Security Now! #1003 - 12-03-24

A Light-Day Away

This week on Security Now!

Microsoft makes very clear what data they are NOT using to train their AI models. What's a "Digital Epileptic Seizure"? What induces them? And why you don't want your self-driving car to have one! A public plea for help in the form of volunteer bridge servers from the Tor Network. If you are one of 140 million Zello users, heed their notice to change your password. The U.S. Federal Trade Commission opens a broad antitrust investigation into whether Microsoft has been naughty or nice. A new form of Android smartphone "scareware" simulates a seriously malfunctioning, cracked and broken screen. It's almost certainly positively and completely safe to leave Wireguard open and listening for incoming connections. Is "almost certainly positively and completely safe" safe enough? If the Internet fills with AI output, what happens when AI starts training on that? It seems we know. Last week, Australia passed the social media age restriction law. Now what? And finally, not only is Voyager 1 nearly an entire light-day away, it's beginning to have some harder to remotely repair problems. How much longer will we be in touch with it?

"Irony" defined:

Security News

"We are NOT training OUR AI on YOUR data"

Microsoft felt the need to clarify what had become the widespread misapprehension that they would be training their AI models against the private and personal data of their Office product users. Last Wednesday, BleepingComputer did a great job of summing up the situation:

Microsoft has denied claims that it uses Microsoft 365 apps (including Word, Excel, and PowerPoint) to collect data to train the company's artificial intelligence (AI) models. This comes after a Tumblr blog post spread on social media, claiming that Redmond used their "Connected Experiences" feature to scrape customers' Word and Excel data for AI training.

*A Microsoft spokesperson told BleepingComputer: "Microsoft does not use customer data from Microsoft 365 consumer and commercial applications to train large language models. Additionally, the Connected Services setting has no connection to how Microsoft trains large language models." The company also told BleepingComputer that this optional setting has been **on by default** since it was first made available in April 2019. BleepingComputer was told:*

"The Connected Experiences feature enables features like co-authoring, real-time grammar suggestions, and web-based resources. These features are on by default because they're features people naturally expect in a cloud-connected productivity tool. However, customers always have control and can adjust their Connected Experiences settings at any time"

As Microsoft explains on its support website, the feature is used to:

- *Provide design recommendations, editing suggestions, or data insights based on the Office content, through features like PowerPoint Designer or Translator,*
- *Or download online content templates, images, 3D models, videos, and reference materials, including but not limited to Office templates or PowerPoint QuickStarter.*

To toggle this feature on or off, Microsoft 365 users have to open their Office apps (like Word or Excel) and choose whether to enable or disable experiences that download online content or analyze their content under "Connected experiences" after going to the File > Account > Account Privacy > Manage Settings menu.

<quote> "The Connected Experiences setting enables cloud-backed features designed to increase your productivity in the Microsoft 365 apps like suggesting relevant information and images from the web, real-time co-authoring and cloud storage, and tools like Editor in Word that provide spelling and grammar suggestions. Microsoft has been using AI in Microsoft 365 for years to enhance productivity and creativity through features like Designer in PowerPoint, which helps create visually compelling slides, and Editor in Word, which provides grammar and writing suggestions. These features do not rely on generative AI or Large Language Models but rather use simpler machine learning algorithms."

Microsoft added that the setting has been available since April 2019, with enterprise admins having the option to choose if connected experiences are available to users within their organizations using multiple policy settings designed to manage privacy controls for Microsoft 365 Apps and Office on Mac, iOS, and Android devices.

I'm glad for the clarification. Whatever Microsoft is doing exactly, and unless anything has changed recently, it's been doing whatever it is for the past five years. It's always been "on" by default and anyone who isn't comfortable with this is free to turn it off if they wish. If nothing else, it seems pretty clear that this has nothing whatsoever to do with CoPilot+ and any of the recent concerns over Microsoft's AI being used to otherwise enhance their users' experiences.

It's one thing to be mistrustful and another thing to accuse them wrongly. We can certainly have one without the other. Given what I 've witnessed first hand, of what they've done to Window's Start menu, tray and Edge – none of which enhances my own use of Windows – I'm obviously not a big fan of the direction they're taking their consumer desktop. Nevertheless, make no mistake, I love Windows. For my purposes it's far better than any alternative. I'm hopeful that when I set up my next Windows desktop, my Microsoft Developer access to the Enterprise edition of Windows 10 will provide me with the cleaner experience I look for in a **tool** rather than from a **toy**.

"Digital Epileptic Seizures"

I was put onto some new research from our friends at the Ben-Gurion University of the Negev and Fujitsu, by one of the researchers who is also one of our listners, Ben Nassi. The title of their 21-page paper is *"Securing the Perception of Advanced Driving Assistance Systems Against Digital Epileptic Seizures Resulting from Emergency Vehicle Lighting."* I suppose it's unavoidable to anthropomorphize driving assistance systems. But calling this problem "*Digital Epileptic Seizures*" somehow rubs me the wrong way. While the overlap is flashing lights, the underlying mechanisms have no similarity whatsoever. I'm unsure what bothers me about it, but something does.

