Zombie Software

Description: This week we answer questions which arose during the past week: When is an attack not an attack? When our AI overlord arrives, how shall we call him? Why has Italy said NO to ChatGPT? What does Twitter's posting of its code to GitHub tell us? Why is India searching for commercial spyware less well-known than Pegasus, and what does the Summit for Democracy have to say about that? Has the FDA finally moved on the issue of medical device security updates? And seven years after the first "Hack the Pentagon" trial, the Pentagon remains standing; or does it? Then, after addressing a quick bit of miscellany, listener feedback, and an update on my ongoing work on SpinRite, we use CISA's KEV database to explore the question of how exactly we define "Zombie Software," and answer the question of whose brains will the zombies eat?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-917.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-917-lq.mp3

SHOW TEASE: This week on Security Now!, Mr. Steve Gibson joins me to share some great information about the wonderful world of cybersecurity. And I've got to tell you, it was a little bit scary at first; but it got better, huh. So what's going on with AI? Are we going to name our AI overlords once they arrive? Italy decides to say no, there's not going to be any more ChatGPT. We're shutting that down. Oh, boy. What does that mean? The Pentagon said, "Hey, hack us." And you know what? It worked. It was not a bad idea. So more details on that and doing bug bounties. And lastly, Zombie Software. Oh, this is such a great story from Mr. Gibson. You don't want to miss it. Got to stay tuned.

ANT PRUITT: This is Security Now!, Episode 917, recorded Tuesday, April 4th, 2023, hosted by Ant Pruitt and Steve Gibson: Zombie Software. Hey, what's going on, everybody? Hey, this is Security Now! here on TWiT.tv.

Steve Gibson: Wait a minute. You're not Leo.

ANT: Yeah, yeah, yeah. I'm not Leo Laporte. I am Ant Pruitt, y'all. I'm going to be hosting today with the one-and-only Mr. Steve Gibson. Steve, how are you, sir?

Steve: Hey, Ant. Great to be with you for the first of our three weeks while Leo is off on some trip. And we're going to hold the fort down here and just keep the ball rolling. So this is...

ANT: We shall do our best, sir.

Steve: Yeah. Security Now! #917 for the 4th of April. And as we've been doing recently, we're going to answer a bunch of questions which have arisen over the past week. When is an attack not an attack? When our AI overlord arrives, how shall we call him? Why has

Italy said no to ChatGPT? What does Twitter's posting of its code to GitHub tell us? Why is India searching for commercial spyware less well-known than Pegasus? And what does the Summit for Democracy have to say about that? Has the FDA finally moved on the issue of medical device security updates? And seven years after the first Hack the Pentagon trial, the Pentagon remains standing. Or does it?

ANT: Oh, boy.

Steve: Then, after addressing a quick bit of miscellany, some listener feedback, and an update on my ongoing work on SpinRite, we use the CISA's KEV, the Known Exploited Vulnerabilities database, to explore the question of how exactly we define "zombie software," and answer the question of whose brains will the zombies eat.

ANT: Oh, boy. Oh, boy.

Steve: Those questions and more answered this week.

ANT: I've got to go ahead and ask, Mr. Gibson. You're saying what happens when our AI overlords arrive, what are we going to call them?

Steve: That's right.

ANT: I thought they were already here.

Steve: Well, there's some concern about that. And in fact unfortunately Washington has been caught flat-footed and are trying to say, wait a minute, uh, we don't understand what any of this stuff is, so everybody should stop. Anyway, we'll be talking about that today.

ANT: Oh, joy. Oh, joy. Well, folks, this is going to be a fun episode. Mr. Steve, the Picture of the Week.

Steve: Our Picture of the Week. I got a kick out of this. Actually, we have a couple things that are sort of about this topic in general. So we have a single-frame cartoon. And a woman is staring at her computer screen, talking to someone standing behind her. And the balloon above her head says: "Congress finally has a better understanding of how TikTok works. Almost half of them now know it's not a breath mint." So, yes, not Tic Tac, folks. We didn't say Tic Tac, we said TikTok. Oh. Oh, not - okay, yeah. So that's right. We're going to give the politicians control of technology and see how well that works.

ANT: Does a younger person know what a Tic Tac is, or a TikTok is?

Steve: Yeah, Tic Tacs are still around. Yeah, probably.

ANT: They still are? Okay. All right.

Steve: Yeah. I think in general where we have old fogies doing podcasts, there is a danger of a generation gap in some of the terminology. Sometimes when I say somebody was asleep at the switch, I think, does anyone know what that means? I mean, old people know that we're talking about train tracks. But I'm not sure that - what? What? Asleep at the switch? What switch? What are you talking about?

ANT: What switch? Oh, goodness.

Steve: Okay. So recall that two years ago all of the fuss, bother, and more than a little bit of horror arose over the attack at a Florida water treatment plant in Oldsmar, Florida. Now, the report that the world received was that a caustic and potentially poisonous concentration of lye had been dumped into the facility's municipal water treatment plants, remotely over the Internet, that is, the control of the water treatment plant's equipment had been done remotely. And that was conducted, we were told, by a disgruntled former employee whose remote access credentials had not been canceled and deleted after his departure.

And the story had so much flair and detail. Recall, because we talked about this a lot then, that the reports of the incident described how a worker at the plant saw his computer being remotely accessed and controlled, like right in front of him. His mouse, we were told, moved to open functions to control the plant's water treatment protocols, and then the amount of sodium hydroxide, which is the chemical formula for lye, in the water was changed from 100 parts per million, which was what it was supposed to be, to 11,100 parts per million. Whereupon the operator, who saw this horror, immediately reduced lye's level back to its proper level and alerted his supervisor.

And everyone got in on the act. The hack immediately gained worldwide notoriety after the local Pinellas County Sheriff held a press conference which, in turn, prompted an investigation led by the FBI and the U.S. Secret Service, as well as a joint federal advisory warning water treatment facility operators throughout the country of the dangers they faced from hackers, and urging them to upgrade their security systems.

Okay, now, gristy as that story was for our mill, what if that's not what happened at all? That's right. According to former Oldsmar City Manager Al Braithwaite, who was with the city at the time, the incident was never a hack or an attack at all. It was just a case of an employee, still with the water treatment facility, mistakenly clicking the wrong button.

ANT: Wait, what?

Steve: I know. I know. Then alerting his superiors to his error. Now the former City Manager Braithwaite describes the incident as a total "non-event," which was resolved in two minutes. But he said law enforcement and the media for some reason seized on the idea of this being a cyberattack and just "ran with it," he said. The attention resulted in a four-month FBI investigation, which Braithwaite said reached the same conclusion, that employee error alone was to blame. And somehow we never heard the result of that after four months. That sort of slipped through the cracks.

ANT: Wow.

Steve: Now, okay, yeah, wow. I mean, this was a big deal. Everyone was, like, you know, alarms were ringing, and water treatment plants elsewhere were being put on alert. And actually there might have been some upside to all of this because all municipal water treatment plants likely did and probably do still need to be more alert and aware of threats to keep their security tight because the thing that was so galvanizing about this, it was like, whoa, you know, poisoning the municipal water supply, that's a big deal.

