home, and it wouldn't cost me anything more.

Leo: Right. Except your time.

Steve: No. No. I mean, I maybe visit Level 3 annually. My servers are typically up for three or four years at a time.

Leo: Right, right.

Steve: So, yeah, I mean, it's just not - it's not a problem for me. But, you know, I built the stuff right once, so I don't have to be continually nursing them. And Leo, we're an hour and a half in, and we haven't even gotten to our main topic. Let's take our third break and talk about TETRA:BURST.

Leo: Okay. You didn't want to play...

ROBOT: Danger, danger.

Leo: Okay, that's fine. That's fine. And now, back to Security Now!.

Steve: Leo, we do need to explain the "Danger Will Robinson" sound effect.

Leo: Okay. Go ahead.

Steve: Steven Perry, he sent a note. He said: "Hi, Steve."

Leo: He's obviously a regular, by the way, in our Discord. We love Steven.

Steve: Ah. He said: "I was listening to yesterday's Security Now! episode and wondered if anyone had ever shared with you and Leo a little bit of trivia about the show 'Lost in Space.'" Which of course we both cut our teeth on.

Leo: Mm-hmm.

Steve: You know, as kids. "Everyone knows and uses the catchphrase 'Danger, Will Robinson.'" Of course one of our faves. He says: "But did you know that it was only ever said once in the entire run of the show?"

Leo: Wow.

Steve: "It was Season 3, Episode 11 when it happened."

Leo: But who's counting.

Steve: "It was never said again." Yeah. "But that is the phrase we all know and love about the show. Thought I'd pass it along. Have a good day." Well, I was astonished by that. I did a little bit of looking around. The Internet agrees with Steven. And apparently one of the reasons is that the robot was always waving its arms around, saying "Danger! Danger!"

Leo: That's what he said was "Danger! Danger!" We add the Will Robinson so you know what it means. If I just said "Danger! Danger!" you wouldn't know. But anyway.

Steve: Yeah, you'd think, what?

ROBOT: I am sorry, Will Robinson. I am afraid I goofed.

Leo: I have many, by the way, many robotic quotes.

ROBOT: A robot does not live by programming alone. Some culture is require to keep my tapes in balance.

Leo: Little did we know, in the future they're going to still use tapes in the robots.

Steve: Yeah, yeah. Actually, it's funny how the use of that term has hung on. I mean, people are still saying "Did you tape..."

ROBOT: Danger, danger. Danger, danger.

Steve: Okay. So by far the news that was most forwarded to me this past week was that the encrypted security of a globally used, "secure" in air quotes, radio communication system whose security has been trusted and relied upon worldwide, turns out not to be as secure as everyone hoped and was led to believe. And moreover, the system's insecurity was well known and kept secret by those whose commercial interests depended upon the system being trusted, when it was not trustworthy.

Wired did a beautiful job of describing the situation in their story last week titled "Code Kept Secret for Years Reveals Its Flaw - a Backdoor." And they followed that with "A secret encryption cipher baked into radio systems used by critical infrastructure workers, police, and others" - meaning lots of military - "around the world is finally seeing sunlight. Researchers say it isn't pretty."

Now, I'm going to share Wired's coverage of this while liberally interjecting my own commentary. So here's what Wired described. They said: "For more than 25 years, a technology used for critical data and voice radio communications around the world has been shrouded in secrecy to prevent anyone from closely scrutinizing its security properties for vulnerabilities." Now, okay. Anybody, if you've listened to this podcast for only one of our almost 18 years, you know that anytime you hear the technology was kept private to prevent anyone from scrutinizing its security properties for vulnerabilities is not good news.

Anyway, but Wired said: "Now it's finally getting a public airing, thanks to a small group of researchers in the Netherlands who got their hands on it and found serious flaws, including a deliberate backdoor. The backdoor, known for years by vendors that sold the technology, but not necessarily by customers, exists in an encryption algorithm baked into radios sold for commercial use in critical infrastructure. It's used to transmit encrypted data and commands in pipelines, railways, the electric grid, mass transit, and freight trains. It would allow someone to snoop on communications to learn how a system works, then potentially send commands to the radios that could trigger blackouts, halt gas pipeline flows, or reroute trains.

"Researchers found a second vulnerability in a different part of the same radio technology that is used in more specialized systems sold exclusively to police forces, prison personnel, military, intelligence agencies, and emergency services, such as the C2000 communication system used by Dutch police, fire brigades, ambulance services, and Ministry of Defense for mission-critical voice and data communications. The flaw would let someone decrypt encrypted voice and data communications and send fraudulent messages to spread misinformation or redirect personnel and forces during critical times.

