
# Artificial General Intelligence (AGI)

**Description:** How Microsoft lured the U.S. government into a far deeper and expensive dependency upon its cybersecurity solutions. Gmail to offer native throwaway email aliases like Apple and Mozilla. Russia to ban several additional hosting companies and give its big Internet disconnect switch another test. Russia uses a diabolical Windows flaw to attack Ukrainians. The value of old Security Now! episodes. TrueCrypt's successor. Using Cloudflare's Tunnel service for remote network access. How to make a local server appear to be on a remote public IP. How to share an "impossible to type" password with someone. How to find obscure previous references in the Security Now! podcast. What are the parameters for the expected and widely anticipated next generation Artificial General Intelligence (AGI)? What do those in the industry and academia expect? And is OpenAI's Sam Altman completely nuts for predicting it next year? Is it just a stock ploy?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-1001.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-1001-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. He says there's not a lot of news, so we're going to do a lot of questions from the audience, feedback and so forth. And then Steve will explain in his understanding of what is going on with AI, the search for artificial general intelligence, and how close we are coming. I think you're going to like this episode. Security Now! is next.

**Leo Laporte:** This is Security Now! with Steve Gibson, Episode 1001, recorded Tuesday, November 19th, 2024: Artificial General Intelligence.

It's time for Security Now!, the show where we cover your security, privacy, safety, how computers work, what's so intelligent about artificial intelligence, all that jazz, with the most intelligent guy I know, this cat right here, Mr. Steve Gibson.

**Steve Gibson:** I am not that, Leo.

**Leo:** You are not that?

**Steve:** No. I'm what we call a domain expert.

**Leo:** Ah, yes.

**Steve:** I have some expertise in a couple places.

**Leo:** When it comes to Sudoku you're just like the rest of us.

**Steve:** And when it comes to artificial intelligence I'm claiming no expertise. I wanted to talk about, as I said last week, artificial general intelligence (AGI), because everyone's throwing the term around. We're hearing people talking about it. What caught my attention was when Sam Altman, the infamous and famous CEO of OpenAI, he claimed, oh, yeah, we'll have that next year.

**Leo:** Any day now.

**Steve:** He said 2025.

**Leo:** Yeah.

**Steve:** And it's like, what?

**Leo:** But he's kind of a salesman.

**Steve:** Oh, well, yes. So maybe this was just a nice little stock price boosting ploy. But I wanted to take some time. I found a couple interesting articles with a lot of other people in the industry interviewed and some academics interviewed. And I thought, let's, you know - so today is like, no one's going to find out some great revelation about AGI because I don't have it. But, you know, it's clearly a thing. And I just thought we should kind of put a marker down and say, okay, here's where it is.

**Leo:** You've done it before. You did it with blockchain. It's very frequent that you're able to, because that's how you work, digest all this stuff. You're kind of our retrieval augmented generation. You digest all this stuff and give it back to us so we can understand it. So I'm very much looking forward to this episode.

**Steve:** Well, in the fullness of time, if I spend some time, you know, digging in...

**Leo:** You're good at it, yeah.

**Steve:** ...then that would be interesting. But we've got a bunch of stuff to talk about. We're going to look at - oh, this is a great story - how Microsoft lured the U.S. government into a far deeper and expensive dependency upon its own proprietary cybersecurity solutions than the Biden administration expected. Also Gmail will be offering native throwaway email aliases, much like Apple and Mozilla.

**Leo:** Oh, good.

**Steve:** We'll touch on that. Oh, my god, and Russia, well, they're banning additional hosting companies. They're going to give their big Internet cutoff switch another trial next month, and some other things that we'll talk about. Oh, and they used a diabolical Windows flaw to attack Ukrainians. It was found by a security group. And, boy, when our old-timers find out that something we assumed was safe might not be safe to do, that's going to raise some hair.

Also, we're going to look at, oh, I have a note from our listener about the value of old Security Now! episodes. We're going to touch on TrueCrypt's successor. Also using Cloudflare's tunnel service for remote network access. Another of our listeners said, hey, this is what I'm doing. So we're going to share that. Also answer the question about how to make a local server appear to be on a remote public IP, which in this case is coming in handy for pretending to be a remote command-and-control server when testing malware.

Also, how to share an impossible-to-type password with someone else. Oh, and another listener asked, and I answered, and then he confirmed about finding obscure previous references in the Security Now! podcasts. So that, and then we're going to dig into this whole question of what is artificial general intelligence and how is what we have today failing that. What are the recognized and widely agreed-upon characteristics that AGI has to have, and when might we get some?

So I think a great podcast. There was not, as you could tell, there was not a huge amount of news. I looked everywhere for good stuff. But, boy, I added it up. I think I have 4,300 plus some inbound pieces of email from our listeners.

**Leo:** Holy cow.

**Steve:** So like since this began. So I'm not starving at all for listener feedback. And, you know, I think it's fun. Actually we've got - changing this from Twitter to email completely changed the feel of the feedback.

**Leo:** Good.

**Steve:** Since it no longer needs to fit into 280 characters.

**Leo:** Yes, I'm not surprised, yeah.

**Steve:** You know, and so it's a lot more interesting.

**Leo:** Excellent, excellent.

**Steve:** So a great podcast. Oh, and Leo, we're starting in on our second thousand.

**Leo:** Oh.

**Steve:** This is Podcast #1001.

**Leo:** I hadn't really thought of it quite that way, but...

**Steve:** The second thousand. That's right.

**Leo:** You put that into perspective, I guess we are.

**Steve:** It's what everybody wants. They want another thousand. It's like, okay.

**Leo:** Oh, god.

**Steve:** Here we go.

**Leo:** Okay. Well, you and I are going to work on it. We're going to do our best. That's all we can promise, just our best.

**Steve:** I look different than I did 20 years ago, but you look about the same.

**Leo:** You're being very kind.

**Steve:** You've got your hair still. It's a nice silver.

**Leo:** Haven't lost the badger. I still have the badger on top. Steve, I'm ready with the Picture of the Week. It's a good one this week.

**Steve:** It is a good one. And I've had some feedback from our listeners already who really liked it. I was, again, on the ball. And just a reminder to our listeners that those - we had just shy of 13,000 people who now subscribe to the Security Now! mailing list.

**Leo:** That's so great.

**Steve:** 12,979.

**Leo:** That's almost exactly the same number of Club TWiT members we have. So I think there's maybe a correlation there.

**Steve:** Yeah, I think there may be. And there was - that was the count when the mailing went out around 3:00 p.m. yesterday. So just saying that 24 hours ahead of time, anybody who subscribed to the list got this stuff. So, okay. Anyway, so the point was that many people wrote back and said, wow, that's terrific.

So what we have is a residential staircase going up, you know, as they do along one wall with a handrail, and then a banister on the outside so that the stairs are not open. Now,

this family has a couple of toddlers. And looks like maybe sister's a little older than brother.

**Leo:** She was first up.

**Steve:** He's in diapers still and looks like maybe he's two. She might be maybe two and a half or three. I don't know. But across the bottom of the stairs is a screen. Mom and Dad have said kids are not going upstairs. They stay downstairs.

**Leo:** It's a child gate. And I think it's a brand new one. It looks like it because it's still got the sales tag on it.

**Steve:** You're right. And I noticed also that behind it are a couple of stacks of stuff that, you know...

**Leo:** Yeah, they don't want the kids to get into.

**Steve:** They don't want the kids to get into, exactly. Well, now, I gave this picture the caption "The bottom of this staircase may have been blocked, but these future hackers are not deterred." Because the stairs protrude out from the banister supports, and both of the kids have walked up the outside of the stairs, like seeing whether there's a way they can get in there because they're going to find a way. And it looks like maybe, if I'm right, the oldest sibling looks like she's sort of trying to squeeze herself in because she sort of ran out of runway there.

**Leo:** We got to the top of that [crosstalk]. Now how do we get in?

**Steve:** She's going to have to - so, yeah. So there are - we hope the analogy is not that they're behind bars because, you know, the banister does look a little bit like that, too. But, you know, these guys, they're determined to find a way past Mom and Dad's blockade of the stairs, so future hackers.

**Leo:** Oh, boy. That's pretty accurate, yeah.

**Steve:** Future hackers. Okay. So some recent reporting by ProPublica raised some interesting questions. And I got a kick out of this. I'm sure that our listeners will, too. So ProPublica, and I'll be interrupting a few times here with some of my own comments, they said: "In the summer of 2021" - and we covered this at the time - "President Joe Biden summoned the CEOs of the nation's biggest tech companies to the White House. A series of cyberattacks linked to Russia, China, and Iran had left the government reeling." And of course some of that was Microsoft's fault; right? "And the administration had asked the heads of Microsoft, Amazon, Apple, Google, and others to offer concrete commitments to help the U.S. bolster its defenses. Biden told the executives gathered in the East Room: 'You have the power, the capacity, and the responsibility, I believe,'" he said, "'to raise the bar on cybersecurity.'"

Now they said: "Now, Microsoft had more to prove than most. Its own security lapses had contributed to some of the incursions that had prompted the summit in the first place, such as the SolarWinds attack, in which Russian state-sponsored hackers stole sensitive data from federal agencies, including the National Nuclear Security Administration. Following the discovery of that breach, some members of Congress said the company should provide better cybersecurity for its customers. Others went even further. Senator Ron Wyden, who chairs the Senate's finance committee, called on the government to 'reevaluate its dependence on Microsoft' before awarding it any more contracts."

Now, as we're going to see shortly, what happened is not exactly what Ron as looking for. This was not the kind of reevaluation that Ron had in mind.