So if the report of this dastardly attack-which-wasn't happened to serve to get some inattentive management at other treatment plants in the nation to tighten up their grip, change employee discharge policies, perhaps delete a bunch of old and unused access credentials and so forth, then this bizarre inflation of this minor event probably helped security overall in the long term. So it's not the way any of us want to be used for having security improved, but it's better to know I think ultimately what happened than not.

ANT: Yeah, somebody had to really take a sacrifice, if you will, to help everybody else raise awareness. And this is good here. My only thought on that, though, Mr. Gibson,

when that happened, I want to say I heard it was like through RDP or something like that.

Steve: Yeah, apparently it was Team Viewer. And they had...

ANT: Team Viewer, that's what it was, Team Viewer. And so when you're using Team Viewer, if someone pops onto your screen and moves the mouse around, can't you move your mouse back to try to take control and kill the session? I'm like, what's going on here? So that story just kind of...

Steve: So your point is that it never really made that much sense to you; right?

ANT: Right. Right.

Steve: Right. Right.

ANT: Geez.

Steve: Okay. So it's not going to be news to anyone that the sudden explosive popularity, adoption, and use of Large Language Neural Network Models such as ChatGPT - and I know, Ant, you've been following along a lot of this stuff - has caught pretty much everyone flat-footed. Everyone's running around trying to figure out what it means, whether it's some sort of apocalypse, and like the subject that the public should be protected from.

So a couple of weeks ago, the U.S. Chamber of Commerce issued a report from the "AI Commission." Apparently we have an AI Commission now. The report claims to highlight the promise of AI - you almost want to put "promise" in air quotes because they're, like, they're not so sure - while calling for a "Risk-Based Regulatory Framework." And the subhead of the report says "Report Finds Policymakers Must Enforce Existing Laws," okay, whatever they would be, and "Develop Policies to Steer the Growth of Responsible, Ethical AI." Okay. Ethical AI? We've got a lot of undefined terms here. So what's actually happening is kind of clear to see. The politicians and the bureaucrats have no idea what any of this is. I mean, most of us don't; right?

ANT: Right.

Steve: I mean, it's kind of like a "wait and see what happens next" because this is so new.

ANT: Right.

Steve: But not knowing scares the crap out of them. And these people, you could say, confusing TikTok and Tic Tac, they're not the brightest bulbs, as our Picture of the Week explained. So these people are apparently the ones who have appointed themselves to shepherd us into the land of ethical AI. Again, as opposed to unethical AI? I don't know.

ANT: Right, right.

Steve: Okay. So here's what the U.S. Chamber of Commerce is thinking. They wrote: "The use of artificial intelligence is expanding rapidly." They like to start off with the obvious. "These technological breakthroughs present both opportunity and potential peril." Oh. "AI technology offers great hope for increasing economic opportunity, boosting incomes, speeding life science research at reduced costs, and simplifying the lives of consumers.

"With so much potential for innovation, organizations investing in AI-oriented practices are already ramping up initiatives that boost productivity to remain competitive. Like most disruptive technologies, these investments can both create and displace jobs." Uh-oh. "If appropriate and reasonable protections are not put in place, AI could adversely affect privacy and personal liberties, or promote bias. Policymakers must debate and resolve the questions arising from these opportunities and concerns to ensure that AI is used responsibly and ethically.

"This debate must answer several core questions: What is the government's role in promoting the kinds of innovation that allow for learning and adaptation while leveraging core strengths of the American economy in innovation and product development? How might policymakers balance competing interests associated with AI - those of economic, societal, and quality-of-life improvements - against privacy concerns, workforce disruption" - there that is again - "and built-in biases associated with algorithmic decision-making? And how can Washington establish a policy and regulatory environment that will help ensure continued U.S. global AI leadership while navigating its own course between increasing regulations from Europe and competition from China's broad-based adoption of AI?"

ANT: Oh, boy.

Steve: And of course that's a concern.

ANT: There's a buzzword.

Steve: Uh-huh. You hog-tie U.S. innovators with needless fuzzy and ill-conceived regulations, and the rest of the planet is going to leapfrog the U.S. with their high-speed plans for unethical AI, apparently.

So they said: "The United States faces stiff competition from China in AI development. This competition is so fierce that it is unclear which nation will emerge as the global leader, raising significant security concerns for the United States and its allies." Oh. We might not win this. "Another critical factor," they said, "that will affect the path forward in the development of AI policymaking is how nations historically consider important values, such as personal liberty, free speech, and privacy.

"To maintain its competitive advantage, the United States and like-minded jurisdictions such as the European Union" - you know, we like them - "need to reach agreement to resolve key legal challenges that currently impede industry growth. At this time it is unclear if these important allies will collaborate on establishing a common set of rules to address these legal issues, or if a more competitive and potentially damaging legal environment will emerge internationally." Well, that's interesting. So, like, are we going to have different rules and regulations and different definitions of ethical AI? We don't know.

ANT: Right. Who knows?

Steve: Nobody knows any of this. "AI has the capacity to transform our economy, how individuals live and work, and how nations interact with each other. Managing the potential negative impacts of this transition should be at the center of global public policy. There is a growing sense that we have a short window of opportunity" - some people would say narrow, but okay - "to address key risks while maximizing the enormous potential benefits of AI," which we don't understand at all. No, they didn't say that at the end. Benefits of AI.

Wow. So, you know, I doubt that anyone knows what any of that really means; right? But, like, okay, somebody must have said you need to publish a position paper on, uh, something.

ANT: We need to make a statement. Whatever that statement is, we need to make a statement.

Steve: Right.

ANT: That's what they did; you know?

Steve: Yeah. So it's clear it's coming on like a speeding train. And it's unclear to me that there's any way to slow it down. Obviously their concern is, if they impede the innovation in the U.S., then countries that don't impede their researchers, and we know who they are, they're going to get ahead of us and take over, and we're going to all be working for them. So that's not good. Anyway...

ANT: And again, to their credit, like most innovative technology, things can be great for mankind. But then in the hands of the wrong actor it could cause a lot of problems. I'm glad they're not that naive.

Steve: Oh, my god, it's what this podcast is all about; you know? Technology which is inherently neutral, we get a lot of leverage and benefit from it. And, boy, is there an awful lot of bad guys who are spending so much industry on just, you know, trying to subvert it.

ANT: Unfortunately.

Steve: So I should mention that at least one country, bizarrely, has chosen to adopt a much more, like an immediate kneejerk position on AI. Believe it or not, Italy has banned ChatGPT. Okay, I'm not kidding. The news is that the Italian Data Protection Authority has issued a temporary ban on ChatGPT as the agency investigates a possible breach of its GDPR regulations, with the agency accusing the OpenAI service of "unlawful collection of data."

Okay, now, and that appears to be the point of this. Someone has apparently freaked out Italian legislators because they believe that "The ChatGPT" is sucking up all of the Internet's data on their Italian citizens without any regard for their GDPR-guaranteed privacy.

ANT: Hmm.

Steve: Now, yeah, it's odd. The only thing I could find to bring some clarity to this was a press release originally written in Italian and poorly translated into English. Now, of course, the translation would be great if they had just let ChatGPT do it. But noooooo. So here's what was announced, from which you can get a sense for their somewhat hysterical concern. And again, pardon the translation, but this was a dumber bot did the translation.