"Three Dutch security analysts discovered the vulnerabilities, five in total, in a European radio standard called TETRA" - which stands for Terrestrial Trunked Radio - "which is used in radios made by Motorola, DAMM, Hytera, and others. The standard has been used in radios since the '90s, but the flaws remained unknown because encryption algorithms used in TETRA were kept secret until now.

"The technology is not widely used in the U.S." - well, not widely, but it is here - "where other radio standards are more commonly deployed. But Caleb Mathis, a consultant with Ampere Industrial Security, conducted open source research for Wired and uncovered contracts, press releases, and other documentation showing TETRA-based radios are used in at least two dozen critical infrastructures in the U.S. Because TETRA is embedded in radios supplied through resellers and system integrators like PowerTrunk, it's difficult to identify who might be using them and for what. But Mathis helped Wired identify several electric utilities, a state border control agency, an oil refinery, chemical plants, a major mass transit system on the East Coast, three international airports that use them for communications among security and ground crew personnel, and a U.S. Army training base.

"The researchers with Midnight Blue in the Netherlands discovered the TETRA vulnerabilities - which they're calling TETRA:BURST - in 2021." Okay? So two years ago they discovered this - "but agreed not to disclose them publicly until radio manufacturers could create patches and mitigations." And we know how that typically goes. "Not all of the issues can be fixed with a patch, however, and it's not clear which manufacturers have prepared them for customers. Motorola, one of the largest radio vendors, did not respond to repeated inquiries from Wired.

"The Dutch National Cyber Security Centre assumed the responsibility of notifying radio vendors and computer emergency response teams around the world about the problems, and of coordinating a timeframe for when the researchers should publicly disclose the issues." And as I said at the top of the show, next week is Black Hat, and all will be revealed there.

"In a brief email, NCSC spokesperson Miral Scheffer called TETRA 'a crucial foundation for mission-critical communication in the Netherlands and around the world,' and emphasized the need for such communications to always be reliable and secure, 'especially during crisis situations.' She confirmed the vulnerabilities would let an attacker in the vicinity of impacted radios 'intercept, manipulate or disturb' communications and said the NCSC had informed various organizations and governments, including Germany, Denmark, Belgium, and England, advising them how to proceed. A spokesperson for DHS's CISA [here] said they're aware of the vulnerabilities, but would not comment further. The researchers say anyone using radio technologies should check with their manufacturer to determine if their devices are using TETRA, and what fixes or mitigations are available.

"The researchers plan to present their findings at the Black Hat security conference in Las Vegas, when they will release detailed technical analysis as well as the secret TETRA encryption algorithms that have been unavailable to the public until now. They hope others with more expertise will dig into the algorithms to see if they can find other issues. So TETRA was developed in the '90s by the European Telecommunications Standards Institute, or ETSI. The standard includes four encryption algorithms, TEA1, TEA2, 3, and 4" - so I'll just call them TEA 1, 2, 3, and 4 - "that can be used by radio manufacturers in different products, depending on their intended use and customer."

Okay. So, as I said, whoa, wait, what? The four different encryption algorithms can be used by radio manufacturers in different products depending upon their intended use and customer. So if that doesn't smell fishy, I don't know what does. So Wired explains this. Wired says: "TEA1 is for commercial uses. For radios used in critical infrastructure in Europe and the rest of the world, though, it is also designed for use by public safety agencies and military, according to an ETSI document, and the researchers found police agencies that use it.

"TEA2 is restricted for use in Europe by police, emergency services, military, and intelligence agencies." Okay. So TEA1 is for commercial uses, whereas TEA2 is restricted

for use in Europe by police, emergency services, military and intelligence agencies? What's the difference? "TEA3," Wired writes, "is available for police and emergency services outside Europe, in countries deemed friendly to the EU, like Mexico and India. Those not considered friendly, such as Iran, only had the option to use TEA1. TEA4, another commercial algorithm, is hardly used, the researchers said.

"The vast majority of police forces around the world, aside from the U.S., use TETRA-based radio technology. After conducting open source research, TETRA is used by police forces in Belgium and the Scandinavian countries; East European countries like Serbia, Moldova, Bulgaria, and Macedonia; as well as in the Middle East in Iran, Iraq, Lebanon, and Syria. Additionally, the Ministries of Defense in Bulgaria, Kazakhstan, and Syria use it. The Polish military counterintelligence agency uses it, as do the Finnish defense forces, and Lebanon and Saudi Arabia's intelligence service, to name a few.

"Critical infrastructure in the U.S. and other countries use TETRA for machine-to-machine communication in SCADA and other industrial control system settings, especially in widely distributed pipelines, railways, and electric grids where wired and cellular communications may not be available."