ProPublica said: "In response to the President's call for help, Microsoft's CEO Satya Nadella pledged to give the government $150 million in technical services to help upgrade its digital security." Well, isn't that nice. "On the surface," they wrote, "it seemed a political win for the Biden administration and an instance of routine damage control from the world's largest software company. But the result of ProPublica's subsequent investigation suggests that Microsoft's seemingly straightforward commitment to provide a bunch of free technical services belied a more complex, profit-driven agenda.

"As time has since revealed, Microsoft's apparent generosity was a calculated business maneuver designed to bring in billions of dollars in ongoing revenue, lock competitors out of lucrative government contracts, and even further tighten the company's grip on federal business." And as I'm reading this, I thought, you know, if I didn't know better, I would think Gates was still around since this turned out to be a recognizably classic Bill move.

So they wrote: "The White House Offer, as it was known inside Microsoft, would dispatch Microsoft consultants across the federal government to install Microsoft's cybersecurity products, which as part of the offer were provided free of charge for a limited time." That's right. What a bargain. What's wrong with this picture?

Okay. So they said: "Well, how about once the consultants installed the upgrades, federal customers would be effectively locked in because shifting to a competitor after the free trial would be cumbersome and costly, according to former Microsoft employees involved in the effort, most of whom spoke on the condition of anonymity because they feared professional repercussions. At that point, the customer would have little choice but to pay for the higher subscription fees.

"In fact, two former sales leaders involved in the effort likened it to a drug dealer hooking a user with free samples. 'If we give you the crack, and you take the crack, you'll enjoy the crack,' one said. 'And then when it comes time for us to take the crack away, your end users will say, "Don't take it away from me." And you'll be forced to pay.'

"Former salespeople said that Microsoft wanted more than those subscription fees. The White House Offer would lead customers to buy other Microsoft products that ran on Azure, the company's, of course, their cloud platform. This carried additional charges based on how much storage space and computing power the customer used. These former salespeople said that the expectation was that the upgrades would ultimately 'spin the meter,' quoting them, 'spin the meter' for Azure, helping Microsoft take market share from its main cloud rival, Amazon Web Services.

"In the years after Nadella made his commitment to Biden, Microsoft's goals became reality. The Department of Defense, which had resisted the upgrades for years due to their steep cost, began paying for them once the free trial ended, laying the groundwork

for future Azure consumption. So did many other civilian agencies. Former Microsoft salesperson Karan Sondhi, who had knowledge of the deals, said that 'The White House Offer' got the government hooked on Azure, 'and it was successful beyond what any of us could have imagined.'

"While Microsoft's gambit paid off handsomely for the company, legal experts told ProPublica the White House Offer should have never come to pass, as they sidestep or even possibly violate federal laws that regulate government procurement. Such laws generally bar gifts from contractors and require open competition for federal business.

"Eve Lyon, an attorney who worked for four decades as a procurement specialist in the federal government, said that accepting free product upgrades and consulting services collectively worth hundreds of millions of dollars is not like a free sample at Costco, where I can take a sample, say 'Thanks for the snack,' and go on my merry way. Here, you have changed the IT culture, and it would cost a lot of money to switch to another system."

Microsoft, for its part, defended, of course, its conduct. Steve Faehl, that's F-A-E-H-L...

**Leo:** Good name, yeah.

**Steve:** Yeah, I thought I should spell it, F-A-E-H-L. Steve Faehl, the security leader for Microsoft's federal business, said in a statement: "The company's sole goal during this period was to support an urgent request by the Administration to enhance the security posture of federal agencies who were continuously being targeted by sophisticated nation-state threat actors. There was no guarantee that agencies would purchase these licenses, and they were free to engage with other vendors to support their future security needs."

"Pricing for Microsoft's security suite was transparent," he said, "and the company worked 'closely with the Administration to ensure any service and support agreements were pursued ethically and in full compliance with federal laws and regulations.' Faehl said in the statement that Microsoft asked the White House to 'review the detail for antitrust concerns and ensure everything was proper, and they did so.'"

**Leo:** I love the phrase "hooked on Azure."

**Steve:** Hooked on Azure.

**Leo:** That's a nice ad campaign right there.

**Steve:** There's only one little problem with this, of course. As we know, it really is surprisingly difficult to switch vendors. And of course it gets worse. ProPublica found: "The White House summit ushered in a new form of concentrated reliance, as well as the kind of anticompetitive behavior the Biden administration has pledged to stamp out. Former Microsoft salespeople told ProPublica that during their White House Offer push, they advised federal departments to save" - get this, Leo - "to save money by dropping cybersecurity products they had purchased from competitors. Those products," they told them, "were now 'redundant.' Salespeople also fended off new competitors by explaining to federal customers that most of the cybersecurity tools they needed were included in the free upgrade bundle.

"Today, as a result of the deals, vast swaths of the federal government, including all of the military services in the Defense Department, are more reliant than ever on a single company to meet their IT needs. ProPublica's investigation, supported by interviews with eight former Microsoft employees who were involved in the White House Offer, reveals for the first time how this sweeping transformation came to be a change that critics say leaves Washington vulnerable, the very opposite of what Biden had set out to achieve with his summit." Because of the monoculture; right? It's like, oh, everybody's using Microsoft. Unfortunately, we've seen Microsoft making some significant mistakes.

**Leo:** Well, wasn't this in kind of response to SolarWinds?

**Steve:** Yes.

**Leo:** Yeah.

**Steve:** Yes, this was three years ago when it was like, oh my god, what are we going to do? And so Microsoft said, hey, how would you like some free stuff?

**Leo:** It was free for the first year.

**Steve:** There's $150 million of stuff for free.

**Leo:** It was only free for the first year. I mean, it wasn't even free for - it was a trial offer, basically.

**Steve:** It was, I mean, okay. So the ProPublica article. I've got a link in the show notes. It goes into much greater detail. That was just like the introduction quarter of it. So I have a link to it, as I said, for anyone who wants more. But I'm sure that all of our listeners get the idea. At one point, Microsoft was asked to provide this enhanced security support to the federal government at no charge indefinitely, which they flatly declined. Then of course it became a negotiation over well, then, how long would the services be free?

And of course what adds even more salt to this wound is that for many years these same federal and military agencies had been steadfastly refusing to go with Microsoft solutions due to their cost. But they could not say "no" to "free." So this allowed Microsoft to get their solutions in the door, to remove any previous "reasonably priced" competitive solutions. And then, once the free offer expired, the choice was either pay up or go without. It's at least mildly disgusting. And what's more, you know, this didn't just fall into Microsoft's lap; right? Former insiders made it clear that this was their intention all along, from the beginning. Microsoft's CEO Satya Nadella knew exactly what he was doing. Basically it was a Trojan horse.

**Leo:** How hard is it, if you've upgraded your security to Microsoft G5 level, is it to go back? Like if they go, oh, we don't want to pay for it, so we're going to go backwards.

**Steve:** If Elon Musk is going to do anything...

**Leo:** This is something he might want to weigh in on, yeah.

**Steve:** This is the kind of thing, I mean, it takes holding your breath and pinching your nose. And, I mean, it's an upheaval. And so anyone in IT understands that. But it's not their money they're spending, it's our money they're spending. And so it's always less expensive to pay for the incremental cost of another, you know, another three months than it is to say, okay, we're on the wrong path. We're going to just - we're going to dead-end this path. Because it does then mean going out and getting competitive bids, and literally having downtime while all of this changes because, you know, you have to remove all of this junk and put in new stuff.

**Leo:** Plus if the whole motivation for doing this was, oh my god, we've got a big security problem, you're not going to tear out the security fix you just installed to fix that so that you can do something else. You're going to be in a lot of pressure just to keep on keeping on.

**Steve:** Well, and Leo, you and I and the old-timers who are listening to the podcast, we all remember Gates. I mean, Bill...

**Leo:** Oh, yeah, he would...

**Steve:** Bill was much, you know, he's reversed as some technical genius. I mean, he's a genius. But he was much more of a businessman...

**Leo:** Oh, yeah. He was a shark.

**Steve:** ...than he was a coder. And he says that now, too. You know, I mean, so we watched all of the early shenanigans that Microsoft got up to, you know, things like, oh, you can't remove our browser. We built it into Windows. No, it's part of the operating system. What?

**Leo:** Right, right.

**Steve:** No, it's not. Until the EU said take it out, and they said, well, okay. You know. Since you're not giving us any choice.

**Leo:** In other words, same old, same old.

**Steve:** But this is just - this just struck me as so Gatesian. It was just like, oh, boy.

**Leo:** Yeah.

**Steve:** Yeah. So, ouch. Okay. So Apple has "Hide My Email." Mozilla offers their "Firefox Relay." And, you know, these are email services that create throwaway aliases for a user's primary account. The recent news is that Google is reportedly working on adding something which they call "Shielded Email" to Gmail, for their two billion Gmail users. So as with the other services, users will be able to quickly generate random-looking usernames for use, you know, filling out online forms and subscribing to things and so forth, which hide their real email addresses. So those are just aliases. And then you'll have some means of managing the aliases so that, for example, if you started to get spammed on one, first of all, it would be interesting to know who, you know, which email address is spamming you. And then you're just able to delete it, and you'll get rid of it.

So I've noticed that a large percentage of the subscribers to GRC's mailing lists are Gmail domain users. So I imagine this will come as a welcome service. Unfortunately, I use Gmail as my trashcan already because I've got, you know, GRC.com email addresses. So it's a little late for me. I don't think it would serve much purpose using, you know, shielding what is already my throwaway account. But still, for people whose main, whose primary email is Gmail, I think this sounds like a good thing. And, you know, better late than never. It certainly took them a while. On the other hand, Leo, can you imagine the infrastructure that Google must have in order to give two billion users, like, email that works as well as Gmail does?