So it reads: "No way for ChatGPT to continue processing data in breach of privacy laws. The Italian SA imposed an immediate temporary limitation on the processing of Italian users' data by OpenAI, the U.S.-based company developing and managing the platform. An inquiry into the facts of the case was initiated, as well.

"A data breach affecting ChatGPT users' conversations and information on payments by subscribers to the service had been reported on March 20th." Okay, now, that's like neither here nor there. That's got nothing to do with this, right, and AI. But okay.

Then they said: "ChatGPT is the best known among relational AI platforms that are capable to emulate and elaborate human conversations. In its order, the Italian SA highlights that no information is provided to users and data subjects whose data are collected by OpenAI; more importantly, there appears to be no legal basis underpinning the massive collection and processing of personal data in order to 'train' the algorithms on which the platform relies. As confirmed by the tests carried out so far, the information made available by ChatGPT does not always match factual circumstances, so that inaccurate personal data are processed.

"Finally, the Italian SA emphasizes in its order that the lack of whatever age verification mechanism exposes children to receiving responses that are absolutely inappropriate to their age and awareness, even though the service is allegedly addressed to users aged above 13, according to OpenAI's terms of service."

Okay. Whew. So this is all a bit weird, right? So far as anyone knows, no private data has been or is being harvested by ChatGPT. The system is simply ingesting vast quantities of publicly available information and merging it into an evolving large language model which has an interactive conversational interface. In a very similar fashion, Google's "Googlebots" crawl the web, following links and indexing everything they encounter. So there's really no difference between what Google collects and what ChatGPT collects.

ANT: Exactly.

Steve: The difference, yeah, is in the accessibility and presentation of the information. Google maintains a massive index, whereas ChatGPT and similar systems maintain a massive model. Anyway, I imagine that things will likely calm down in Italy once someone who knows something actually talks to those who need to know more than they apparently do at the moment. One thing seems clear, though, which is that Italy would not want ChatGPT to be ignorant of everything that country has to offer. It would represent a huge blind spot that would ultimately hurt them economically as the world's inevitable adoption of and reliance upon this emerging technology continues. If someone asks their favorite AI bot for recommendations about where to stay in Rome, you don't want the world's AIs asking, "What's a Rome?"

ANT: And totally, totally screwing up that economy there from all of those visitors.

Steve: That's right. Where did all the tourists go? Well, you told ChatGPT to forget about us, and it did.

ANT: We spoke about another AI offered by Adobe here on the network last week. And the key difference between what Adobe's doing and the other AIs out there is Adobe is scraping data based on information submitted to them directly from their subscribers to Adobe so they could figure out text phrases, text strings or what have you to apply it to art. I wonder if that would be a problem in Italy or somewhere else over there in Europe with concerns of GDPR. Even though there's a consent by the users of Adobe products to say, you know what, whatever you create in our software, we're going to use it to train our models. But you can't opt out of it. So I wonder if there's any connection there between what's going on in Europe there versus here.

Steve: Also this raises an interesting question, or you have, which is I wonder if the users' input into ChatGPT influences the model. I mean, we know that it influences in the

short time; right? And after a while it goes insane, and you have to restart it, you know, because it loses its mind, or it starts hallucinating.

ANT: Right. They say it hallucinates or something. That's what they're saying now.

Steve: Yes, exactly. And so I wonder like if children were talking to it and giving it information which is not on the public Internet, if the information they voluntarily disclose ends up being captured in the long term. I don't know one way or the other. But anyway, it just seems like everyone's having a different reaction to the craziness of this.

But, you know, the reason I was put in mind of this, if you ask the AI where's a good place to stay in Rome, is I was watching the live stream sometime in the last week, I don't remember which podcast it was on, but Leo was talking about ChatGPT-4 and demonstrating it, and he literally, he said, "We're going to be" - I think it was in Rome. "We're going to be in Italy. What would you recommend as an itinerary?" Well, it spit out, like, half a page. And Leo said, "Wow, I've got to copy that. That looks really good."

ANT: Right.

Steve: And so, like, literally, it had some good ideas that he hadn't considered. And in fact it estimated how long the various things would take.

ANT: That's awesome.

Steve: So that it shows you, like, where during the day you would be. And Leo said, "Oh, I've got to make a copy of this." So the world is clearly changing.

ANT: That would be awesome to have that, though, sort of a predetermined itinerary. And I'm sure all of the merchants there in the locale would appreciate your patronage, too.

Steve: That's right. So last Friday, Twitter officially posted the source code that it uses for selecting which tweets users will see. This of course has been the source of much controversy and handwringing and, you know, social media postings everywhere but Twitter and on Twitter, since we've learned that the tweets people see turns out to have an outsized effect upon the beliefs they hold. So for anyone who may be interested, it's now there on GitHub. Since I'm not a Twitter surfer, I'm not very curious.

But I'm sure we'll see some analysis of the algorithm once those who do care have invested the time in learning what it all does. Ars Technica shared their take after a very quick look, and from what they saw it mostly looked like what any student of social science would predict. Watch what each person does, what they click, what they search for, how long they remain, how far down they scroll, and so forth. Collect every possible meaningful scrap of data from their interaction, figure out what sorts of things they want to see, then show them more of that.

ANT: Well, that seems logical; right?

Steve: That's basically the algorithm; right. And of course, you know, we've talked a long time ago that the so-called "filter bubble" effect, right, is where if you show people things that they've seen, that they've already demonstrated they want to see, you end up amplifying their positions and hardening their positions and essentially radicalizing them to a position, as opposed to just showing them everything going on and exposing them to all possible positions.

So one observer did note that Twitter is actively burying tweets about Russia's invasion of Ukraine. Apparently that's in the algorithm there. So that mystery is resolved for anyone who might have wondered where those tweets went after Twitter became an Elon property because it's been observed that that suddenly disappeared. We talked about how Elon's Starlink satellite system was turning out to be hugely instrumental in helping to keep Ukraine connected to the rest of the world as Russian missiles continued to batter Ukraine's communications and other utility infrastructures.

So, you know, he seems to be helping Ukraine. I don't understand why burying those tweets, like specifically biasing against that is a thing that he would want to do. But then it's a little hard to be in Elon's head.

ANT: Right.

Steve: So why he does anything. So Ant, I think we should tell our listeners about our second sponsor.

ANT: Yeah.

Steve: And then we're going to talk about Israel's Pegasus and why India doesn't want to use it.

ANT: Oh, boy. Another, 'nother fun story coming up.

Steve: Okay. So this one, I guess if I was going to title this story, I'd say, okay, we know it's illegal. How much will it cost? Now, if that was a criminal saying that he knows it's illegal, what'll the price be, that would be one thing. But this is the government of India. Who is reportedly seeking bids from as many as a dozen, I guess we would call them "me-too smartphone surveillance software purveyors" because the current preferred go-to spyware, the well-known Pegasus from Israel's NSO Group, has become too well known. The use of Pegasus now is being frowned on by several annoyed nations, including the current U.S. administration, which has been quite vocal about it. So now it appears to matter which super-secret hidden illegitimate illegal surveillance spyware a country chooses to use to implant into the smartphones of its surveillance targets. Which again, like, what world are we in?