And now, get a load of this blast from the past: "Although the standard itself is publicly available for review," meaning the paper printed standard saying this is what we're going to offer you for your radio to use, "the encryption algorithms are only available under a signed NDA to trusted parties, such as radio manufacturers. The vendors have to include protections in their products to make it difficult for anyone to extract the algorithms and analyze them." Oh, boy.

"To obtain the algorithms, the researchers purchased an off-the-shelf Motorola MTM5400 radio and spent four months locating and extracting the algorithms from the secure enclave in the radio's firmware. They had to use a number of zero-day exploits to defeat Motorola protections, which they reported to Motorola to fix. Once they reverse-engineered the algorithms, the first vulnerability they found was a backdoor in TEA1."

Okay. So first of all, huge props to these guys. No one made it easy for them to obtain the information they needed. In fact, their efforts were deliberately being thwarted at every turn by the use of requiring a signed NDA, which they were not able to agree to because they wanted to disclose it, and a secure enclave. And they needed to find zero-day exploits, brand new zero-day exploits and then use them to crack the lid off the code.

And let's also just pause for a moment to thank our lucky stars that this reverse engineering conduct has been deemed legal. If white hat hackers like these guys could be jailed for conducting research in the interest of improving the security of the products they're examining, even when doing so is not in the interest of those who are working hard to keep those secrets, the world would be far less secure, and only the bad guys would be pursuing such reverse engineering. They would not be agreeing to keep their secrets quiet. They would never be disclosing them because they would then be turning around and leveraging them. And all of the stuff that we talk about on this podcast constantly which is being reverse engineered at significant effort and cost by good guy researchers, none of that would be happening because doing so would be illegal. Thank goodness that decision was made, making it so that this kind of research is safe.

So here's what they found. All four TETRA encryption algorithms use 80-bit keys, which the researchers say, and I would agree, even more than two decades after their release, still provides sufficient security to prevent someone from cracking them. And I'll note that the keys are rotated, and they're dynamically changing. So it's not like they're just fixed 80-bit keys. They're ephemeral, so they're not around long enough for that to be a problem. But they are around for a while.

TEA1 has a "feature," in quotes, that reduces its encryption key length to just 32 bits, which the researchers were able to crack in less than a minute using a standard laptop and samples of just four cipher texts, which of course you get by putting a radio up in the air and receiving some of this encrypted communication. Brian Murgatroyd, the chair of the technical body at ETSI, the people behind this, responsible for the TETRA standard, objects to calling this a backdoor. He says when they developed the standard they needed an algorithm for commercial use that could meet export requirements - now, remember this is more than two decades ago - to be used outside Europe, and that in 1995 a 32-bit key still provided security, although he acknowledges that with today's computing power that's no longer the case. Remember, these guys, the researchers cracked the key in less than a minute.

Matthew Green, our well-known Johns Hopkins University cryptographer and professor, calls the weakened key "a disaster." He said: "I wouldn't say it's equivalent to using no encryption, but it's really, really bad." Gregor Leander, a professor of computer science and cryptographer with a security research team known as CASA at Ruhr University Bochum in Germany, says it would be "stupid," not mincing any words, for critical infrastructure to use TEA1, especially without adding end-to-end encryption on top of it. He said: "Nobody should rely on this." Murgatroyd insists the most anyone can do with the backdoor is decrypt and eavesdrop on data and conversations. TETRA has strong authentication, he says, that would prevent anyone from injecting false communication.

"That's not true," says Wetzels, one of the researchers. TETRA only requires that devices authenticate themselves to the network, but data and voice communications between radios are not digitally signed or authenticated. The radios and base stations trust that any device that has the proper encryption key is authenticated, so someone who can crack the key as the researchers did can encrypt their own messages with it and send them to base stations and other radios.

While the TEA1 weakness has been withheld from the public, it's apparently widely known in the industry and governments. In a 2006 U.S. State Department cable leaked to WikiLeaks, the U.S. embassy in Rome describes an Italian radio manufacturer asking about exporting TETRA radio systems to municipal police forces in Iran. The U.S. pushed back on the plan, so the company representative reminded the U.S. that encryption in the TETRA-based radio system they planned to sell to Iran is less than 40-bits - indeed, 256 times less than 40 bits because it's 32 bits, implying that the U.S. should not object to the sale because the system isn't using a strong key.

The second major vulnerability the researchers found isn't in one of the secret algorithms, but it affects all of them. All of them. The issue lies in the standard itself and how TETRA handles time syncing and keystream generation. When a TETRA radio contacts a base station, they initiate communication with a time sync. The network broadcasts the time, and the radio establishes that it's in sync. Then they both generate the same keystream, which is tied to that timestamp, to encrypt the subsequent communication.