In any event, it turns out that driving assistance systems have a problem with the flashing lights used by emergency vehicles. WIRED has a nice summary of the very good research this group has conducted. Under WIRED's headline *"Emergency Vehicle Lights Can Screw Up a Car's Automated Driving System"* with the subhead: *"Newly published research finds that the flashing lights on police cruisers and ambulances can cause "digital epileptic seizures" in image-based automated driving systems, potentially risking wrecks."* WIRED wrote:

> *Carmakers say their increasingly sophisticated automated driving systems make driving safer and less stressful by leaving some of the hard work of knowing when a crash is about to happen—and avoiding it—to the machines. But new research suggests some of these systems might do the virtual opposite at the worst possible moment.*
>
> *A new paper from researchers at Ben-Gurion University of the Negev and the Japanese technology firm Fujitsu Limited demonstrates that when some camera-based automated driving systems are exposed to the flashing lights of emergency vehicles, they can no longer confidently identify objects on the road. The researchers call the phenomenon a "digital epileptic seizure"—**epilepticar** for short—where the systems, trained by artificial intelligence to distinguish between images of different road objects, fluctuate in effectiveness in time with the emergency lights' flashes. The effect is especially apparent in darkness, the researchers say.*
>
> *Emergency lights, in other words, could make automated driving systems less sure that the car-shaped thing in front of them is actually a car. The researchers write that the flaw "poses a significant risk" because it could potentially cause vehicles with automated driving systems enabled to "crash near emergency vehicles" and "be exploited by adversaries to cause such accidents."*

While the findings are alarming, this new research comes with several caveats. For one thing, the researchers were unable to test their theories on any specific driving systems, such as Tesla's famous Autopilot. Instead, they ran their tests using five off-the-shelf automated driving systems embedded in dashcams purchased off of Amazon. (These products are marketed as including some collision detection features, but for this research, they functioned as cameras.) They then ran the images captured on those systems through four open source object detectors, which are trained using images to distinguish between different objects. The researchers aren't sure whether any automakers use the object detectors tested in their paper. It could be that most systems are already hardened against flashing light vulnerabilities.

While this might appear to render the value of this research questionable, there was at least good reason to wonder, and the researcher's findings bore this out:

The research was inspired by reports that Teslas using the electric carmaker's advanced driver assistance feature, Autopilot, collided with some 16 stationary emergency vehicles between 2018 and 2021, says Ben Nassi, a cybersecurity and machine learning researcher at Ben-Gurion University who worked on the paper. "It was pretty clear to us from the beginning that the crashes might be related to the lighting of the emergency flashers," says Nassi. "Ambulances, police cars and fire trucks are different shapes and sizes, so it's not the type of vehicle that causes this behavior."

A three-year investigation by the US National Highway Traffic Safety Administration into the Tesla-emergency vehicle collisions eventually led to a sweeping recall of Tesla Autopilot software, which is designed to perform some driving tasks—like steering, accelerating, braking, and changing lanes on certain kinds of roads—without a driver's help. The agency concluded that the system inadequately ensured drivers paid attention and were in control of their vehicles while the system was engaged. (Other automakers' advanced driving assistance packages, including General Motors' Super Cruise and Ford's BlueCruise, also perform some driving tasks but mandate that drivers pay attention behind the wheel. Unlike Autopilot, these systems work only in areas that have been mapped.)

In a written statement sent in response to WIRED's questions, Lucia Sanchez, a spokesperson for NHTSA, acknowledged that emergency flashing lights may play a role. "We are aware of some advanced driver assistance systems that have not responded appropriately when emergency flashing lights were present in the scene of the driving path under certain circumstances."

Tesla, which disbanded its public relations team in 2021, did not respond to WIRED's request for comment. The camera systems the researchers used in their tests were manufactured by HP, Pelsee, Azdome, Imagebon, and Rexing; none of those companies responded to WIRED's requests for comment.

Although the NHTSA acknowledges issues in "some advanced driver assistance systems," the researchers are clear: They're not sure what this observed emergency light effect has to do with Tesla's Autopilot troubles. Ben Nassi said: "I do not claim that I know why Teslas crash into emergency vehicles. I do not know even if this is still a vulnerability."

The researchers' experiments were also concerned solely with image-based object detection. Many automakers use other sensors, including radar and lidar, to help detect obstacles in the road. A smaller crop of tech developers—Tesla among them—argue that image-based systems augmented with sophisticated artificial intelligence training can enable not only driver

assistance systems, but also completely autonomous vehicles. Last month, Tesla CEO Elon Musk said the automaker's vision-based system would enable self-driving cars next year.

Indeed, how a system might react to flashing lights depends on how individual automakers design their automated driving systems. Some may choose to "tune" their technology to react to things it's not entirely certain are actually obstacles. In the extreme, that choice could lead to "false positives," where a car might hard brake, for example, in response to a toddler-shaped cardboard box. Others may tune their tech to react only when it's very confident that what it's seeing is an obstacle. On the other side of the extreme, that choice could lead to the car failing to brake to avoid a collision with another vehicle because it misses that it is another vehicle entirely.

The Ben-Gurion University and Fujitsu researchers did come up with a software fix to the emergency flasher issue. It's designed to avoid the "seizure" issue by being specifically trained to identify vehicles with emergency flashing lights. The researchers say it improves object detectors' accuracy.

Earlence Fernandes, an assistant professor of computer science and engineering at University of California, San Diego, who was not involved in the research, said it appeared "sound." "Just like a human can get temporarily blinded by emergency flashers, a camera operating inside an advanced driver assistance system can get blinded temporarily," he says.

For researcher Bryan Reimer, who studies vehicle automation and safety at the MIT AgeLab, the paper points to larger questions about the limitations of AI-based driving systems. Automakers need "repeatable, robust validation" to uncover blind spots like susceptibility to emergency lights, he says. He worries some automakers are "moving technology faster than they can test it."