The news coverage of this in the Financial Times stated that India is hunting for new spyware with a lower profile than the controversial Pegasus system which has been blacklisted by the U.S. government. Indian defense and intelligence officials have decided to acquire spyware from less well known and less publicly exposed competitors of the NSO Group, and they plan to spend around $120 million, they figure it's going to cost, for the replacement software which they feel they need. And again, this is, like, illegal spying on other people's phones.

ANT: I'm sitting here. I know I'm not supposed to be laughing at this, sir, but this is hilarious.

Steve: It's mind-boggling.

ANT: Spyware for hire.

Steve: Yeah. These lesser-known alternatives are doubtless gleeful that Pegasus became so popular that it's now being frowned upon by those with a reputation to maintain. Anyway, this all seems so bizarre since it's, as I said, it's illegal. But it's what everyone does.

The Financial Times noted that India's move shows how demand for sophisticated, unregulated, and illegal technology remains strong despite growing evidence that governments worldwide have abused spyware by targeting dissidents and critics. Right? That's only supposed to be for terrorists and, like, you know, you're only supposed to use it against criminals. But it's being used against non-criminal critics of the government.

India, for their part, has never publicly acknowledged ever being a customer of NSO, but they would like their money back. However, Pegasus spyware has been found on the phones of secular journalists, left-leaning academics, and opposition leaders around India, and this sparked a political crisis, ultimately. As we know, Pegasus turns phones into surveillance devices which can collect encrypted WhatsApp and Signal messages surreptitiously. And of course it doesn't crack their encryption, it's just there on the phone, so it gets it before it's encrypted and sends it off to wherever Pegasus sends things. This is what the spooks inside governments want, claim they need, and are obviously willing to pay for.

And I loved this: India's Modi government officials have grown concerned about the "PR problem" caused by the ability of human rights groups to forensically trace Pegasus - it's a PR problem, yeah - to forensically trace Pegasus, as well as warnings from Apple and WhatsApp to those who have been targeted, according to two people familiar with the discussions. So, yeah, we're looking for something, they're saying, that is just as effective, but is also much less well known. So where do we wire the funds?

ANT: Wow.

Steve: Unbelievable. You know, I mean, so we were talking, you know, the promise of being able to spy on people who are critical to your government has turned these nations into outlaw nations where they're explicitly using illegal, forbidden, you know, who's worried about ethical AI? What about ethical nations?

ANT: Right.

Steve: And governments. We already don't have that. Wow.

ANT: Oh, man.

Steve: Okay. Meanwhile, with timing that you couldn't make up, last Wednesday the so-called "Summit for Democracy" was held, after which a joint 12-nation statement condemning the proliferation of exactly the sort of commercial spyware that India is currently on the market for, happened. The governments of, in alphabetical order, Australia, Canada, Costa Rica, Denmark, France, New Zealand, Norway, Sweden, Switzerland, the UK, and the U.S. published a joint statement on planned efforts to counter the proliferation of this commercial spyware.

Agreed countermeasures include tightening export controls, better information sharing to track the proliferation of such tools, and engaging with additional partner governments to reform and align policies on the use of spyware by their agencies. Maybe they ought to give India a call. Given how widespread the practice has become at the highest levels of nation-state intelligence tradecraft, you really wonder how many of the signatories of that letter are condemning on one hand and deploying exactly that spyware from their own agencies on the other. So it's just, you know, it's just too tantalizing.

ANT: Tantalizing, scary, sad, all at the same time.

Steve: So on a happy front.

ANT: Thank you.

Steve: Guidance has been issued - yes, we need some happy news. Guidance has been issued by the U.S. Food and Drug Administration last Thursday, March 30th, which explains now that any and all medical devices being submitted to the FDA for approval will from now on be required to meet specific cybersecurity requirements. This is a first. We've never had this before. These new requirements are part of the so-called "Consolidated Appropriations Act" which was signed into law late last year, 2022. That new law contained a section titled "Ensuring Cybersecurity of Medical Devices," which is exactly what it sounds like.

ANT: Right.

Steve: According to the FDA, submissions for new medical devices will need to include specific cybersecurity-related information such as the description of a plan for identifying and addressing vulnerabilities and exploits within a reasonable time. Companies must also provide details on the processes and procedures for releasing post-market updates and patches that address security issues, including through regular updates and out-of-band patches in the case of critical vulnerabilities. The information provided to the FDA must also include, and this is something that we're seeing more and more, a software bill of materials, an SBOM.

ANT: SBOMs.

Steve: For commercial - yes - commercial, open source, and off-the-shelf components. So that they will be disclosing what tools and components and subsystems were used to build this whole thing. So again, to create some visibility into the individual offering. The requirements apply to any "cyber device" that runs software, that's defined as a device that runs software, has the ability to connect to the Internet, and would be vulnerable thereby to cyber threats. The new cybersecurity requirements do not apply to submissions prior to March 29th, 2023, so like not prior to last Thursday. And the FDA will not reject out of hand applications solely failing this new requirement until October 1st.

So you've got some, you know, a grace period, essentially, from this announcement and this going into effect kind of semi-officially last Friday until October 1st. And in the interim it will provide assistance to companies up until that date. However, starting with October 1, the agency may start rejecting premarket submissions that do not contain the required information. So...

ANT: You know, I love this, that they make this very, very clear for everybody. So there's no wiggle room.

Steve: Yup. Nope.

ANT: Just simplify this. We're trying to make things better for everybody. And here we have the government stepping in and making it just a simple, simple process. Thank you. Thank you.

Steve: Right. We're announcing it now. We're giving you till October 1st. We'll help you until then, if anything about this is confusing. And no reason why it would be. You know?

ANT: Right.

Steve: In our industry, in the security industry, the stories, for example, about insulin pumps being hacked and taken over, you know, have almost become a meme in the

industry. So this welcome legislation means that devices can no longer be sold and forgotten. Now, this works easily today for medical devices, the marketing and sales of which are already highly regulated. So like there's already FDA regulation mumbo-jumbo and paperwork and analogies and all that in place. But it does seem likely that somewhere in the future, and I don't know how far in the future we're talking, anything that can connect to the Internet may need similar mandatory functionality. Note that in the FDA's definition of "cyber device" it meant could connect to the Internet because...

ANT: It doesn't get any simpler than that.

Steve: Right.

ANT: You know, it doesn't get any simpler than that. Just keep it that way. None of that like legal jargon mumbo-jumbo that you see in all the other types of services. This is great.

Steve: Yeah. The problem, of course, is that with non-medical devices, we don't already have any kind of regulatory framework in place. It's just like, you know, it's the Wild West out there. And so the question is will that eventually happen. It would be better if it was done voluntarily by manufacturers, and if consumers were to insist upon those features. Unfortunately, most consumers feel that updates to Windows are a big pain in the ass. It's interfering with their work and their job.

No one who listens to this podcast thinks that updates are a problem. It's like, the only question is when do we update? Do we wait to see if it causes any problems and so forth. But most consumers don't have anything like that appreciation. It's like, because they don't see a problem that's being fixed. They just see that, you know, their computer had something spinning on the screen for an hour, and then they finally were able to get back to work.

ANT: With this requirement in place, what do you think would be the biggest pushback, if any, from the OEMs? Would it be listing the SBOM or something else? Would there be any pushback for this security requirement?