**Leo:** And they use their own server. They're not using, you know, an open source server or anything like that. So if you were, you might be a simple plugin. But, yeah, it's a big deal. It's a lot to move.

**Steve:** Yeah.

**Leo:** Yeah. Plus it's old. Let's not forget, Gmail is not a brand new service by any means.

**Steve:** Correct.

**Leo:** It was one of the very first web services.

**Steve:** Correct. In fact, I remember - do you remember a guy named Steve Bass, who was - he was the - he ran the Pasadena IBM PC User Group.

**Leo:** Oh, yes, okay.

**Steve:** PIBMUG was the...

**Leo:** Yeah, yeah.

**Steve:** ...if you tried to pronounce the - anyway.

**Leo:** Yeah, yeah.

**Steve:** And I think he wrote for PC World also.

**Leo:** Yeah, I remember his byline, I do, yes.

**Steve:** Yeah. Neat guy. And he had early access to Gmail and so sent me...

**Leo:** An invite.

**Steve:** ...an invite that allowed me to get a special email account at Gmail.

**Leo:** Yeah, which you're not going to tell anybody because otherwise it would be completely useless.

**Steve:** Believe me, it's next to that now anyway.

**Leo:** It's useless now, yeah.

**Steve:** It's just, you know...

**Leo:** I have laporte@gmail, which was - because I was also early on.

**Steve:** Very nice, yup.

**Leo:** And everybody's decided apparently, the spam world has decided that I'm French. And I get a lot of French spam, almost exclusively French spam. And I also, because people - probably this happens to you. I'm sure it happens to our listeners. They don't really understand that you can't put a space in a Gmail address. So a lot of people named Francois Laporte and Abigail Laporte, they type a space in there, and it all goes to laporte@gmail.

**Steve:** Right.

**Leo:** So I get all sorts of stuff like your tickets are ready, I mean, just endless. Your reservations for tonight in Paris, I mean, it's - I'm tempted; but no, I'm not.

**Steve:** Well, and you're right. The problem with it being that big, like all those domains, or all those names in a single domain is that, if it is not like, you know, bzqrt79 or something, if it is Leo or Fred...

**Leo:** It's the end of the world, yeah.

**Steve:** You're just like, you know, goodbye.

**Leo:** There's a story about jim@aol.com. Poor Jim never really did get to use that email address.

**Steve:** Wow.

**Leo:** Do you want me to take a break, or do you want to continue on?

**Steve:** I think now is a good time. We're half an hour in. And then we're going to talk about it's definitely not love coming from Russia.

**Leo:** "From Russia With Love."

**Steve:** So we're going to talk about some - and we do get to talk about Roskomnadzor.

**Leo:** Roskomnadzor. Thank you, Steve.

**Steve:** So Russian officials...

**Leo:** Roskomnadzor. I'm sorry. I jumped the gun.

**Steve:** No, no, we're going to get there in a second, have recently announced via Telegram - which I thought was interesting. Oh, yeah, let's use Telegram...

**Leo:** Isn't that interesting.

**Steve:** ...while punishing them.

**Leo:** Wow.

**Steve:** ...that they plan to expand Russia's ban on foreign web hosting providers who are hosting content that discredits the "glorious Russian Army," their words. So Akamai and CDN77 may soon find themselves added to the banned list for being naughty. Overall, Russia appears to feel that the Internet is at best a mixed blessing. It's unclear to me how it's possible to even function within today's globalized economy without it. I think they're nuts. But Russia seems poised...

**Leo:** [Clearing throat] I'm sorry, I'm getting ready. I'm getting ready for the - go ahead.

**Steve:** That's right. Russia seems poised to at least explore getting along without the Internet. To which end, Russia's illustrious Internet watchdog, none other than Roskomnadzor...

**Leo:** Roskomnadzor [with echo]. I'm sorry.

**Steve:** ...has announced its plan to conduct another test next month of Russia's big Internet disconnect switch, when pulled, does what it says. It severs all ties between Russia and the rest of the global Internet.

**Leo:** Wow. They did it once before; didn't they? They tried this.

**Steve:** Yes. And they've been working on it for years. They have to do things like figure out what to do with DNS queries that resolve to IP addresses that are no longer available. I mean, they just don't want everything to hang and crash and, like, you know, with the hourglass spinning. So it turns out that disconnecting from the Internet is not an easy thing to do. And of course as I was thinking about this, I thought, what about Starlink? Because, you know, it's no longer the case that useful Internet connectivity requires landlines and fiber optic trunks and all of that. You know, Starlink is a thing.

**Leo:** I would guess Starlink is banned in Russia. That would be my guess.

**Steve:** Is it?

**Leo:** Or doesn't offer it. Let me see. It's available in Ukraine, of course.

**Steve:** And you're right, Russia is sanctioned right now.

**Leo:** Yeah, that's what I thought.

**Steve:** So, yeah.

**Leo:** So that just works in their favor; doesn't it.

**Steve:** That's right. Easier to disconnect.

**Leo:** Oh, man.

**Steve:** Easier to pull the switch. So anyway, so they're going to do another test in December. And again, you know, it's like, is there some big long-term plan here? Is it just so that they, like, are worried they're going to get attacked? I don't know. We would know if our country was doing the same thing because it would have an effect. I mean, pulling the switch on global connectivity will have an effect.

**Leo:** Yeah.

**Steve:** So, really interesting. We'll have to see what they've got planned. But while we're on the topic of Russian antics, get a load of this. One of the zero-days, it was CVE-2024-43451, that Microsoft patched this past week, you know, in Patch Tuesday last week, was used in a Russian hack of Ukrainian organizations earlier this year. According to the security firm ClearSky, the zero-day was part of an exploit chain that exposed NT LANMAN, you know, NT LAN Manager, credential hashes, also known as NTLM credential hashes, when victims interacted with .URL files that were received in phishing emails.

But here's the part that really caught my attention. ClearSky said that right-clicking, deleting, or moving the file established a connection with the attacker's server, exposing authentication data. The report suggests that the campaign also used social engineering to convince victims to run executables.

Okay, but hold on. Right-clicking on a file to display its context menu and examine its properties, deleting it or dragging it to another directory, was all that's needed to cause the victim's machine to establish a remote connection to a malicious server? What? So I went over to ClearSky to see what was up. And I've got a link in the show notes for anyone who wants to see, too.

The ClearSky Research Team posted their write-up last Wednesday, writing: "A new zero-day vulnerability" - oh, by the way, it was posted Wednesday because the patches were pushed on Tuesday, the day before, you know, closing this down. They said: "A new zero-day vulnerability, 43451..."

**Leo:** Ironically, ClearSky Security's sent an invalid response. I don't know if it's blocked or can't provide a secure connection. So it might be my browser. Sometimes this happens.

**Steve:** Interesting.

**Leo:** I think it's me, probably.

**Steve:** Maybe do an explicit HTTPS?

**Leo:** Yeah, no. Because I think the Ubiquiti blocks certain things. I don't know why.

**Steve:** Ah, okay.

**Leo:** I was just clicking the link you provided.

**Steve:** Yeah, yeah. Let me try clicking it here.

**Leo:** Yeah, I'm sure it's fine. It's just me. Yeah, I also have that from Safari.

**Steve:** Yup, it just came right up for me.

**Leo:** Yeah, so it's a - I've noticed this, there are certain places I can't go, and I think it's the security, I do use security in Ubiquiti.

**Steve:** Oh. Okay. So they wrote: "A new zero-day vulnerability, 43451, was discovered by ClearSky Cyber Security in June of this year, 2024. This vulnerability affects Windows systems and is being actively exploited in attacks against Ukrainian entities. The vulnerability activates URL files containing malicious code through seemingly innocuous actions." Then they have three bullet points.

First, a single right-click on the file in all Windows systems will do this. Deleting the file in Windows 10 or 11 will do this. Dragging the file to another folder in Windows 10 or 11 and some Windows 7, 8, and 8.1, they wrote. The malicious URL files were - and I should note that a URL file is just text. So it's kind of pushing it to call it malicious, but okay.

**Leo:** Yeah, it's just a link.

**Steve:** It's just, yeah, it's got - it looks like an INI file. So they wrote: "The malicious URL files were disguised as academic certificates and were initially observed being distributed from a compromised official Ukrainian government website." What actually happened was that the Russians compromised an email server in Ukraine and then used the email server's credentials to send, you know, DKIM, SPF, you know, DMARC-approved email to others in Ukraine. So the email that was coming in looked like it was verifiably authentic from the compromised server. But in fact, unfortunately, it was phishing email.

So they said: "The attack begins with a phishing email sent from a compromised Ukrainian government server. The email prompts the recipient to renew their academic certificate. The email contains a malicious URL file. When the user interacts with the URL file by right-clicking, deleting, or moving it, the vulnerability is triggered." So I'll just say this is like, this is the first time I've seen that, like, you know, dragging a file and dropping it in the trash or right-clicking to learn more about it, that's all it takes under Windows 10 and 11 in order to, well, and right-clicking in all versions of Windows in order for this thing to happen. Anyway, I've got more detail.

So they said: "When the user interacts with the URL file by right-clicking, deleting, or moving it, the vulnerability is triggered. This action establishes a connection with the attacker's server and downloads further malicious files, including SparkRAT malware. SparkRAT is an open-source remote access trojan that allows the attacker to gain control of the victim's system. The attackers also employed techniques to maintain persistence on the infected system, ensuring their access even after a reboot."

Okay. So the culprit here is a .URL file, which is a Windows Internet URL shortcut. It's a text file. And anyone who's ever looked at like the original .INI, you know, config files back in the early days of Windows will recognize the format here. It's got sections that are surrounded by square brackets, and then just simple name=value pairs, all in text. The key is that the file contains a URL= line where the scheme of the URL is "file://" followed by the IP of the malicious remote server.