Okay. So my own take on this is that this sort of independent research is vitally important. It needs to be done by some. It's obvious that the various car manufacturers are holding their cards – and their cars – very close to their vests. They understandably consider their future auto-driving technology to be ultra proprietary and no one else's business. Yet flesh and blood human beings and beloved pets are moving within the same space as these autonomous high-speed rolling robots. It's a recipe for disaster and this has the feeling of being driven by the same sort of gold rush mentality as the push for Artificial General Intelligence. So the headlines that these researchers have generated will doubtless, of nothing else, induce all of the developers of similar self-driving technology to consider and test the effects of bright flashing lights on their driving AI. The lives of people and pets have probably been saved.

If anyone's interested in digging deeper, the show notes has links to their overview and detailed research paper: https://sites.google.com/view/epilepticar
https://drive.google.com/file/d/1uMk3IFfggywK7aUlpzmK-0jTg6bD-Vtp/view

A Public Plea from The Tor Network

Last Thursday the Tor Network posted their plea for volunteer help. They wrote:

Recent reports from Tor users in Russia indicate an escalation in online censorship with the goal of blocking access to Tor and other circumvention tools. This new wave includes attempts to block Tor bridges and pluggable transports developed by the Tor Project, removal of circumvention apps from stores, and targeting popular hosting providers, shrinking the space for bypassing censorship. Despite these ongoing actions, Tor remains effective.

One alarming trend is the targeted blocking of popular hosting providers by Roscomnadzor. As many circumvention tools are using them, this action made some Tor bridges inaccessible to many users in Russia. As Roscomnadzor and internet service providers in Russia are increasing their blocking efforts, the need for more WebTunnel bridges has become urgent.

Why Webtunnel bridges?

Webtunnel is a new type of bridge that is particularly effective at flying under a censors's radar. Its design blends itself into other web traffic, allowing a user to hide in plain sight. And since its launch earlier this year, we've made sure to prioritize small download sizes for more convenient distribution and simplified the support of uTLS integration further mimicking the characteristics of more widespread browsers. This makes Webtunnel safe for general users because it helps conceal the fact that a tool like Tor is being used.

We are calling on the Tor community and the Internet freedom community to help us scale up WebTunnel bridges. If you've ever thought about running a Tor bridge, now is the time. Our goal is to deploy 200 new WebTunnel bridges by the end of this December (2024) to open secure access for users in Russia.

Tor's posting goes on to explain how to set up and run a WebTunnel. Among other things it can be as straightforward as hosting a Docker image. I have a link to this posting in the show notes: https://blog.torproject.org/call-for-webtunnel-bridges/

Since we haven't looked closely at Tor's WebTunnel technology I wanted to share a bit about it from their description last March which was titled *"Hiding in plain sight: Introducing WebTunnel"* https://blog.torproject.org/introducing-webtunnel-evading-censorship-by-hiding-in-plain-sight/

Today, March 12th, on the World Day Against Cyber Censorship, the Tor Project's Anti-Censorship Team is excited to officially announce the release of WebTunnel, a new type of Tor bridge designed to assist users in heavily censored regions to connect to the Tor network. Available now in the stable version of Tor Browser (which is, as we know, based upon Firefox), WebTunnel joined our collection of censorship circumvention tech developed and maintained by The Tor Project.

The development of different types of bridges are crucial for making Tor more resilient against censorship and stay ahead of adversaries in the highly dynamic and ever-changing censorship landscape. This is especially true as we're going through the 2024 global election megacycle, the role of censorship circumvention tech becomes crucial in defending Internet Freedom.

If you've ever considered becoming a Tor bridge operator to help others connect to Tor, now is an excellent time to get started! You can find the requirements and instructions for running a WebTunnel bridge in the Tor Community portal.

So what is WebTunnel and how does it work?

WebTunnel is a censorship-resistant pluggable transport designed to mimic encrypted web traffic (HTTPS). It works by wrapping the payload connection into a WebSocket-like HTTPS connection, appearing to network observers as an ordinary HTTPS (WebSocket) connection. So, for an onlooker without the knowledge of the hidden path, it just looks like a regular HTTP

connection to any web server giving the impression that the user is simply browsing the web.

In fact, WebTunnel is so similar to ordinary web traffic that it can coexist with a website on the same network endpoint, meaning the same domain, IP address, and port. This coexistence allows a standard traffic reverse proxy to forward both ordinary web traffic and WebTunnel to their respective application servers. As a result, when someone attempts to visit the website at the shared network address, they will simply perceive the content of that website address and won't notice the existence of a secret bridge (WebTunnel).

WebTunnel's approach of mimicking known and typical web traffic makes it effective in scenarios where there is a protocol allow list and a deny-by-default network environment.

Consider a network traffic censorship mechanism as a coin sorting machine, with coins representing the flowing traffic. Traditionally, such a machine checks if the coin fits a known shape and allows it to pass if it does or discards it if it does not. In the case of fully encrypted, unknown traffic, as demonstrated in the published research How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic, which doesn't conform to any specific shape, it would be subject to censorship. In our coin analogy, not only must the coin not fit the shape of any known blocked protocol, it also needs to fit a recognized allowed shape--otherwise, it would be dropped. WebTunnel traffic resembling HTTPS traffic, a permitted protocol, will pass.