Steve: Certainly the software bill of materials would require some disclosure which otherwise would not be necessary.

ANT: Right.

Steve: So I remember, what was it, I think it was maybe it was the Kindle when I first - because the Kindle was from Amazon, was based on Linux, and a whole bunch of public stuff. I remember seeing - oh, I know. It was the requirement to disclose the licenses of the open source software that the project was using. And you looked through it, and it was like, well, Amazon, did you write anything?

ANT: Nothing.

Steve: So glue this together from stuff you've got on GitHub. I mean...

ANT: Skinned it.

Steve: It was crazy. Yeah. Exactly. Oh, stamp a label on it.

ANT: Wow.

Steve: Okay. So seven years ago, in 2016, we covered here on the podcast the U.S. Department of Defense's first halting, experimental, we're not sure about this, "Hack the Pentagon" bug bounty effort. No one was sure it would work, least of all the skeptical government bureaucrats, but work it did. Since that program's launch, ethical hackers have helped the DoD find and fix more than 2,100 vulnerabilities, many of them which were very scary and good to find, which were identified by more than 1,400 hackers.

The news today is that the continually useful and successful program now has a dedicated website at www.hackthepentagon.mil, which the DoD will be using to continue to educate additional branches of the sprawling U.S. government about the benefits available from allowing good-guy hackers to take a crack at cracking. The site will also continue to seek and sign up talented new hackers for the continually expanding program. So bottom line, this was a win. The government's networks and the DoD military government networks are more secure as a direct result of inviting ethical hackers, not anybody, you just can't go attack the government, you've got to sign up first and identify yourself and say, uh, may I please...

ANT: At least you'd better identify yourself.

Steve: Yeah, exactly. It doesn't work if you say, oh, but wait a minute, I was hacking the Pentagon, and they said, "We know." Anyway, no, you need to let them know you're going to do that first. So it's cool that it is succeeding. And in general that's what we're seeing. We're seeing, and I talked about this a couple weeks ago, in general bug bounty programs are a win for the security industry. They have become, just as much as security updates have become, you know, they are now part of the ecosystem, and we're better for it.

Speaking of being better for it, since Firefox 110, which everyone should have by now - I'm at 111.0.1. But ever since 110, Firefox has provided a new built-in facility for showing third-party DLLs, meaning those not signed by either Mozilla or Microsoft, which have arranged, sometimes by hook or by crook - and I mean "crook" in that sense - to have themselves injected into Firefox's address space. If you're a Firefox user, you can type up in the address bar "about:third-party," hit ENTER, and you'll be looking at this new page. When I did that I only had three DLLs, all which were signed by Intel. Since I'm running on an Intel NUC machine, that wasn't surprising. By clicking on the little folder icon to the right of each of the DLLs' names, Windows Explorer will be opened on that file, and it's then possible to right-click to look further into its properties. I did that. And sure enough, in my case, the three DLLs were Intel graphics drivers signed by Intel. And actually they were also co-signed by Microsoft, so they probably shipped along with Windows 10 in that instance.

So this is very useful from a security standpoint since injecting DLLs into another process's process space and very usefully into a browser's process space is something that malware likes to do because it allows it to steal things like the passwords that are being entered, either by you or your password manager, into the browser fields, and the crypto addresses that you may be copy-and-pasting in order to send money to various places. But Mozilla's primary motivation was to help identify misbehaving DLLs that might be causing Firefox to become unstable and to crash. So in that sense it was self-defense is really what was going on.

ANT: Right.

Steve: You know, like people were reporting crashes. And so Mozilla said, uh, and I'm sure they initially said type some bizarre command so that you can tell us what DLLs have been injected into the browser. And they decided, okay, let's just make this a form, you know, a fully public known UI, about:third-party. Now you can see the DLLs. And for

DLLs that are not signed by Mozilla or Microsoft, and in this case these were co-signed, mine were, you'll find a little - another option there, a little red cross-out symbol that allows you to deny that DLL's ability to inject itself. So you're able to turn off, to block the injection by DLL in order to see whether that might cure a problem that you've having. So they've also made it nicely diagnostic.

ANT: End-user remediation is what that sounds like; right?

Steve: Right, right. Of course, nobody would know it's there until, you know, you've complained. Or, you know, certainly in forums with other knowledgeable Firefox users. They might say, oh, yeah, type about:third-party, and you'll find out. And then, you know, turn that stuff off. You don't know what it's for. And remember, from a security standpoint, if something looks suspicious, then definitely go explore what that is that's been injected into your browser's process space.

Okay. So not surprisingly, a great many of our listeners wrote after last week's Microsoft rant on my part. Rather than share with everyone here what was essentially the same sentiment expressed many times in the feedback that I received, I'll just share one which is representative of all. Matthew Hile tweeted me. He said: "Steve. Listener since Episode 1 and SpinRite alpha tester. Thanks." He said: "I think you need to reconsider last week's rant. For years you have decried outdated, unpatched servers stuck in the closet." I absolutely have. "You have spoken positively of governments' recent efforts to locate and report vulnerable servers. You have asked what it would take to get those servers updated, and would it ever happen. Well, Microsoft's refusal to accept email from unpatched and unsupported Exchange servers is clearly a very dramatic way to accomplish your goal."

ANT: Huh.

Steve: "Regardless of the dangers from email from those systems, this will force users to either upgrade or move from known vulnerable systems to a less vulnerable one, making the Internet safer for all." Okay.

ANT: That's an interesting point. An interesting perspective, I should say.

Steve: Of course. And I'm bringing it up because I completely agree with what Matthew wrote, and with what everyone who wrote something similar said. I agree 100%. Really. This is without question an absolutely powerful and doubtless effective means for Microsoft to force the upgrading of their older Exchange servers. As we said last week, refuse to accept any email from them, and those who are still running them will get the message. And toward the end of last week's rant I did say exactly that, though I would not fault anyone for missing it since it certainly wasn't my main thrust.

What I said toward the end of my rant was: "No one using any Microsoft Exchange Server software will ever again be able to fail to keep it updated, nor to avoid the purchase of future licenses forever. I celebrate that idea from here forward, since keeping software updated, especially Exchange Server, is a good thing." So anyway, I just wanted to make sure that everyone understood that I really do appreciate that aspect of this mess. But in my mind, the fact that it would be effective still doesn't make it right because Microsoft is rendering otherwise useful servers unuseful, claiming that they represent a danger to them, and clearly what would generate revenue for them. So, you know...

ANT: Let me push back on you.

Steve: Okay.

ANT: Let me push back on you here. So you said "clearly useful servers." What would those uses be if they're unpatched and not necessarily secure? How would they be considered useful?

Steve: Well, okay. So there are Exchange Server 2007, 2010, and soon to be 2013, which actually expires now we're in April. It's end-of-life. So they're useful because they are online, and they are receiving and sending email. They're doing the job for their licensors. And so it is certainly the case that - and in fact this is actually what we're going to - this is the gist of our conversation we're about to have about zombie software, is what does it mean when software has left its useful service life, you know, like what does that really mean? My argument is Microsoft is claiming that those servers represent a danger to them because they can send malicious email. Well, every email server can send malicious email.

ANT: Okay. Good point.