In Windows, the file:// scheme is handled by SMB, which is of course Server Message Blocks, which underlies Windows original file and printer sharing which, as we know, was never up to snuff security-wise. So that's where NTLM credential hashes come in because

Windows has always been extremely generous handing out its, like, ID'ing its users by sending their credential hashes around, long before it was realized that, you know, that's not a good idea, to be sending somebody's hashed credentials, because there's all kinds of mischief you can get up with them, including just a replay of the credential hash in order to impersonate them. Which is exactly what this thing does.

So apparently upon even extremely innocuous contact with these files in Windows - and it's worse in more recent Windows 10 and 11 - Windows Explorer will, without any prompting, reach out to the file server that's indicated in the shortcut, even without its recipient executing the shortcut. The researchers wrote: "When examining the URL file, ClearSky's team exposed a new vulnerability. Right-clicking the file establishes a connection to an external server. In addition, execution in a sandbox raised an alert about an attempt to pass the NTLM hash through the SMB protocol. After receiving the NTLM hash, an attacker can carry out a Pass-the-Hash attack to identify as the user associated with the captured hash without needing the corresponding password. In other words, the credential hash that NTLM's SMB protocol sends out to identify its Windows user can simply be captured and subsequently used to impersonate the user as if they were logged in."

The researchers wrote: "Further investigation yielded that in Windows 10 and 11 operating systems, the action of dragging the file from one folder to another, or deleting the file, caused the file to communicate with the target server and only then be deleted or moved. Under Windows 7, 8, and 8.1, the file did not initiate communication when dragged or deleted unless the target folder was open at the time of dragging." They said: "This did not happen on the first attempt, but was observed only after two to three attempts. That is," they concluded, "the newly detected vulnerability is somewhat more exploitable on Windows 10 and 11 operating systems."

So I'm sure that it must be a bit unnerving to those old pros among our listeners here to learn that the actions that any of us might take to dispose of something we may have inadvertently received could themselves lead directly to a compromise of our machine. That's new. So Microsoft reportedly patched and closed this flaw in last Tuesday's patch updates, so that's good. But it should serve to remind us that those of us using Windows are using an extremely complex operating system that is still dragging a ton of legacy code forward. That code was written, that NTLM SMB file and printer sharing code was written, and its protocols were designed, long before the world had an appreciation for just how secure our future systems would need to be.

What came to mind as I was thinking about this, the classic example of this was the original design of the Windows metafile format. Windows draws on the screen through a series of drawing primitives, you know, invoking a circle or a rectangle or a line function with parameters and so forth. A Windows metafile (WMF) is just the capture of those drawing primitives. It's essentially a script. Then later, when that metafile is opened, those primitives are replayed onto a new blank canvas to recreate the original drawing. So the metafile contents are interpreted. But the designers of the original metafile format thought, what if we want to do something more, you know, something more than just replaying something that was previously recorded? Why can't the file contain some code that's executed? And remember, this was Windows 3.0.

So among all of the interpreted tokens, they specified a META_ESCAPE code, which is what it was called, that would cause the system to execute, to essentially escape from interpreting a GDI, graphics device interface tokens, and execute the code contained within the Windows metafile, starting at the bytes immediately following the special escape code.

And so it sat there in the metafile specification for years, until much later - oh, and it was copied, as like from 95 to 98 to - what was the last 16-bit version? It was ME, Windows

ME. And then it made the jump to Windows NT and so on. So later, years later, in the era of NT and networking and Internet connectivity, it was suddenly rediscovered and labeled as a horrible exploitable flaw. At the time, when I calmly stated that it was obviously there all along by design, many people misunderstood me. They thought I was saying that Microsoft had deliberately planted a backdoor in Windows metafiles. It was, you know, it was originally deliberate, but it was never malicious.

**Leo:** It was convenience.

**Steve:** Yes. Yes. It was a +reasonable thing to do back when we could trust every image our machines might try to render. But let's just say it didn't age well. And neither was Microsoft's original NT LAN Manager and their SMB protocol.

**Leo:** Right.

**Steve:** You know, they have not aged well, either. And they were also designed back before we really understood security. So this wasn't deliberate on Microsoft's part. And what was really interesting was that a week or two ago we were just talking about how Microsoft has decided not to keep patching NTLM problems, yet the 0patch guys are. So there's another reason why 0patch is worth looking at. Oh, and I should mention, I got a bunch of feedback from our listeners who said, you know, Steve, you should mention that there's a free tier also.

**Leo:** Ah.

**Steve:** So it's not necessary to subscribe to 0patch in order to get some of the benefits of it. So I just wanted to mention that, along with all the others. And thank you, everybody who wrote to say, uh, you know, there's a freebie available. So there is a free tier for 0patch.

Okay. So not a lot happened this week, and we've just covered it all. So I'm going to spend some time with some feedback from our amazing listeners.

**Leo:** Good.

**Steve:** I believe he would pronounce his name Ayiko, A-Y-I-K-O. I'm sorry if that's wrong, but I'll say Ayiko Fred is in Uganda. And he said: "Hey, Steve and Leo. This is Ayiko Fred from Uganda. I've been listening to Security Now! since 2021, starting around the 800s." As in, you know, episode number. He said: "I occasionally miss a few episodes when things get busy, sometimes up to a month; but I'm thoroughly enjoying the show!"

He said: "I do not have a formal background in computer science, but I developed an interest in programming in 2020 and learned some Erlang and Elixir," he said, "my first and only languages, which I'm now using at work." He said: "It made me realize I had only a blurry understanding of many key concepts. I'd never thought to go back to the earlier episodes from 2005, but a few episodes ago a listener recommended going back to the earlier episodes. So I decided to give it a try. And, wow!" He said: "The way you explain topics like how the Internet works, cryptography, and VPNs really clicked for me." He said: "I was blown away by how much easier it was to understand these concepts

through your explanations. Now I feel like I've been 'programming by superstition' all along."

**Leo:** I know that feeling.

**Steve:** He said: "Each episode has left me wanting more, and I've even re-listened to some episodes three to four times, especially those on cryptography and Internet fundamentals. I'm now on Episode 58, and I'd encourage anyone with a shaky grasp on these topics to check out the earlier episodes. They won't regret it."

**Leo:** Isn't that nice. That's so nice.

**Steve:** So I wanted to share that just to remind our listeners about that. But he finishes, saying: "One episode made me think 'This is exactly what I need,'" he said. "That was Episode 41, 'TrueCrypt.'" He said: "Unfortunately, I learned that TrueCrypt's development was discontinued in 2014. Do you have any recommendations for alternative tools with similar features to TrueCrypt that are compatible with Linux? I'd love something with the same level of privacy and security. Thank you again for all your work. I really appreciate it. Looking forward to Episode 1000. Best regards."

So I mentioned this bit of feedback last week that I wanted to share this week because I know that this podcast has been discovered by many people years after we recorded those early fundamental technology podcasts. We've heard from others who, after discovering this podcast, had the idea of going back to start from scratch and catch up. And those people have invariably found that it was worth their time. So, frankly, part of me is tempted to just stop and recreate some of that work from the early days so that they're put back into everyone's feeds.

But that doesn't make any sense because they're already there. Every podcast we've ever recorded remains available to everyone. And reproducing content we've already created would displace our new content, for which we often barely have enough time as it is. So from time to time I'll take a moment, as I have here, to remind our listeners that back in the early days we laid down many of the fundamentals of the way everything we're talking about today works. And it was done in a way that many people have found to be extremely accessible.

Also, another thing we often hear is that while our listeners enjoy the content today, they feel that there's much they don't understand. You know, they say, like, well, I understand maybe 20% of what you're talking about. We just mentioned that a week or two ago. You know, it is true that I consciously build upon the foundation that we have laid down before, using what's come before. That's the only way it's possible for us to move forward. So to those who feel that they've been tossed into the deep end of the pool by showing up here late, let me note that all of that knowledge that's missing and assumed was once covered in detail back in the earlier days of this podcast. Really. I mean, all of the stuff we have talked about and sort of zip over when we're talking about something new, that's all been discussed in detail in the past, and it's all there, waiting and free for the asking for anyone who wants it.

**Leo:** At some point I'd love to make a playlist of foundational episodes that people should listen to.

**Steve:** Yeah.

**Leo:** But just for Ayiko Fred, there is a replacement for TrueCrypt. Steve talks about it in Episode 582. You'll get there. It's VeraCrypt, and he talks about it in this episode and many other episodes.

**Steve:** Yup. And I have a link to VeraCrypt in the show notes, VeraCrypt.fr, VeraCrypt.fr. I went over and took a look; and, yep, I mean, it was updated a month or two ago.

**Leo:** Oh, yeah, yeah, yeah.

**Steve:** So it is being kept current, and it is platform agnostic. It'll work beautifully for Linux and encrypt your drive just like TrueCrypt once would have.

**Leo:** Very nice.

**Steve:** Yes, we've got our...

**Leo:** See, we've covered it all. We've covered it all over the years.

**Steve:** We really, we have...

**Leo:** We really have.

**Steve:** Well, Leo, how many thousands of hours?

**Leo:** That's right, several, at least.

**Steve:** Okay. Scott Gottfried wrote to share his powerful solution for accessing his network from home. But Leo, let's take a break, and then we're going to find out what Scott is using in order to get roaming access. And it's not something we've ever talked about.