It is so cool. What that means is that any regular website can also be hosting a Tor WebTunnel at the same IP and port, and no one would ever be the wiser. Since, in this case, Russia or any other censoring regime, would be unable to detect that someone is not just visiting that website, the traffic would not be blocked. But this also makes it clear that the more pseudo websites are available, the better. So if any of our listeners is moved to help the Tor project, and specifically Russian citizens who are unable to see out past their country's censorship that's being enforced for propaganda purposes, the Tor Project needs you! To make following up on this easier, I created a GRC shortcut link: https://grc.sc/tor

Zello asks their 140 million users to please change their passwords

"Zello" is a mobile PTT – Push-To-Talk – service used by 140 million first responders, hospitality services, transportation, and family and friends to communicate via their mobile phones using a simple push-to-talk app. The news is that over the past two weeks, starting on November 15th, Zello's customers have been receiving legitimate notices from Zello asking them to change their passwords. The notice reads: *"Zello Security Notice - As a precaution, we are asking that you reset your Zello app password for any account created before November 2nd, 2024. We also recommend that you change your passwords for any other online services where you may have used the same password."*

Well... it certainly doesn't take a rocket scientist, nor anyone who's been following this podcast for more than a few months to know what must have happened over at Zello headquarters. And it's not good news. But Zello is also not saying. BleepingComputer has reached out to Zello and been rebuffed. Customers who received that notice told BleepingComputer that they had not received any further information from Zello, and BleepingComputer's repeated attempts to contact the company have gone unanswered.
So, at this point it's unclear whether Zello may have suffered a data breach or a credential stuffing attack, but the notice certainly does imply that threat actors may have access to the

passwords of any users who had accounts before November 2nd. BleepingComputer noted in their reporting of this that Zello had previously suffered a data breach in 2020, which also required users to reset their passwords after threat actors stole customers' email addresses and hashed passwords.

In any event, 140 million users is a substantial user base. So if our listeners or anyone they know may be affected it might be a good idea to heed this notice.

Microsoft in the crosshairs — again

Meanwhile, the U.S. Federal Trade Commission has opened an antitrust Microsoft probe, announcing a broad antitrust investigation into Microsoft's business practices. The investigation will cover the company's software licensing practices, cloud computing, cybersecurity, and AI business units. The FTC allegedly received complaints that Microsoft was locking-in customers (Gee, like the U.S. Government, perhaps?), preventing them from moving to competitors. In September, Google filed an official antitrust complaint against Microsoft's cloud business in the EU. So this will be something to keep an eye on.

New scareware tactics

Under the headline *"You mean this actually convinces someone?"* security researcher Lukas Stefanko has identified a new form of Android scareware uses what he refers to as "convincing full screen images" that resemble cracked or malfunctioning screens which trick users into calling tech support numbers or downloading malware on their devices.

I've included a photo of this malware in action in the show notes. I can see how a neophyte might be led to believe that something has gone very wrong with their phone because the screen looks like it's no longer even remotely able to display an image.

The only problem with this is that it's having no problem whatsoever – apparently despite the cracked and malfunctioning screen – displaying the malware's warning pop-up notice claiming that a virus has been detected on the handset.

I suppose we'll give them points for coming up with something new. And, actually, I imagine it would probably succeed.