Steve: Well, that's the nature of email. And so Microsoft is singling them out, for one thing. Why? Oh, well, because what's the shortest path for somebody who has an out-of-date server that suddenly is no longer able to send email to any Office 365 or Outlook.com user? Well, the shortest path is to upgrade that server from whatever they have to the current Exchange Server. And that's not free.

ANT: Right, right, yeah.

Steve: So what pisses me off here is that Microsoft is essentially rendering those functioning servers - yes, they're old. Yes, they may be buggy. Yes, they have security vulnerabilities. They're not vulnerabilities to Microsoft. They're vulnerabilities that affect the users of those servers. And it's up to them if they want to update, not up to Microsoft to force them to do so.

ANT: Yeah, that is a bit of an extortion. I give you that. I give you that.

Steve: Okay. We already did this last week, and my blood pressure rises every time I start talking about this. So on to a happier subject.

ANT: Yeah.

Steve: It was Episode 887, which we recorded on September 6th of last year. And before I went back to look it up, I assumed it was longer ago because that means that it's only been seven months since I shared my discovery of this new-to-me author. And in that time I've read his first two series totaling 24 books. Our listeners will know that I'm referring to The Silver Ships series by Scott Jucha, J-U-C-H-A.

ANT: J-U-C-H-A, mm-hmm.

Steve: Now, Leo was not a fan of the series since he felt that the main character, Alex Racine, was unrealistically portrayed as being too perfect. Okay, I understand that. Not everyone is going to like everything. Not everyone likes the same sort of music, nor the same food. But for what it's worth, I had a wonderful time reading the series; and I have heard from many of our listeners who felt the same way I did. And I mean, rave reviews. There was some terrific science fiction in there. Scott is a great storyteller.

Now, when I look back over the 24 books, my main complaint would be that 20 books spent with the same characters is a lot of time with them. I knew that was the case because, when I switched - I don't remember now, like maybe it was after Book 16 of the first 20 - I switched to the four-book Pyreans sideline series, I was a bit relieved to

be meeting some completely new characters in an entirely different set of worlds. And inevitably, any 20-book storyline is going to have some spots where you're wading through detail that just seems to be taking up time.

I'm mentioning the series again because I did love it, and I'm glad to have recommended it to everyone here. If it wasn't your cup of tea, then no harm done. And if it was, then you already know how much fun was contained within those pages. Since the beginning of the year, the amount of time I've allowed myself for recreational reading has been significantly curtailed, so I must have really been reading a lot more during the first months of the series because today all I really want to do is to get SpinRite v6.1 finished and published. So pretty much whenever I'm awake, except for the time spent assembling this podcast each week, SpinRite has my full attention. But I did manage to finish the final 20th book in the Silver Ships series.

Since I really think that Scott is extremely pleasant to read - it's just comfortable reading, he writes well, he's a great storyteller, I like the way he develops characters - I have opened the first of his next eight-book Gate Ghosts series. And so far I love it, too. New people, new worlds, but the same very acceptable writing style and terrific storytelling.

Anyway, I should also note that, while I've been doing this reading of this new author, another favorite of ours, Ryk Brown, has been cranking away on the third of his 15-book arcs. Remember he's laid out five arcs, each of 15 books. And many of us have all read the first two. Many of you may already be into the third arc of 15. I read a couple, and then I was turned onto the Silver Ships series while I was waiting for the next book of Ryk's to happen. I switched, and then I got sucked into that. So anyway, he's continuing to tell the story of Nathan, Jessica, Cameron, Telles, Josh, Loki, and the rest. Anyone who's dipped their toe into the Frontiers Saga knows all of those names quite well. So I'm sure that once I eventually finish Jucha's eight-book Gate Ghosts series, I'll switch back to catch up with the ongoing Frontiers Saga, which is underway and continuing.

And speaking of SpinRite, very briefly, I'll just say that we're getting very, very close. SpinRite has not misbehaved in a long time, since I adopted full protection from any of the many ways that BIOS firmware can misbehave and was misbehaving. And as I mentioned before, there were many. But I think that, if anything now, SpinRite is probably over-insulated. But you can never have too much insulation. The few problems remaining mostly surround how hard SpinRite - I guess they're not really problems, there are questions remaining - around how hard SpinRite should try to work with badly damaged drives. When a drive is really very badly damaged, it's even difficult for SpinRite to verify that it is a drive, or that it's established communication with the drive.

So that's where we are now. We're sorting through a few issues surrounding that. On one hand, it's mostly of academic interest since none of those clearly dead and dying drives that SpinRite's testers have would ever be considered useful any longer for data storage. But we're all curious, right, to know, like, what's going on exactly. And for me, I want to make SpinRite as good as it can be, and I don't ever want to return to work on SpinRite once I declare that it's finished.

So I'm still willing to give it a little more time. But, you know, we're talking a week or two, not much more than that because, I mean, it has actually for quite a while been working just great for I think we're up to 668 people that have been running the alpha releases through the wringer. And there are a few people with clearly dead drives that, you know, SpinRite still - it's still fighting itself about whether it should wait a little bit longer for the drive to come back with a response or not. Or somebody, like, when it's chugging along and gets stuck on a problem, how long should it - how hard should it work on that before it finally gives up. So it's that sort of...

ANT: It's trying to figure out the difference between a malfunctioning drive versus a broken drive is what you're saying; right?

Steve: Yes, yes, yes. Okay. Let's talk about our last sponsor, and then Zombie Software.

ANT: All right. So we're going to talk about "The Last of Us." No, not "The Last of Us," zombies.

Steve: Yeah. So we know that a few years ago CISA, after its formation, created their KEV, which is the, you know, KEV, Known Exploited Vulnerabilities list or database. The idea was for this list to serve as a prioritization for any new entity exposed to the Internet. So this wasn't like every vulnerability ever known. The idea was that only vulnerabilities whose active exploitation has recently been observed in the wild would make it onto the list.

And as a prioritization mechanism, when you have 20 things, okay, that would be feasible. Unfortunately, because so many vulnerabilities are under exploitation, and CISA committed to adding anything that they see being exploited to the list, the list has grown. We talked about how the size of CISA's KEV ballooned last year, not because of many new vulnerabilities discovered with patches being made available in 2022, but rather because older vulnerabilities, in some cases almost ancient, were still seen in use and were therefore, as is KEV's charter, added to the list.

Okay. Now, today, we have another shoe dropping, with a rather breathtaking report from a security firm known as, I guess it's Rezilion, R-E-Z-I-L-I-O-N, Rezilion?

ANT: It could just be Rezilion, but a fancy way of spelling it.

Steve: Yeah, I mean, I would have two L's if it was Rezilion. You know?

ANT: Okay, yeah.

Steve: I guess. So anyway, I like the way Rezilion sounds, so we'll go with it as if it was million, but it's Rezilion. Anyway, so Rezilion writes: "Do you know KEV? You should," as if it's a person, I guess. Do you know Kev? "You should, because hackers do. Rezilion's research team just released a new report which highlights the critical importance of Known Exploited Vulnerabilities (KEV). Specifically, our research," they write, "uncovers that although the KEV catalog vulnerabilities are frequent targets of APT (Advanced Persistent Threat) groups, many organizations are still exposed and at risk from these vulnerabilities because they're not patching them. This gap in patching may be due to a lack of awareness, or a lack of patching resources, or both." Or maybe priorities, who knows.