**Leo:** Oh, how fun.

**Steve:** Something new.

**Leo:** Yeah, like Hamachi, or we've talked about a lot of different ways of doing stuff like that, yeah.

**Steve:** Yup. And you know Hamachi still exists?

**Leo:** Really. But it was bought by LogMeIn.

**Steve:** LogMeIn bought them, and so it's a commercial service. But it's still there.

**Leo:** And it was a great idea, using, what, five-dot; right?

**Steve:** Mm-hmm. Exactly.

**Leo:** Well, I can't wait to hear what else there is out there. Okay. On we go. More Q&A.

**Steve:** So Scott leaves to the end that everything he describes is all a free service provided by Cloudflare.

**Leo:** Ah.

**Steve:** Which is really interesting.

**Leo:** I've used their Pages. They have a lot of free services, actually.

**Steve:** Yeah. So I wanted to mention that upfront, that is, the freeness, so that while I'm sharing what Scott wrote, everyone who might have a similar need will be taking it seriously and thinking, oh, this is interesting.

So Scott said: "Hi, Steve. Congrats on 1,000. I've listened for all 20 years, every episode. Thank you and Leo." He said: "I've heard several questions from your listeners about how to access their home network while traveling. VPN? Overlay network? I had the same question. My primary requirement for accessing my home network was that I did not want to open any ports on my router." Amen to that. He said: "I researched solutions for several months until I happened upon a blog post at Cloudflare. The solution for me is the Cloudflare Tunnel." And that's at www.cloudflare.com/products/tunnel, T-U-N-N-E-L.

And he said: "I run an old Intel NUC from inside my network that creates an outgoing tunnel to Cloudflare. The Cloudflare dashboard lets me add my own domains, has a firewall, provides authentication, and allows me to configure routing for my four internal home subnets." He said: "It's awesome. I run two separate photo sharing apps for the family. The apps run in Docker containers on the NUC which has Linux and CasaOS. But the tunnel could run on a NAS or ZimaBoard.

"When traveling, I use the Cloudflare Warp app on my laptop and connect to my home network. I can then RDP to my Windows NUC. I can access my Ubiquiti cams. And I can access my TrueNAS. Nothing on the home network is exposed to the Internet. It all happens through the tunnel. The family accesses my shared photo apps, Jellyfin and Piwigo, using a web browser pointed to my custom domain. I add authorized family member email addresses to the Cloudflare dashboard. When a family member tries to log

onto one of the apps, they just enter their email address. They are sent a PIN for access. All of that is handled by Cloudflare.

"It's a little bit of a propeller beanie kind of stuff, but one could just start with the tunnel to access the home network without sharing apps and dealing with authentication. Oh," he says, "I forgot to mention, all of the stuff I use at Cloudflare is FREE!" He said: "I hope this might help anyone searching for this type of solution. Best, Scott."

So thank you, Scott, for sharing that. It was news to me, so I went over to take a look. Cloudflare's Tunnel page says: "Protect your web servers from direct attack. From the moment an application is deployed, developers and IT spend time locking it down, configuring ACLs (Access Control Lists), rotating IP addresses, and using clunky solutions like GRE tunnels. There's a simpler and more secure way to protect your applications and web servers from direct attacks: Cloudflare Tunnel. Ensure your server is safe, no matter where it's running: public cloud, private cloud, Kubernetes cluster, or even a Mac mini under your TV."

So from Scott's description, it sounds like an extremely powerful and capable solution. For simple safe remote connections to an internal network, it may be more than many of our listeners need. But I wanted to put it on everyone's radar, you know, because it really does sound like a power user's tool, you know, being able to set up authentication, have registered email addresses where someone is able to receive a PIN, provide that back, and then automatically get access through the tunnel back to the network. You know, there's a lot there. It does a lot. But anyway, it looks like a potentially very interesting solution.

At the same time I got a note from Jeff Price, who also happened to write: "Thanks for the emails. They are very helpful," he said, meaning the weekly Security Now! preview of the podcast. HE said: "I have a medium-sized network at home with Synology NAS, dozens of IOT devices, et cetera. I've been using Tailscale for all remote connections. This means no open ports or port forwarding. I also set up a system inside my home as an exit node, which means even when I am traveling I can encrypt all of my traffic back to my home and then exit from there." In other words, anything he's doing while he's traveling believes he's still at home, which can be useful for, you know, access to streaming services and so forth that have specific geographic boundaries. And he said: "Tailscale has worked great, and it is much faster than OpenVPN."

So just another reminder that the overlay network solution is almost drop-in easy to use, and there are Tailscale and ZeroTier, and there's also Nebula and Netmaker. There are clients for all of the various OSes that we're using, and even for the various NASes. So, you know, there's probably a, well, it is far less flexible and capable. It's also sort of more of a homegrown solution than Cloudflare's Tunnel. So, you know, your mileage may vary. Pick the solution that seems best for you.

Adam B. has an intriguing problem. He said: "Hi, Steve. I'm a long-time listener to the show. I'm not sure how long, but I definitely remember when you used to alternate episodes between topics and news." And he means news and feedback. He says: "I'm a proud SpinRite owner and, thanks to you and Leo getting me interested in HackerOne, a few hundred dollars better off, having found a couple of Local Privilege Escalation vulnerabilities during some poking around on my weekends." That's very cool. So he's a little bit of a white hat hacker, helping people.

He says: "I have a question that I have not been able to find an answer to online, and I thought might interest you and my fellow listeners. I'm a hobbyist malware analyst."

**Leo:** Interesting hobby, yeah.

**Steve:** Clearly, based on the experience he shared. He said: "And as part of that I often run the samples in a network that's isolated from the Internet, just to see what happens. Sometimes the samples will try to communicate with a command-and-control' server. Often, the hard-coded C2 server is a Fully Qualified Domain Name, but sometimes it's a public IP address." He says: "It can often be useful to pretend to be the command-and-control server, just to see what the sample sends. When the C2 server is a Fully Qualified Domain Name, it's easy enough to use my own DNS server in the isolated network to answer the DNS request with an A record IP address of my choosing."

Meaning that, right, so the malware says I need the IP address of badguys.ru. And because he's created an isolated network, he's got his own DNS server. So the machine running the malware generates a DNS query to badguys.ru, and the DNS responds with, you know, 192.168.0.20 or something, which is a machine on that network, so that's where the malware attempts to connect to, which is his own server, so he can see what's going on.

He said: "However, when the C2 server is a public IP address, this becomes more troublesome. I think I have two choices," he wrote. He said: "One, patch the sample to change the IP address to one on the LAN. Or, two, somehow get my LAN to answer the ARP request with a MAC address of my choosing." He said: "The problem with choice number one is that this isn't practical at scale." Meaning, you know, patching the malware in order to point it to something local. And I agree. And he said: "As in, you know, sometimes I like to run 10, 20, or 50 versions of the same malware family." He said: "I don't want to have to manually patch 50 different samples. It also seems like the less satisfactory choice.

"The problem with choice two is that I simply can't figure out how to do it. How can I configure my network so that if a sample makes a request for a public IP address, in other words, one that isn't in the /24 of my LAN, the request is handled by my C2 server? The best answer I could find online was concerned with ARP poisoning, but this seemed very unreliable and likely to cause an unstable network. It feels like the answer will be something to do with the default gateway, but I can't figure it out. I hope that makes sense. I would really appreciate your thoughts on the subject. A big thank you to you, Leo, and the whole team. Kind regards, Adam."

Okay. What Adam wants to do can definitely be done in a highly robust fashion. It would be possible to manually add static routes to the routing table of the machine that's hosting the malware. This would cause the traffic bound for that target IP to override the normal, non-local default route, which would send the traffic out to the network's gateway interface, and instead to another local network interface. But doing that is tricky and messy.

The more straightforward solution, and it's really slick, would be to obtain a router that has some extra hardware interfaces. That little Netgate SG-1100 which I'm using here has an AUX network connection, you know, it's got WAN and LAN and AUX, as in auxiliary. And it's not a simple switch using the same network as the LAN. It's a separate network interface, and that can be given its own LAN. Or, for example, one of those Protectli (P-R-O-T-E-C-T-L-I), Protectli Vault devices, I'm using one of those at my other location. Those are nice also, and Amazon has those for sale, or you can get them directly from Protectli.

The idea is to have an extra physical network interface. You would use the router's software such as pfSense or OPNsense to define another small LAN network for that extra interface. And instead of using one of the normal private networks like 192.168.x.x or 10.x.x.x, you would create a network that includes the target IP of the command-and-control server. You then attach a machine, this C2, your command-and-control spoof

server. You attach a machine to that interface and manually assign it the IP of the command-and-control server that the malware's looking for.

Now, whenever the malware in the host machine addresses Internet traffic to that remote public IP, your local router's routing table will see that the IP matches within that extra network and will send the traffic to it rather than out onto the public Internet. So you wind up with a very straightforward, robust, and easily adjusted and maintained solution. And...

**Leo:** Yes?

**Steve:** Dale Myers.

**Leo:** Okay.

**Steve:** Has a problem. I've forgotten how many breaks we've taken.

**Leo:** I thought there was something going on. We have one more. So you could put that anywhere you want.

**Steve:** Okay. Only one left, good.

**Leo:** Only one more, yeah.

**Steve:** And then we'll finish our feedback. And before we get into what is AGI...

**Leo:** Perfect, yeah.

**Steve:** Thank you. Dale Myers has a problem no one should ever face. He said: "Hi, Steve. I never thought, when I started listening at 0001 that there would ever be a thousand, and still counting, Security Now! podcasts." He said: "I started at the beginning, right after Fred Langa suggested that your podcast might be worthwhile. He was right.

"At the time I was a volunteer in the IT department of a parochial school. The things I learned from Security Now! led to important improvements in our system over the years. In those days there were not so many listeners, and you took time to answer two of my questions submitted in the 'Feedback' dialog box at the bottom of the Security Now! page. Now I have a new question that relates to using a password manager." He said: "I've been doing a bit of traveling by air lately, and the last time I was in my travel agent's office I decided to use some of the accumulated points. She said she could not access my account without my password. There was a place for it on her screen, but I could not figure out how to get the password to there from my password manager. Any thoughts? Signed, Dale Myers."