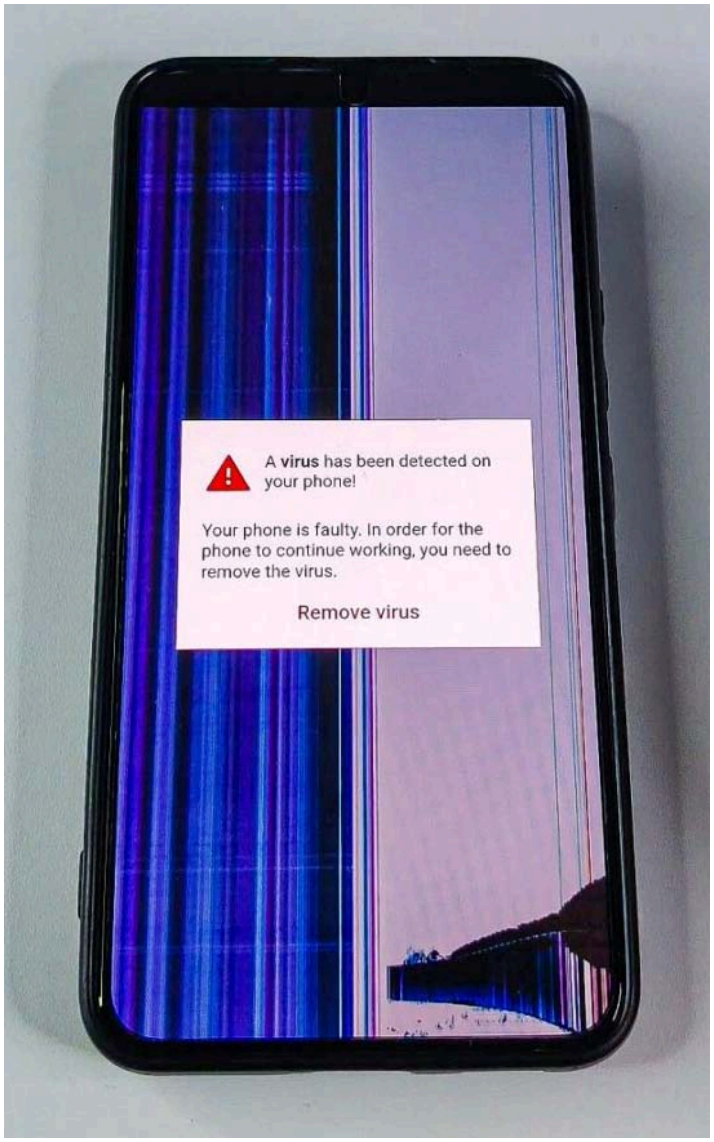

Closing The Loop

Matt Warner

Hi Steve, Regarding your comment about Wireguard's static ports in ep. 1002: I run WireGuard on an OPNsense firewall with Suricata and Crowdsec watching my WAN interface. Neither ShieldsUp! nor any other port scanner could find an open port, even when I specify the port number. I don't have WireGuard mapped to a specific allowable IP because that changes depending on my location. I'm happy to leave this as it is for now, but will certainly change my setup if a new vulnerability surfaces in any of the tools I use. Love the podcast. I look forward to it every week.

There is no reason to believe that it is not completely safe to leave a Wireguard VPN server running on a firewall, such as OPNsense, listening for incoming connections from a WireGuard client. There's no reason to believe that's a problem ... until there is. Everything we know tells us that this **COULD** flip from "absolutely safe" to "Oh my god!" within a single heartbeat of a skilled hacker who, while studying Wireguard's open source code, notices something no one else has. That's one of the ways these things happen.

Or perhaps the hacker is throwing nonsense packets at Wireguard's listening service port and one of them suddenly crashes the Wireguard server. That's another way this could happen. The specific packet that crashed the server is then examined and the source of the crash is reverse-engineered to create a repeatable working exploit.

But it's every bit as true that **none** of this may ever happen. It's also true that perhaps it can't. The conundrum of security is that "could happen" does not necessarily mean "could happen." Perhaps it really can't. The trouble is, today's systems have become so complex that it's no longer possible for us to be **absolutely and mathematically provably certain** about the behavior of anything above a distressingly low level of complexity. Today, we just can't know. That's one of the things I'm hoping future AI might be able to help us with. My intuition suggests that this is the sort of thing that ought to be right in AI's backyard.

But we don't have that today. What we have today ... is hope. Hope's better than nothing, but "hope" is not enough for me. I fully respect Matt's decision and position. It's one that's shared by 10's of thousands of others. But my network is not the typical residential network. It's both the development and production arms of GRC. So the stakes, for me, are higher. I'm not suggesting that my network is utterly impervious to attack. But it's as utterly impervious as I've been able to make it – ***without exception***. So deliberately exposing a WireGuard process, no matter how safe I hope it is, to the public Internet would be an exception I will not make.

Another listener "**An On**" reminds us why we trust, and should trust, Wireguard's design:

Hi Steve, Regarding the discussion of Wireguard and port knocking on this week's Security Now episode (2024-11-26), I just wanted to let you know that it's not really necessary. With Wireguard, the server will not respond to client connection requests AT ALL unless the client provides a public key that the server knows and trusts. This, in addition to the fact that the protocol is UDP based, means that it's not possible to even know if there is a Wireguard server listening on a specific IP and port unless you already have public key credentials to connect.

While it technically would still be possible to have a bug where this can be bypassed, this is very unlikely because this is the first thing the server checks, so the code surface for bugs is minimal. This technicality would also apply to any port knocking techniques which can have

their own bugs in implementation. Regards, non

"Non" is 100% correct. And this is why Wireguard represents the best of the best today. Is that good enough? Almost certainly. And his point about the possibility that adding port knocking, to introduce an additional layer of pre-Wireguard security, might itself introduce a new vulnerability is also a keen observation. That could happen.

My defense of the use of port knocking is that from an implementation standpoint, unlike anything like Wireguard that necessarily invokes a huge amount of complexity in order to validate a cryptographic certificate, Port Knocking adds an appealingly trivial layer of complexity while providing virtually absolute protection. In other words what might be termed its "security gain" is nearly infinite. And the Port Knocking service is inherently sitting behind the firewall which it's monitoring. So it's much more difficult to see how its failure could do anything other than fail to open a port. And all of this is, of course, what makes the study of "security" so interesting!

One of our listeners, **Richard Craver in Clemmons North Carolina**, pointed me at something that was so interesting it needed sharing. First of all, here's what Richard wrote:

Hi Steve and Leo, I just finished the AGI episode, interesting to ponder. I personally am not a particular fan of AI in general as I see it as crowdsourcing knowledge that may or may not be correct. Science is based on challenging and testing prevailing assumptions and thought. AI, in my humble opinion, discourages critical thinking; but for good or bad, it is here.

Below is a link to Tom Fishburne the Marketoonist with a thought provoking cartoon and short viewpoint message:

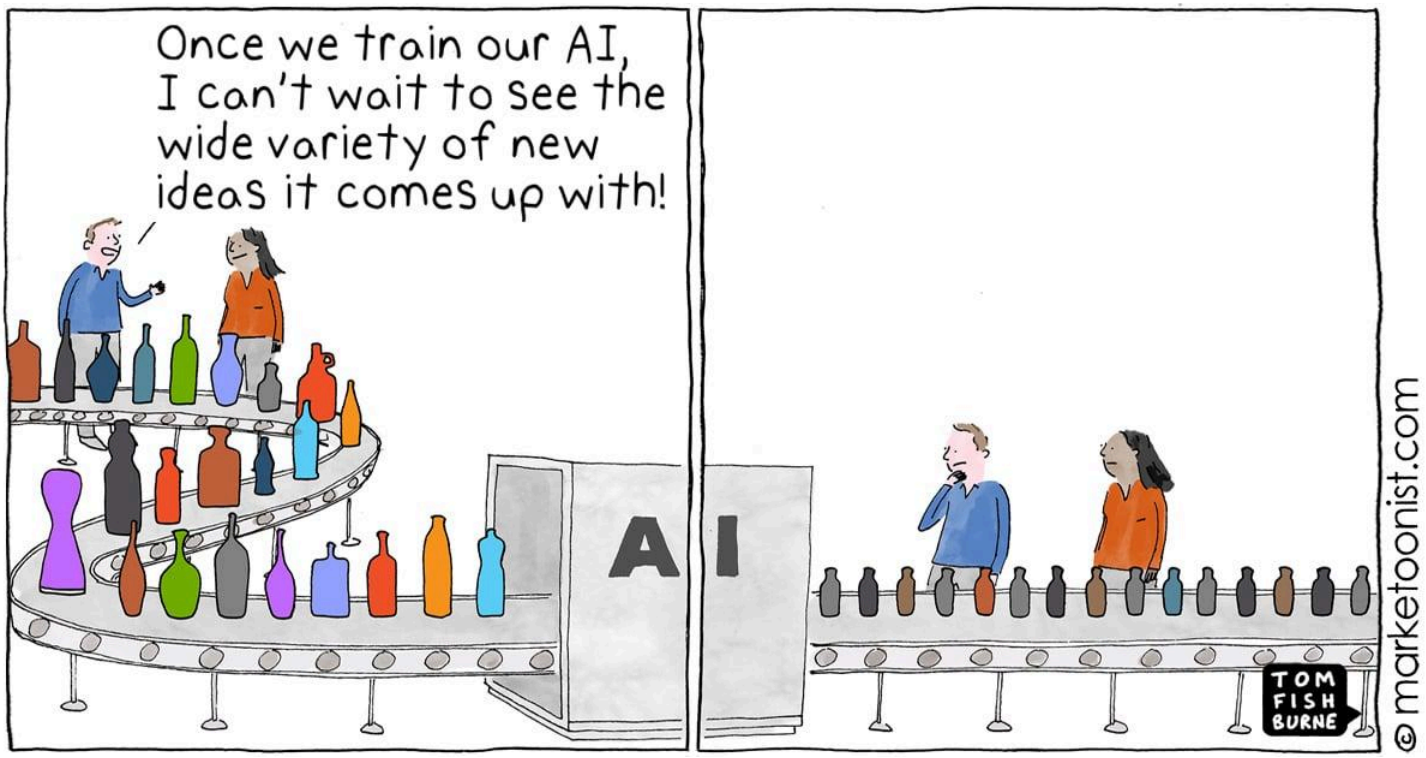

https://marketoonist.com/2024/11/ai-generated-homogeneity.html

Tom Fishburne writes:

It's still early days with AI Generation tools. We're all still learning potentials and limitations. One watch-out is the bias toward homogeneity — the tendency for AI results to look alike. As AI predicts what to generate, the path of least resistance is an averaging of the content in its source material. Ian Whitworth once referred to this as "The Great Same-ning", writing:

"ChatGPT, Jasper and all the rest are powerful conformity machines, giving you the ability to churn out Bible-length material about yourself and your business that's exactly the same as your competitors."

[Tom continues:] A couple months ago, Oxford and Cambridge researchers illustrated the risk of homogeneity in a study of AI Generated content in Nature magazine. The risk increases as AI gets trained not only on human-created content, but on other AI-generated content.

As an example, the researchers studied an AI model trained on images of different breeds of dogs. The source material included a naturally wide variety of dogs (French Bulldogs, Dalmatians, Corgis, Golden Retrievers, etc. "the works"). But when asked to generate an image of a dog, the AI model typically returned the more common dog breeds (Golden Retrievers) and less frequently the rarer breeds (French Bulldogs).

Over time, the cycle reinforces and compounds when future generations of AI models are trained on these outputs. It starts to forget the more obscure dog breeds entirely, soon only creating images of Golden Retrievers.

Eventually, the researchers found, there's "Model Collapse", where the LLM is trained so much on AI-generated Golden Retriever images that the results turn nonsensical and stop looking like dogs at all.

Now, substitute dog breeds for whatever you're trying to create — new products, packaging, advertising, communication, and the risk is that all outputs devolve to look the same.

A related study from the University of Exeter found that AI Generation tools have the potential to "boost individual creativity", but with a "loss of collective novelty." The good news is that this baseline situation creates opportunities for those who can push against this status quo. Homogeneity is ultimately at odds with distinctiveness. As with all tools, it's all in how you use them. You can't break through the clutter by adding to it.

These conclusions feel intuitively correct, and the research cited above supports that intuition. Also, it's certainly true that there's an unrealized danger as the Internet's content becomes more and more AI Generated while our AI models are being continuously trained against the Internet's content. Future historians may wonder... *"what happened to all the French Bulldogs?"*

Greg Haslett has an intriguing problem:

Steve, I have an EdgeRouter and created a IOT network. My problem is I can't reach my ASUS RT-66 to update the firmware that is on the IOT network. Any quick ways to allow temporary access to the ASUS router? My last ditch answer would be to backup the edgerouter and reset to original settings, hopefully find the IP address of the ASUS and update the firmware. Then restore the edgerouter from backup with IOT. Long time listener and met you at the squirrel take in Irvine.