They said: "The KEV catalog, maintained by the Cybersecurity and Infrastructure Security Agency (CISA), is a reliable source of information on vulnerabilities that have been exploited in the past or are currently under active exploitation by attackers. According to our new research, there are over 15 million vulnerable instances in the KEV catalog, with the majority being vulnerable Microsoft Windows instances. That's a massive number of systems exposed to attacks, leaving organizations vulnerable to exploitation from threat actors and Advanced Persistent Threat groups."

Okay. So Rezilion scanned the Internet, checking specifically for systems exposing vulnerabilities to any of the now 896 individual vulnerabilities currently listed in the CISA KEV database. What they found is what any bad guys who took the time to do the same could also find, which was more than 15 million systems worldwide currently in need of patching to prevent their easy exploitation. Rezilion's report listed the top 10 most frequently exposed vulnerabilities in a table which I've included in the show notes. I was

curious to know more about the specifics of a few of these top 10 still most prevalent exposed vulnerabilities.

Okay. So the list is sorted by the number of IP addresses expressing that vulnerability, where that vulnerability is now, today, present. So as I said, I was curious to know more about the specifics of a few of these top 10 still most prevalent exposed vulnerabilities. Number one of the top 10, which is currently present in 6,453,785 IP addresses, is CVE-2021-40438. So that one is only two years old.

ANT: Wow. And a third of the cases.

Steve: Yes, yes, is one vulnerability. And listen to what Rapid7 wrote about this back on November 30th of 2021: Rapid7 said: "On September 16th, 2021, Apache released version 2.4.49 of HTTP Server, which included a fix for CVE-2021-40438, a critical server-side request forgery (SSRF) vulnerability affecting Apache HTTP Server 2.4.48 and earlier versions," like all the way back. "The vulnerability resides in mod_proxy and allows remote, unauthenticated attackers," meaning anyone on the Internet, "to force vulnerable HTTP servers," all six-plus million of them, "to forward requests to arbitrary servers, giving them the ability to obtain or tamper with resources that would potentially otherwise be unavailable to them."

They wrote: "Since other vendors bundle HTTP Server in their products, we expect to see a continued trickle of downstream advisories as third-party software producers" - here we come to the software bill of materials we were talking about before, right, like what stuff does your product have in it - "as third-party software producers update their dependencies. Cisco," they said, "for example, has more than 20 products they are investigating as potentially affected by CVE-2021-40438, including a number of network infrastructure solutions and security boundary devices. To be exploitable, CVE-2021-40438 requires that mod_proxy be enabled. It carries a CVSS score of 9.0."

Finally, they wrote: "Several sources have confirmed that they have seen exploit attempts of CVE-2021-40438 in the wild. As of November 30th, 2021, there is no evidence yet of widespread attacks. But given HTTP's prevalence and typical exposure levels, and the fact that it's commonly bundled across a wide ecosystem of products, it's likely exploitation will continue and potentially increase."

Okay. Now, this particular vulnerability would not be leveraged for any sort of widespread attack. Maybe you could use it in an HTTP query reflection, but that's not clear due to HTTPS certificates being needed. But it is exactly what a high-level determined and targeted attacker might be looking for. Imagine a publicly exposed Apache web server on an enterprise's network boundary, which is where they are. The web server sits on the enterprise's network, as it pretty much has to, and it fields public HTTP requests over port 443 incoming from the public Internet.

And that enterprise also maintains other servers, naturally, of various types for strictly internal use. And those servers have no public exposure. They're simply sitting on the internal network. However, they are visible, because they're on the internal network, to the vulnerable Apache web server, since as I said they reside on the same internal corporate network.

ANT: Right.

Steve: This vulnerability presents the perfect means for allowing a remotely located attacker to bounce queries off of the enterprise's public Apache web server and into the private corporate network behind it. It turns the Apache server into an unregulated and unrestricted public proxy with access to the enterprise's internal private network. If and when used in that way, it is a horrifying vulnerability. And right now, of the 896 known

vulnerabilities on CISA's KEV list, it holds the number one slot by being present at 6,453,785 known IP addresses spread around the world.

The presence of this vulnerability came to light only about 18 months ago. And one of the problems, as noted by Rapid7, is that Apache is embedded in many networked appliances. They noted that Cisco has at least 20 different devices that were, and probably to a distressing degree still are, subject to this vulnerability. No fault of Cisco's. I'm sure they created patches for all of them. But as we know, having a patch and having the patch applied are two very different things. So it's not just instances where someone is running a standalone Apache web server on a Unix or Linux box. Those would be obvious and easily updated.

No. The real challenge is all of the places where Apache has been embedded inside an appliance that appears to be humming along just fine. Hey, Cisco has some updates for our network thingamajig? Well, okay, but it's working fine just now. We'll get back to it when we have some time. Then suddenly your entire network is encrypted, the ransom note arrives, and you really don't have any time. You wonder after the fact how they got in. Well, maybe they pivoted off of a known but unpatched vulnerability, six-million-plus instances of which currently exist and were recently enumerated on the public Internet.

Okay, now, I'm not going to go through all 10 of these, but let's look closely at one more that's halfway down the list, occupying slot number five. That's CVE-2015-1635. The first thing we notice is the 2015, okay, so eight years ago. And it's in slot number five of how many servers are currently vulnerable.

ANT: How?

Steve: I know.

ANT: Eight years old?

Steve: And the other thing you notice is the 1635, you know, a reminder of how quaint things were eight years ago when CVE numbers only had four digits. Now we need five. I hope we never need six.

ANT: It was a simpler time, sir.

Steve: It was a simpler time. So how many known vulnerable instances of that now eight-year-old problem are still presently exposed on the public Internet? Well, 120,156. And get this. It's one of those quite rare vulnerabilities that's managed to earn itself a CVSS score of 10.0. That's right. It doesn't get any worse or more frightening for those who rate these things. We have seen how difficult it is to score a perfect 10. There are plenty of wannabe 9.8's out there, but the 10.0 remains a rare beast. And who managed to bring that one home? None other than Microsoft in their own web server, where its successful exploitation, simply by making the proper remote query, allows the attacker's code to run in the user's system with full system kernel level privilege, thus earning the full monty 10.0 score. Again, a full eight years downstream, and yet today 120,156 Microsoft web servers remain sitting ducks for this remote code execution vulnerability. I know. Okay. Now, I'm a software developer.

ANT: Why? I mean, why? This is known information. Is it just lack of resources? And when I say "resources" that means time and money where none of this stuff is getting patched?

Steve: You know, the thing that's missing from these scans, because there's really no way to get them, we can get demographics by country, and sometimes that is

illuminating. But we don't get demographics by size of company or number of employees or corporate annual revenue. And it would be really interesting to see what the correlations were there. And to answer your question, Ant, it's like, okay, what explains this? Because at this point all we know is that that's the number.