Okay. So my first thought was, huh. That's a really good question. How would you do that securely? And then I thought, I wonder why this isn't a problem we've heard about

before? And then the question answered itself, since no one should EVER have this problem. No one should ever be asked to give their password to someone else, like a travel agent, so that she could access their account. So it's not a bigger problem because it should never be required of anyone, ever. The whole thing, you know, seems like a fundamentally bad idea. But that doesn't help Dale, who apparently does have this problem, even if everyone agrees he should never have had this problem in the first place. Given that Dale has been listening since Episode 1, we know that his travel account is currently protected by a ridiculously gnarly, long, random, and impossible to manually enter or even communicate, password.

So my advice would be not to even try. Briefly change your password to something ridiculously simple to type which meets the travel system's password policies, but otherwise minimal in every way. You know, it's only going to be that way for a few minutes, so its security doesn't really matter. Once the travel points have been transferred, the account's password can either be restored to what it was before, or set to something new. Now, a workable alternative would be to just send the account's initial gnarly password via email or text to the travel agent, let her login, do whatever she needs, then change the account's password to something new and super secure once the points have been moved.

Now, having said that, I did get a piece of feedback from a listener about an incredibly cool-looking device. I've got it on the way to me because I want to understand and be able to talk about it. It is a little dongle which has a USB port, and it is a Bluetooth keyboard dongle, meaning that what Dale could do, if he had this, or if any of our listeners had this problem, Dale could have this with him, give it to the travel agent and have her plug it into her computer, you know, just any USB port.

Now, very much like the original YubiKey, this thing looks like a USB keyboard. So then there are Android and iOS and other apps for this thing. So Dale would be able to send his password through this app, and it would type into the password field on the travel agent's computer, which is kind of a cool hack. Anyway, I'll know more about it. I'll have all the details in next week's podcast for anybody who wants to jump ahead. It was not cheap. It was $37, and it's being shipped from Poland, as I recall. But still, I thought it was kind of...

**Leo:** That's a clever idea, though, yeah.

**Steve:** Kind of a cool thing. Chris C. asked: "A while back, you said something about a large company that was fined for not keeping Teams or Slack chats as required by federal law. Do you remember who this was and what the law was?"

So I replied to Chris: "I vaguely recall that in passing, but I have no specific recollection." And I said: "GRC's onsite search in the upper right of every page can be used to search only the podcast transcripts which are fully indexed. So you might be able to track down the reference that way." So that was my reply to Chris. I wanted to share this because I use GRC's search from time to time myself in the same way when I'm looking for something from our own past. You've heard me casually mention that we talked about something, whatever it was, you know, back during podcast number whatever.

So I just don't want anyone to imagine for a second that I recalled that podcast. Like Chris here, I did recall that it was something that was mentioned, but not what or when. Since I get these sorts of questions often, like that Chris asked, I just wanted to pass on to everyone that both the show notes and Elaine's precise transcripts are fully indexed, and that index can be easily searched using GRC's search box.

And I checked a little bit later. Chris had replied. He responded: "Thank you! I didn't know that was there." He said: "I found it in SN #959." He said: "Google did not help me, but the search engine on your site, 'powered by' the same company, did." So again, we do have, you know, essentially, podcast-specific search which will allow anyone to find something that they think they recall that we talked about before, but can't remember exactly where or when. You're free to keep asking me, but I'll do the same thing you could do, which is to use the little search box in the upper right of every page at GRC.

And Leo, we are ready to talk about artificial general intelligence.

**Leo:** Oh, boy.

**Steve:** Whatever that is. We'll at least maybe know what it is, even if we don't know when, about half an hour from now. But let's take our last break, and then we'll plow into that.

**Leo:** I'm excited. I'm really excited. I'm ready to take notes. I've been dying to hear this, Steve Gibson on AGI.

**Steve:** Well, okay. Steve Gibson surveying a bunch of other people's feelings about AGI.

**Leo:** Okay. That's fair. Yeah, that's fair. But I want to know what you think, too, though. I think you'll probably give us some ideas.

**Steve:** Yeah, I do have some feelings. So, okay. I should note that I already have everything I need, thanks to today's ChatGPT 4.0. And it has changed my life for the better. I've been using it increasingly as a timesaver, sort of in the form of a programming language super search engine, and even a syntax checker. I've used it sort of as a crutch when I need to quickly write some throwaway code in a language like PHP, where I do not have expertise, but I want to get something done quickly. I just, you know, I'd like to solve a quick problem. You know, parse a text file in a certain way into a different format, that sort of thing.

In the past, I would take, you know, if it was a somewhat bigger project than that, an hour or two putting queries into Google, following links to Programmer's Corner or Stack Overflow or other similar sites. And I would piece together the language construction that I needed from other similar bits of code that I would find online. Or if I was unable to find anything useful, you know, solve the problem, I would then dig deeper in through the language's actual reference texts to find the usage and the syntax that I needed and then build up from that. You know, because after you've programmed a bunch of languages, they're all sort of the same, largely. I mean, Lisp is a different animal entirely, as is APL. But, you know, the procedural languages, it's just a matter of, like, okay, what do I use for inequality, what do I use for, you know, how exactly are the looping constructs built, that kind of thing.

That's no longer what I do because I now have access to a what I consider a super programming language search engine. Now I ask the experimental coding version of ChatGPT for whatever it is I need. I don't ask it to provide the complete program, since that's really not what I want. You know, I love coding in any language because I love puzzles, and puzzles are language-agnostic. But I do not equally know the details of

every other language. You know, there's nothing ChatGPT can tell me about programming assembly language that I have not already known for decades.

But if I want to write a quick throwaway utility program like in Visual Basic .NET, a language that I spend very little time with because I like to write in assembly language, you know, but I need to, for example, quickly implement an associative array, as I did last week, rather than poking around the Internet or scanning through the Visual Basic syntax to find what I'm looking for, I'll now just pose the question to ChatGPT. I'll ask it very specifically and carefully for what I want. And in about two seconds I'll get what I may have previously spent 30 to 60 minutes sussing out online. It has transformed my work path for those sorts of - for that class of problem that I have traditionally had. It's useful whenever I need some details where I do not have expertise, is I think the way I would put it.

And I've seen plenty of criticism levied by other programmers of the code produced by today's AI. To me it seems misplaced, that is, their criticism seems misplaced, and maybe just a bit nervous. And maybe they're also asking the wrong question. I don't ask ChatGPT for a finished product because I know exactly what I want, and I'm not even sure I could specify the finished product in words, or that that's what it's really good for. So I ask it just for specific bits and pieces, and I have to report that the results have been fantastic.

I mean, it is literally, it's the way I will now code languages I don't know, I think is probably the best way to put it. It is, you know, it's ingested the Internet. And, you know, obviously we have to use the term it "knowing" them very advisedly. It doesn't know them. But whatever it is, I am able to, like, ask it a question, and I actually get, like, really good answers to tight problem domain questions.

Okay. But what I want to explore today is what lies beyond what we have today, what the challenges are and what predictions are being made about how and when we may get more, whatever that "more" is. You know, the "there" where we want to get is generically known as Artificial General Intelligence, which is abbreviated AGI. Okay. So let's start by looking at how Wikipedia defines this goal. Wikipedia says: "Artificial general intelligence is a type of artificial intelligence that matches or surpasses human cognitive capabilities across a wide range of cognitive tasks. This contrasts with narrow AI, which is limited to specific tasks. Artificial superintelligence (ASI), on the other hand, refers to AGI that greatly exceeds human cognitive capabilities. AGI is considered one of the definitions of strong AI."

They say: "Creating AGI is a primary goal of AI research and of companies such as OpenAI and Meta. A 2020 survey identified 72 active AGI research and development projects across 37 countries. The timeline for achieving AGI remains a subject of ongoing debate among researchers and experts. As of 2023, some argue that it may be possible in years or decades; others maintain it might take a century or longer; and a minority believe it may never be achieved. Notable AI researcher Geoffrey Hinton has expressed concerns about the rapid progress toward AGI, suggesting it could be achieved sooner than many expect.

"There's debate on the exact definition of AGI, and regarding whether modern Large Language Models (LLMs) such as GPT-4 are early forms of AGI. Contention exists over whether AGI represents an existential risk. Many experts on AI have stated that mitigating the risk of human extinction posed by AGI should be a global priority. Others find the development of AGI to be too remote to present such a risk.

"AGI is also known as strong AI, full AI, human-level AI, or general intelligent action. However, some academic sources reserve the term 'strong AI' for computer programs that experience sentience or consciousness. In contrast, weak AI, or narrow AI, is able to

solve one specific problem, but lacks general cognitive abilities. Some academic sources use 'weak AI' as the term to refer more broadly to any programs that neither experience consciousness nor have a mind in the same sense as humans.

"Related concepts include artificial superintelligence and transformative AI. An artificial superintelligence is a hypothetical type of AGI that is much more generally intelligent than humans, while the notion of transformative AI relates to AI having a large impact on society" - thus transforming it - "for example, similar to the agricultural or industrial revolutions.

"A framework for classifying AGI levels was proposed in 2023 by Google DeepMind researchers. They define five levels of AGI: emerging, competent, expert, virtuoso, and superhuman. For example, a competent AGI is defined as an AI that outperforms 50% of skilled adults in a wide range of non-physical tasks, and a superhuman AGI, in other words an artificial superintelligence, is similarly defined, but with a threshold of 100%. They consider Large Language Models like ChatGPT or Llama 2 to be instances of the first level, emerging AGI."