I'm not 100% certain that I completely understood Greg's problem and question. But my first thought is to leave the EdgeRouter alone and temporarily rearrange some wires. Rather than get fancy with reverting the EdgeRouter's configuration to its original simple switch, why not plug the ASUS RT-66 into the LAN where a PC is located and update its firmware. I suppose that if Greg doesn't have a spare old wired Ethernet switch lying around somewhere that might be a problem. But it's also possible to plug the ASUS RT-66 directly, point-to-point into a PC's LAN port. So, if I understand Greg's question, it would appear that being less fancy and going old school might be the solution! 👍

And while we're on the topic of old school solutions that are, in this case, obvious in retrospect, our listener Troy wrote:

> *Steve, Congrats on security now! Hey regarding typing long messages on the iphone, I hope you know you can connect a bluetooth keyboard to your iPhone.*

And this is where I use the expression "Doh!!" I confess, that had **completely** escaped me! And I should have remembered that, because one of my first reactions to the loss of the wonderful physical clicky-button keyboard of my beloved Blackberry when I switched to my first iPhone, was the addition of a little add-on keyboard which did, indeed, link to the phone via Bluetooth. And it worked perfectly!

Earl Rodd in North Canton, Ohio shared some facts about social media age restrictions:

> *The recent book by Jonathan Haidt, "Anxious Generation" has extensive discussion of the age limit issue. The main theme of the book is rather convincing evidence that the dramatic (100-200%) increase in teen mental health problems which corresponds to the introduction of smartphones is in fact CAUSED by use of those phones, and, in particular, social media.*
>
> *Haidt's argument rests on his work as a social psychologist combining knowledge of the vulnerability of early teens due to brain development happening at that time of life with research on how social media is carefully designed to "hook" young adolescents. If Haidt is right (and I think he is), the problem is VERY severe. We make a huge mistake equating our (older adults who grew up before the smartphone era) use of various apps and how we handle it with adolescents during critical brain development years. (Note: My adult children have been telling me this for years – that I can't transfer how I use social media for just the few things I want to the experience of youngsters.*
>
> *The book has an extensive discussion of "what to do." In that section Jonathan discusses some technical ideas (not at the technical depth of Security Now), but also the social factors, like parental role, the problem of peers having more access, and how some methods can be neutralized. The book has references to extensive discussions of both social scientists like Haidt, and technical sources by people who have thought through a lot of ideas. While I share some skepticism of the effectiveness of age verification, I think the combination of laws requiring age verification, more parental awareness, and cooperation between schools and parents can have a very positive impact.*

Interestingly, in our recent discussion I also touched upon a number of the same potential

pitfalls, such as parents being pushed by their own children to make exceptions for them, which is then followed by other kids complaining to their more strict parents that their peers have been given access by their parents, so why can't they have the same? And saying "After all, how bad can it be if 16 year olds are able to have access?" Among other things, my wife is an accomplished therapist and while she rigorously honors the privacy of her clients, she's noted on a number of occasions that many of today's parents appear to be afraid of their own children, whom they appease by giving them anything they want. So how are such parents not going to capitulate to their childrens' demands, especially having previously established that pattern?

And Earl's note is all the more salient since when we talked about this recently the Australian legislation was merely pending. Since then, late on Thursday the 28th on the U.S.'s Thanksgiving Day, that legislation was passed.

Since this promises to be a big deal, Australia may be only the first, and since it has a strong and very interesting technology implementation side, I want to set the stage for our inevitable future discussions by sharing the Associated Press's coverage of the event. The AP wrote:

MELBOURNE, Australia (AP) — A social media ban for children under 16 passed the Australian Parliament in a world-first law. The law will make platforms including TikTok, Facebook, Snapchat, Reddit, X and Instagram liable for fines of up to 50 million Australian dollars (currently around $32.5 million USD) for systemic failures to prevent children younger than 16 from holding accounts.

The Australian Senate passed the bill on Thursday 34 votes to 19. The House of Representatives on Wednesday overwhelmingly approved the legislation by 102 votes to 13. On Friday the House endorsed opposition amendments made in the Senate, making the bill law. Prime Minister Anthony Albanese said the law supported parents concerned by online harms to their children. Albanese told reporters "Platforms now have a social responsibility to ensure that the safety of our kids is a priority for them." The platforms have the next year to work out how they will implement the ban before penalties are enforced.

Meta Platforms, which owns Facebook and Instagram, said the legislation had been "rushed."

Digital Industry Group Inc., an advocate for the platforms in Australia, said questions remain about the law's impact on children, its technical foundations and scope. The Group's managing director said: "The social media ban legislation has been released and passed within a week and, as a result, no one can confidently explain how it will work in practice – the community and platforms are in the dark about what exactly is required of them."

The amendments passed on Friday bolster privacy protections. Platforms would not be allowed to compel users to provide government-issued identity documents including passports or driver's licenses, nor could they demand digital identification through a government system. Critics of the legislation fear that banning young children from social media will impact the privacy of all users who must establish they are older than 16.

While the major parties support the ban, many child welfare and mental health advocates are concerned about unintended consequences. Sen. David Shoebridge, from the minority Greens party, said mental health experts agreed that the ban could dangerously isolate many children who used social media to find support.

Shoebridge told the Senate "This policy will hurt vulnerable young people the most, especially in regional communities and especially the LGBTQI community, by cutting them off."

Exemptions will apply for health and education services including YouTube, Messenger Kids, WhatsApp, Kids Helpline and Google Classroom.