Wetzels says: "The problem is that the network broadcasts the time in packets that are unauthenticated and unencrypted." As a result, you can time spoof. An attacker can use a simple device - and actually, Leo, you probably have one in your pocket.

Leo: No, I gave it to Father Robert to take to Black Hat.

Steve: Oh, yeah, that's good, he'll get some use out of it - "to intercept and collect encrypted communication passing between a radio and base station, while noting the timestamp that's initiated the communication. Then he can use a rogue base station to

contact the same radio or a different one in the same network and broadcast the time that matches the time associated with the intercepted communication. Basically, you know, resetting them to the key that he already has that was decrypted earlier. The radio is dumb and believes the correct time is whatever the base station says it is. So it will generate the keystream that was used at the time to encrypt the communication the attacker collected. The attacker recovers that keystream and can use it to decrypt the communication collected earlier.

"To inject false messages, he would use his base station to tell a radio that the time is tomorrow noon, and ask the radio to generate the keystream associated with that future time. Once the attacker has it, he can use the keystream to encrypt his rogue messages, and the next day at noon send them to a target radio using the correct keystream for that time. In other words, it was really badly designed, even in 1995. There were all kinds of holes in the system, not just secret algorithms for encryption.

"Wetzels imagines Mexican drug cartels could use this to intercept police communications to eavesdrop on investigations and operations or deceive police with false messages sent to radios. The attacker needs to be near a target radio, but the proximity is only dependent on the strength of the rogue base station's signal and the terrain." He says: "You can do this within a distance of tens of meters. The rogue base station would cost $5,000 or less."

So ETSI's Murgatroyd downplays the attack, saying TETRA's strong authentication requirements - oh, boy - would prevent a non-authenticated base station from injecting messages. Wetzel disagrees, saying TETRA only requires devices to authenticate to the network, not to each other. The researchers didn't find any weaknesses in the TEA2 algorithm used by police, military, and emergency services in Europe, but they did initially think they found another backdoor in TEA3. Given that TEA3 is the exportable version of TEA2, there was good reason to believe it might also have a backdoor to meet export requirements.

Anyway, we basically have a system which is full of holes, has been used for, what, 28 years, since 1995, is known to be insecure, never received the upgrading that it should have received. But as I said, that never happened. So, as I and Wired noted, in eight days all the wraps will be coming off of this when the research team presents their work and findings during Black Hat in Las Vegas. With TETRA we have a legacy, encrypted, radio communications system being widely used today throughout the entire world, including in the U.S. And it not only contained multiple really exploitable flaws that were only fixed after security researchers cracked it open and shamed its creators with the threat of disclosure, and even now they're not actually saying, okay, yeah, you got us, you're right, it also contained deliberately weakened encryption which most of the world was given to use while some agencies knew of the weakness and were apparently leveraging that knowledge for eavesdropping.

And now we learn that the ETSI group who did all of this has replaced their earlier flawed work with more of the same. Keeping their encryption secret after rotating the original TEA1 through 4 ciphers out, there are now new ones; and they, too, are kept secret. Even though we have well-vetted, well-tested, well-functioning, lightweight, high-performance encryption. Nobody should be rolling their own any longer. It's just crazy. Why would anyone ever trust these people?

Leo: So true. This reminds me of SS7, although SS7 is still around, the sideband that is totally hackable in every phone. It's still around just because you can't - it's too hard to change; right?

Steve: Right. Well, we do have the requirement for encryption intersystem, but that's what has not happened.

Leo: Right.

Steve: Intrasystem encryption has happened, and they're supposed to be doing intersystem. But the problem is apparently they're making too much money out of spam.

Leo: Yeah, right. There you go.

Steve: They really don't want to limit it.

Leo: There you go. They don't want to limit it, yeah. Ah.

ROBOT: Does not compute.

Leo: That's just the way it is. Even the robot has an opinion on that one. Well, that concludes this thrilling, gripping edition of Security Now! as we edge into our 19th year, a couple more weeks.

Steve: Coming up on it.

Leo: Wow. Only 66 episodes left. Guess we're counting down, too. Steve Gibson's at GRC.com, which is proudly not in the cloud. All you have to do is go to GRC.com, and then you will see all sorts of good stuff, including SpinRite, the world's best mass storage maintenance and recovery utility. You need this if you've got hard drives or solid-state drives. Version 6.0 is still there, but it is soon to be replaced by 6.1. You will get a free upgrade when 6.1 comes out if you buy today. GRC.com. You can also get the show there.