Okay. So we're getting some useful language and terminology for talking about these things. The article that caught my eye last week as we were celebrating the 1000th episode of this podcast was posted on Perplexity.ai, titled "Altman Predicts AGI by 2025." The Perplexity piece turned out not to have much meat, but it did offer the kernel of some interesting thoughts, and some additional terminology and talking points, so I still want to share it.

Perplexity wrote: "OpenAI CEO Sam Altman has stirred the tech community with his prediction that Artificial General Intelligence (AGI) could be realized by 2025, a timeline that contrasts sharply with many experts who foresee AGI's arrival much later. Despite skepticism, Altman asserts that OpenAI is on track to achieve this ambitious goal, emphasizing ongoing advancements and substantial funding, while also suggesting that the initial societal impact of AGI might be minimal.

"In a Y Combinator interview, Altman expressed excitement about the potential developments in AGI for the coming year. However, he also made a surprising claim that the advent of AGI would have 'surprisingly little' impact on society, at least initially. This statement has sparked debate among AI experts and enthusiasts, given the potentially transformative nature of AGI.

"And Altman's optimistic timeline stands in stark contrast to many other experts in the field, who typically project AGI development to occur much later, around 2050. Despite the skepticism, Altman maintains that OpenAI is actively pursuing this ambitious goal, even suggesting that it might be possible to achieve AGI with current hardware. This confidence, coupled with OpenAI's recent $6.6 billion funding round and its market valuation exceeding $157 billion, underscores the company's commitment to pushing the boundaries of AI technology. Achieving Artificial General Intelligence faces several significant technical challenges that extend beyond current AI capabilities."

So here we have four bullet points that outline what AGI needs that there's no sign of today. First, common-sense reasoning. AGI systems must develop intuitive understanding of the world, including implicit knowledge and unspoken rules, to navigate complex social situations and make everyday judgments. Number two, context awareness. AGI needs to dynamically adjust behavior and interpretations based on situational factors, environment, and prior experiences. Third, handling uncertainty. AGI must interpret incomplete or ambiguous data, draw inferences from limited information, and make sound decisions in the face of the unknown. And fourth, continual learning. Developing AGI systems that can update their knowledge and capabilities over time without losing previously acquired skills remains a significant challenge.

So one thing that occurs to me as I read those four points - reasoning, contextual awareness, uncertainty, and learning - is that none of the AIs I've ever interacted with has ever asked for any clarification about what I'm asking. That's not something that appears to be wired into the current generation of AI. I'm sure it could be simulated if it would further raise the stock price of the company doing it. But it wouldn't really matter; right? Because it would be a faked question, like that very old Eliza pseudo-therapist program from the '70s. You know, you would type into it "I'm feeling sort of cranky today," and it would reply, "Why do you think you're feeling sort of cranky today?" You know, it wasn't really asking a question, it was just programmed to seem like it was understanding what we were typing in.

The point I hope to make is that there's a hollowness to today's AI. You know, it's truly an amazing search engine technology, but it doesn't seem to be much more than that to me. There's no presence or understanding behind its answers.

The Perplexity article continues, saying: "Overcoming these hurdles requires advancements in areas such as neural network architectures, reinforcement learning, and transfer learning. Additionally, AGI development demands substantial computational resources and interdisciplinary collaboration among experts in computer science, neuroscience, and cognitive psychology.

"While some AI leaders like Sam Altman predict AGI by 2025, many experts remain skeptical of such an accelerated timeline. A 2022 survey of 352 AI experts found that the median estimate for AGI development was around 2060 - also known as Security Now! Episode 2860. 90% of the 352 experts surveyed expect to see AGI within 100 years. 90% expected, so not to take longer than 100 years, but the median is by 2060, so not next year, as Sam suggests."

They wrote: "This more conservative outlook stems from several key challenges. First, the 'missing ingredient' problem. Some researchers argue that current AI systems, while impressive, lack fundamental components necessary for general intelligence. Statistical learning alone may not be sufficient to achieve AGI." Again, the missing ingredient problem. I think that sounds exactly right.

"Also, training limitations. Creating virtual environments complex enough to train an AGI system to navigate the real world, including human deception, presents significant hurdles. And third, scaling challenges. Despite advancements in Large Language Models, some reports suggest diminishing returns in improvement rates between generations. These factors contribute to a more cautious view among many AI researchers, who believe AGI development will likely take decades rather than years to achieve.

"OpenAI has recently achieved significant milestones in both technological advancement and financial growth. The company successfully closed" - and here they're saying again a massive $6.6 billion funding round, valuing it at $157 billion. But, you know, who cares? That's just, you know, Sam is a good salesman.

They said: "This round attracted investments from major players like Microsoft, Nvidia, and SoftBank, highlighting the tech industry's confidence in OpenAI's potential. The company's flagship product, ChatGPT, has seen exponential growth, now boasting over 250 million weekly active users." And you can count me among them. "OpenAI has also made substantial inroads into the corporate sector, with 92% of Fortune 500 companies reportedly using its technologies. Despite these successes, OpenAI faces challenges, including high operational costs and the need for extensive computing power. The company is projected to incur losses of about $5 billion this year, primarily due to the expenses associated with training and operating its Large Language Models."

So when I was thinking about this idea of we're just going to throw all this money at it, and it's going to solve the problem, and oh, look, you know, the solution is going to be next year, the analogy that hit me was curing cancer because there sort of is an example of, you know, oh, look, we had a breakthrough, and this is going to cure cancer. It's like, no. We don't really understand enough yet about human biology to say that we're going to do that.

And, you know, I know that the current administration has been, you know, these cancer moon shots. And it's like, okay, have you actually talked to any biologists about this, or do you just think that you can pour money on it, and it's going to do the job? So that's not always the case. So to me, this notion of the missing ingredient is the most salient of all of this, is like what we may have today has become very good at doing what it does. But it may not be extendable. It may never be what we need for AGI. But I think that what I've shared so far gives a bit of calibration about where we are and what the goals of AGI are.

I found a piece also in Information Week where the author did a bunch of interviewing and quoting of people that I just want to share just to finish this topic off. It was titled "Artificial General Intelligence in 2025: Good Luck With That." And it had the teaser "AI experts have said it would likely be 2050 before AGI hits the market. OpenAI CEO Sam Altman says 2025, but it's a very difficult problem to solve."

So they wrote: "A few years ago, AI experts were predicting that artificial general intelligence would become a reality by 2050. OpenAI has been pushing the art of the possible, along with Big Tech; but despite Sam Altman's estimate of 2025, realizing AGI is unlikely soon.

"HP Newquist, author of 'The Brain Makers' and executive director of The Relayer Group, a consulting firm that tracks the development of practical AI, said: 'We can't presume that we're close to AGI because we really don't understand current AI, which is a far cry from the dreamed-of AGI. We don't know how current AIs arrive at their conclusions, nor can current AIs even explain to us the processes by which that happens. That's a huge gap that needs to be closed before we can start creating an AI that can do what every human can do. And a hallmark of human thinking, which AGI will attempt to replicate, is being able to explain the rationale for coming up with a solution to a problem or an answer to a question. We're still trying to keep existing Large Language Models from hallucinating.'"

And I'll just interrupt to say that I think this is the crucial point. Earlier, I described ChatGPT as being a really amazingly powerful Internet search engine. Partly that's because that's what I've been using it to replicate. For my own needs, as I said, it's been a miraculous replacement for a bunch of searching I would otherwise need to do myself. My point is, this entire current Large Language Model approach may never be more than that. This could be a dead end. If so, it's a super useful dead end. But it might not be the road to AGI at all. It might never amount to being more than a super spiffy search engine.

The InfoWeek article continues: "OpenAI is currently alpha testing advanced voice mode, which is designed to sound human, such as pausing occasionally when one speaks to draw a breath. It can also detect emotion and non-verbal clues. This advancement will help AI seem more human-like, which is important, but there's more work to do." And frankly, that's where we begin to get into the category of parlor tricks, in my opinion. Like, you know, making it seem like more than it is, but it still isn't.

"Edward Tian, CEO of ZeroGPT, which detects generative AI's use in text, also believes the realization of AGI will take time. In an email interview with the article's author, Edward said: 'The idea behind artificial general intelligence is creating the most human-

like AI possible, a type of AI that can teach itself and essentially operate in an autonomous manner. So one of the most obvious challenges is creating AI in a way that allows the developers to be able to take their hands off eventually, as the goal is for it to operate on its own.

"'Technology, no matter how advanced, cannot be human, so the challenge is trying to develop it to be as human as possible. That also leads to ethical dilemmas regarding oversight. There are certainly a lot of people out there who are concerned about AI having too much autonomy and control, and those concerns are valid. How do developers make AGI while also being able to limit its abilities when necessary? Because of all these questions and our limited capabilities and regulations at the present, I do not believe that 2025 is realistic.'

"Current AI - which is artificial narrow intelligence (ANI), performs a specific task well, but it cannot generalize that knowledge to suit a different use case.

"Max Li, the CEO of the decentralized AI data provider Oort and an adjunct associate professor in the department of electrical engineering at Columbia University, said: 'Given how long it took to build current AI models, which suffer from inconsistent outputs, flawed data sources, and unexplainable biases, it would likely make sense to perfect what already exists rather than start working on even more complex models. In academia, for many components of AGI, we do not even know why it works, nor why it does not work.'

"To achieve AGI, a system needs to do more than just produce outputs and engage in conversation, which means that LLMs alone won't be enough. Alex Jaimes, chief AI officer at the AI company Dataminr, said in an email interview: 'It should also be able to continuously learn, forget, make judgments that consider others, including the environment in which the judgments are made, and a lot more. From that perspective, we're still very far. It's hard to imagine AGI that doesn't include social intelligence, and current AI systems don't have any social capabilities, such as understanding how their behavior impacts others, cultural and social norms, et cetera.'