Okay. So I'm a software developer who's done a fair share of Internet programming. SQRL is an Internet authentication system. GRC's ShieldsUP and the DNS Spoofability services are persistently available on the Internet, and GRC's DNS Benchmark, now with 8.2 million downloads, which are now occurring at a rate of around 2,000 new downloads a day, has become the industry standard for measuring DNS server performance. I published that benchmark 13 years ago, back in 2010, and it's never had a bug. So I deeply and fundamentally object to what is clearly a growing presumption in the world that any software that's not being actively maintained is therefore inherently bug ridden.

And I see this. I mean, it's in the air. It's as if no one trusts software anymore that isn't in intensive care with multiple IV bags hanging overhead. It has to be in serious trouble, on continuous life support, being monitored 24 hours a day, for anyone to think that it doesn't have serious problems. It's perverse. And yet the most frustrating thing is I can't argue against that because it is the only sane and rational conclusion to draw from all the evidence that is presented to us every day.

ANT: I mean, we hear from the likes of Microsoft, as big as they are, there's always some type of announcement about a zero-day or some other bug that scares the crap out of everybody. So why would the public folks just trust that, oh, this software doesn't have any bugs? I mean, why would we? It's because there's always something in the news.

Steve: Now, there are, as it turns out, there are some beautiful exceptions to that. SpinRite 7 will be based upon an embedded 32-bit real-time operating system, all ongoing support for which ended at the end of last year when its publisher went out of business; and yet I've selected it for SpinRite's future. Why? They didn't go out of business because the software was no good. They went out of business because the software was too good. It was done. It was a perfectly functioning finished product, and no one was subscribing to updates for it anymore because it was feature complete, and there were no more bugs to be fixed. None. It was done. Take it off the ventilator. Pull the IV lines. Hold your breath and count to 10. Does it still have a pulse? What do you know. It's alive. Works great.

Now, the idea that a company went out of business because its software didn't have any bugs is a bit chilling because nobody wants to use Windows after its life support has been terminated. If it's still moving after that, the presumption is that it's become a zombie that wants to eat your brain. So when we hear that Microsoft shipped a brand new edition of Windows - I don't recall which Windows it was. I remember hearing it, maybe Vista, 7, 8, or 10, which at the time of its release had more than 10,000 bugs known to Microsoft, you have to wonder...

ANT: That sounds like a Vista thing 'cause Vista was so bad.

Steve: Yeah. You have to wonder whether the leaking of what should have been embarrassing information might not have been deliberate on their part. If a company went out of business because their software worked perfectly and had no more bugs, that's not a worry that will ever keep anyone at Microsoft awake at night.

Everyone wants to use Windows, and everyone wants to make sure that Windows IV lines are continuously connected and that vital, life-sustaining fluids are flowing freely into it. And the minute Microsoft cuts off that flow, whoa, time to upgrade. And not because you necessarily want anything that the new Windows has. Quite often, in fact, much of what they've done to it is noxious and unwanted. Just look at Vista or Windows

8. But what choice do you have? Everyone knows that Windows is perpetually so sick that the only safe way to use it is with it lying in ICU, attached to a ventilator, with IV lines flowing and under constant monitoring. You disconnect it from its continuous life support at your peril.

And Microsoft with Windows, or Exchange Server, or, well, I guess anything Microsoft creates, is only a slice of a much larger problem that this industry has. This is the sad reality of today's software ecosystem. Our insatiable desire for features, or in Microsoft's case their unrelenting need to seduce us with new bug-ridden features, means that we're never going to get ahead of this; and that, indeed, almost all software only remains viable while it's on perpetual life support.

So now we loop back to CISA's KEV list, and we see that what we have is a list of what can only be described as Zombie Software. For one reason or another, that software has been disconnected from the life support that was, sadly and unnecessarily, but still in fact vital for its ongoing safe use on the Internet. In the case of that eight-year-old Windows web server critical 10.0 vulnerability, that Windows 7 and Server 2008 era software is now more than three years past its end of support life, which Microsoft terminated in January of 2020.

It's worth noting, however, that the updates for those 120,156 very old servers are still online and are available from Microsoft. So simply clicking "check for updates" on any of those machines, assuming that they're activated and validly licensed, will quickly bring them current with the last build of their server's code and would, in the process, cure any and all vulnerabilities that were known and for which patches were available when support for that server ended three years ago.

ANT: Imagine that.

Steve: Yeah. But that's one of the two oldest entries in the list. The others, the newer eight are much more recent, and ongoing life support still remains available for all of those. A perfect example of that is the other critical vulnerability we examined, Apache. Any of those 18-month-old Apache servers could be updated immediately. Yet 6,453,785 publicly exposed instances of Apache voluntarily remain unpatched today, despite the fact that updates are certainly readily available for every instance of them.

So where does this all leave us? One thing I want to observe is that zombies eat the brains of those who are nearby. So this is about self-responsibility. If you don't want your zombie server to eat your brain, either shut it down or get it patched, if possible. This is one, as I said, of my main complaints with Microsoft's plan to force other people to update their Zombie Exchange Servers. Microsoft contends that their brain was going to be eaten by somebody else's remotely located zombie. That's utter nonsense. We all understand that we're the ones whose brains are in danger if we choose to use software after it's gone zombie.

But taking the broadest view possible, it is an unfortunate and sad truth that, with very few exceptions, almost none of today's software is able to stand alone without the presumption that periodic and continuing updates are required for that software's users to feel comfortable with the performance of their software. The assumption of the availability of ever-present Internet, which provides a lifeline for would-be zombies, allowing them to remain on life support, has enabled a gradually growing laxity on the part of software publishers. Today, few apparently feel the need to get it right the first time, since problems can always be fixed later.

In the extreme case of Microsoft, major products are shipped, despite reportedly being riddled with tens of thousands of bugs known to its developers at the time that it is shipped. And as we saw with last week's Exchange Server issue, Microsoft's own self-

fulfilling prophecy is that, once any of their software is removed from life support, not only everyone nearby, but also everyone else on the Earth is put in danger once it becomes a zombie. It's a sad state of affairs which did not have to be. But as we sometimes observe on this podcast, it's the world we live in. The best we can do is to navigate it safely.

ANT: Now, Mr. Steve, I'm going to come to Microsoft's defense for half a second; okay?

Steve: Okay.

ANT: Google is just as bad on this front. There's been a gazillion times Google will push out a version of Android that's gone through alpha, it's gone through beta testing and so forth. And the second it gets on your device, which happens to be their flagship device, it is a craptastic experience until they run that first patch.

Steve: So actually you're supporting my position.

ANT: Oh, I know, I know.

Steve: You're adding another example of a company that is relying on updates rather than ever getting it right.

ANT: Yeah. I just didn't want it to come out as if you're just only beating up on Microsoft because there's a bunch of them out there.

Steve: Oh, yeah. Yeah. I'm just using them as a good example, and because they had number five on the list with a server that 120-plus thousand of them are right now today vulnerable to a remote code execution exploit which exists and, if anyone cared to, could take them over because you know, bugs were not fixed.

ANT: Unbelievable.

Steve: Yup. It's the world we live in.

ANT: So Mr. Steve, this has been another, 'nother fun chat and discussion. And yes, there was some sad stuff that happened in this, some scary stuff. But you did give us a little bit of good news on this week's episode of Security Now!. And I thank you for all of the information and tidbits that you provided with us this week, sir.

Steve: My pleasure, Ant. And we'll be back for more next week.