"Sergey Kastukevich, the deputy CTO at the gambling software company SOFTSWISS said: 'To get to AGI, we need advanced learning algorithms that can generalize and learn autonomously, integrated systems that combine various AI disciplines, massive computational power, diverse data, and a lot of interdisciplinary collaboration. For example, current AI models like those used in autonomous vehicles require enormous datasets and computational power just to handle driving in specific conditions, let alone achieve general intelligence.'

"LLMs are based on complex transformer models. While they are incredibly powerful and even have some emergent intelligence, the transformer is pre-trained and does not learn in real-time. For AGI, there will need to be some breakthroughs with AI models. They will need to be able to generalize about situations without having to be trained on a particular scenario. A system will also need to do this in real-time, just like a human can when they intuitively understand something.

"In addition, AGI capabilities may need a new hardware architecture, such as quantum computing, since GPUs will probably not be sufficient. Note that Sam Altman has specifically disputed this and said that current hardware will be sufficient. In addition, the hardware architecture will need to be much more energy efficient and not require massive data centers.

"LLMs are beginning to do causal inference and will eventually be able to reason. They'll also have better problem-solving and cognitive capabilities based on the ability to ingest data from multiple sources."

So, okay. What's interesting is the degree of agreement that we see among separate experts. You know, they're probably all reading the same material, so there's some degree of convergence in their thinking. But, you know, Altman is an outlier. And it seems to me as though these people know what they're talking about from the things they've said. Perhaps, you know, maybe Sam has already seen things in the lab at OpenAI that no one else in the outside world has seen, because that's what it would take for Sam to not be guilty of over-hyping and over-promoting his company's near-term future.

Now, I put a picture in the show notes, you had it on the screen there a second ago, Leo, that is not a mockup. That is not a simulation. This is an actual image of a tiny piece of cerebral tissue. Those are neurons and axons and dendrites. The coloration was added. But that is actual human brain tissue in that photo in the show notes. I'm especially intrigued by the comments from the top academic AI researchers in the world who admit that, to this day, no one actually understands how Large Language Models produce what they do. Given that, I'm skeptical that just "more of the same" will result in the sort of qualitative advancement that AGI would require, which is certainly not just more of the same.

When I said in the past that I see no reason why a true artificial intellect could not eventually be created, I certainly did not mean next year. I meant someday. I meant that I believe that a biological brain may only be one way to create intelligence. One thing I've acquired during my research into the biology of the human brain is a deep appreciation for the astonishing complexity, I mean astonishing, of the biological computing engine that is us. The number of individual computing neurons in the human brain is 10^11; okay? So that's 100 billion, 100 billion individual neurons. A billion neurons 100 times over.

So consider that, a billion neurons a hundred times. And not only are these individual neurons very richly interconnected, typically having connections to 20,000 others, each individual neuron is, all by itself, individually, astonishingly complex in its behavior and operation. They are far from being simple integrative binary triggers like, you know, we learned in elementary school. And we have 100 billion of these little buggers in our heads.

So perhaps Sam is going to surprise the rest of the world next year. We'll see. Color me skeptical, but not disappointed. As I said, I'm quite happy to have discovered the wonderful, language-accessible, Internet digest that ChatGPT is. You know, that's more than a simple parlor trick. It's a big deal. And it's, I think, kind of magic. But I suspect that all it is, is what it is. And for me, that's enough for now. I'd wager that we have a long ways to wait before we get more.

> **Leo:** How would you know if something is in an AGI? That's one of the things that's bothered me. The Turing test is not real.

**Steve:** No.

> **Leo:** There's a Chinese room test that may be a little better. I think there's really no way to judge an AGI.

**Steve:** No. I mean, it would, well, another perfect example is chess. Once upon a time you could have easily said, well, humans are like, you know, humans can play chess. No machine could play chess.

**Leo:** Right.

**Steve:** Right? I mean, that was something people were saying for a long time.

**Leo:** Right, right.

**Steve:** Now we just, you know, the computers have blown past us. So, and for me, and I know that you have also used constrained domain Large Language Models which you've trained by dumping all of a bunch of Lisp textbooks into it, and then been able to ask questions. You know, this is a fantastic technology that we have.

**Leo:** Right.

**Steve:** But I think it's very much in the same way that, like, the solution we have for cancer is by using chemotherapy to limit growth of our whole body because cancer cells are a problem because they're able to reproduce at such a high rate, I mean, it's like we haven't even begun to start an actual cure. We just have sort of mitigation that is able to push people into remission. So my feeling is that I agree with the experts who suggest that what we may see today we should regard as nothing more than what it is. And there's no reason to believe that we're going to get some sort of transformation just by getting more of the same.

**Leo:** Yeah. I also think that looking for an AGI is maybe not really the sensible end goal, that machines could be as useful as an AGI or as powerful as an AGI without actually being a general intelligence. I don't know if that's a reasonable thing to be measuring, AI progress.

**Steve:** Well, it is certainly the case that, if we had something where people could describe casually exactly how they wanted a computer program to operate and actually, like, got a functioning error-free, bug-free...

**Leo:** We're close to that, by the way, yeah.

**Steve:** ...thing, that would be transformative for the world of coding.

**Leo:** Right. We're very close to that.

**Steve:** And I would not be surprised, yes, I would not be surprised if we don't have something like that before long.

**Leo:** I asked one of my favorite AIs, Perplexity AI, which is a search, Internet search engine - you should give it a try since that's how you seem to think or seem to like using AI. So I asked is there a test for AGI. It mentions the Turing test, some other tests. But then it mentioned some casual tests like the coffee test. An AI enters an

average American home and figures out how to make coffee. You know what, if a robot could do that, it may not be AGI, but, boy, that's impressive. Or could go to college, enrolls in a university, obtains a degree, passing the same classes as humans. I think we might be close to that. The IKEA test, an AI controls a robot to assemble flat pack furniture correctly after viewing parts and instructions. Many humans can't do that. So that would be an interesting test, as well.

I just, I think that those are obviously kind of silly, but that points out there is no kind of accepted definition for what AGI is. And there are many different ways, just as with humans there are many ways to be intelligent, I think there are many ways for a machine to be usefully intelligent. If a machine could come in my house and make coffee without any, you know, advanced knowledge about that except kind of maybe a basic idea of what coffee is and how to make it, I'd be impressed. I think that would be useful. May not be AGI, but it'd be pretty cool. Anyway, I think that's going to happen in our lifetime.

**Steve:** When we were growing up, there was a game, it was called Nim.

**Leo:** Yeah, loved Nim.

**Steve:** And there was a way to set up a computer using matchboxes and matchsticks...

**Leo:** Right.

**Steve:** ...where you would - basically this thing was like a very early combinatorial computer. And by iterating on this, you were training it to make the right decisions over time about how many sticks to take away when a certain number of matchsticks remained. And, I mean, this is the kind of stuff that fascinated me as I was a kid. I wasn't climbing stairs on the outside of the banister. I was, you know...

**Leo:** But, see, that's combinatorial math. And you can easily see how it would be simple to program something.

**Steve:** Yup.

**Leo:** You know, I have kind of a famous book, a Lisp book, as it turns out, by Peter Norvig called "Paradigms of Artificial Intelligence Programming." And it talks about some of the - this is an early book. I think it's 30 or 40 years old now. It's in public domain, it's that old. But he talks about some of the early attempts to do what he called a GPS, a General Problem Solving machine. And it's basically that. It's a combinatorial thing. We'll try this, and then this, and the this. And if that doesn't work, backtrack and then try this and this and this. And you could see how you could solve chess that way, given a fast enough machine. Or even Go, which is a lot more difficult to play than chess. Or protein folding, a lot of things. Those are useful tools. Maybe not intelligence. But we don't even know what human intelligence is. So I don't know how we [crosstalk].

**Steve:** Yeah, and I think you're right. When you mention protein folding, there are many people who are expecting, like, that what we have now, or could have in a year or two, could make, you know, dramatically change healthcare by, like, you know, looking at mass amounts of data and pulling associations and relationships out of that that we don't see.

**Leo:** Right.

**Steve:** Because it just has a scope that we don't have.

**Leo:** And that's really more a question of...

**Steve:** Applicable.

**Leo:** Yeah, and it has something to do more with capacity, the amount of data it can store, which is so much vaster than a human mind. The amount of speed with which it can process it, again, faster than a human mind. That doesn't make it intelligent. It just makes it faster and bigger and better.

**Steve:** Yeah.

**Leo:** In some ways. I think it's a fascinating subject. And you probably feel the same way. As science fiction fans, I think we both would love to see AGI in our lifetime. Just be fun to talk to an alien intelligence that we created.

**Steve:** It would certainly be the case that creating a conversation would be a next step.

**Leo:** Oh, yeah.

**Steve:** Where if you actually got a sense of, you know, there being something there. I just, you know, I get no sense that it's anything other than...

**Leo:** No.

**Steve:** And it's clearly, you know, it refers to itself in the first person. You know, it's like let me know if there's anything more I can do for you. So they're like, you know, they gave it a bunch of sugarcoating...

**Leo:** Right.

**Steve:** ...that is designed to make us think like, you know...

**Leo:** Exactly. We anthropomorphize it.

**Steve:** ...like we're talking to an entity. There's not an entity.

**Leo:** Even the word "hallucination" really is an inappropriate anthropomorphization of what's really going on.

**Steve:** Yeah, calling it a mistake is a...

**Leo:** It's a mistake.

**Steve:** It's a mistake.

**Leo:** It's an error. Steve, as always, fascinating show. Great information. Lots of food for thought.

**Steve:** Bye.

**Leo:** Bye.