Opposition Senator Maria Kovacic said the bill was not radical but necessary. She told the Senate "The core focus of this legislation is simple: It demands that social media companies take reasonable steps to identify and remove underage users from their platforms. This is a responsibility these companies should have been fulfilling long ago, but for too long they have shirked these responsibilities in favor of profit."

Online safety campaigner Sonya Ryan, whose 15-year-old daughter Carly was murdered by a 50-year-old pedophile who pretended to be a teenager online, described the Senate vote as a "monumental moment in protecting our children from horrendous harms online." She said: "It's too late for my daughter, Carly, and the many other children who have suffered terribly and those who have lost their lives in Australia, but let us stand together on their behalf and embrace this together."

Wayne Holdsworth, whose teenage son Mac took his own life after falling victim to an online sextortion scam, had advocated for the age restriction and took pride in its passage. Holdsworth said "I have always been a proud Australian, but for me subsequent to today's Senate decision, I am bursting with pride."

Christopher Stone, executive director of Suicide Prevention Australia, the governing body for the suicide prevention sector, said the legislation failed to consider positive aspects of social media in supporting young people's mental health and sense of connection. Stone said: "The government is running blindfolded into a brick wall by rushing this legislation. Young Australians deserve evidence-based policies, not decisions made in haste."

The platforms had complained that the law would be unworkable and had urged the Senate to delay the vote until at least June 2025 when a government-commissioned evaluation of age assurance technologies will report on how young children could be excluded.

Meta Platforms, owner of Facebook and Instagram said: "Naturally, we respect the laws decided by the Australian Parliament. However, we are concerned about the process which rushed the legislation through while failing to properly consider the evidence, what industry already does to ensure age-appropriate experiences, and the voices of young people."

Snapchat said it was also concerned by the law and would cooperate with the government regulator, the eSafety Commissioner. Snapchat's statement said: "While there are many unanswered questions about how this law will be implemented in practice, we will engage closely with the Government and the eSafety Commissioner during the 12-month implementation period to help develop an approach that balances privacy, safety and practicality. As always, Snap will comply with any applicable laws and regulations in Australia."

Critics argue the government is attempting to convince parents it is protecting their children ahead of a general election due by May. The government hopes that voters will reward it for responding to parents' concerns about their children's addiction to social media. Some argue the legislation could cause more harm than it prevents.

Criticisms include that the legislation was rushed through Parliament without adequate scrutiny, is ineffective, poses privacy risks for all users, and undermines the authority of parents to make decisions for their children.

Opponents also argue the ban would isolate children, deprive them of the positive aspects of social media, drive them to the dark web, discourage children too young for social media to report harm, and reduce incentives for platforms to improve online safety.

I saw some additional reporting last week which noted that the leaderships of several other countries had congratulated Australia on the passage of this legislation. So, as always, we'll see how it shakes out. It appears that this is going to happen and the means for enforcing these age restrictions should be interesting for us to examine. And since the legislation is due to take effect on November 20th of 2025 next year we're not going to have long to wait.

Dawn appreciates our picture of the week for audio-only listeners:

Hello, Steve and Leo! I've listened to your show for a while now, and, I really enjoy it! I love all things computers, technology, etc, and, there's one thing I can definitely say with 1000% assurance: There will ALWAYS be a need for this podcast, and experts such as yourselves to cover, and explain it all. With the added challenge of putting the cookies on the bottom shelf where the little kids can get them. Which, you are very good at doing!

I wanted to write you an email thanking you for describing the pictures of the week! I have to admit, I got quite a bit of laughs from the one last week, where the little troublesome twosome were finding a way to get upstairs. Even now, as I write this, I'm chuckling! It means a lot to me that you guys describe the pictures of the week, because I'm completely blind. Without your descriptions, I would not be able to get any enjoyment out of them!

Sometimes, I think we do things like this without a second thought, and without knowing the impact that we have, and will have on someone when we do those things. This is one of them.

Please keep the picture descriptions coming! Before you ask, I think one of my favorite pics of the week, was the one that said "Treat your passwords like your underwear." I remember I couldn't stop laughing for a long time after that one, and had to rewind the podcast a couple times just for the laughs! I must admit, I had never heard password safety put that way before! 🙂 😉 Thank you once again for the podcast, and the image descriptions, and, please keep them coming! — Dawn

Dawn, I hope you're listening, thank you for your note, and I can promise that we'll keep the picture-of-the-week descriptions coming.

A Light-Day Away

Our listener **Rob Woodruff** brought this bit of news to my attention. NASA's posting was titled *"NASA's Voyager 1 Resumes Regular Operations After Communications Pause"*. I'm sharing it today since it contained a bunch of interesting and amazing science and engineering information:

NASA's Voyager 1 has resumed regular operations 1 following a pause in communication last month. The probe had unexpectedly turned off its primary radio transmitter, called an X-band transmitter, and turned on the much weaker S-band transmitter. Due to the spacecraft's distance from Earth — about 15.4 billion miles (24.9 billion kilometers) — this switch prevented the mission team from downloading science data and information about the spacecraft's engineering status.

Earlier this month, the team reactivated the X-band transmitter and then resumed collecting data the week of Nov. 18 from the four operating science instruments. Now engineers are completing a few remaining tasks to return Voyager 1 to the state it was in before the issue arose, such as resetting the system that synchronizes its three onboard computers.

The X-band transmitter had been shut off by the spacecraft's fault protection system when engineers activated a heater on the spacecraft. Historically, if the fault protection system sensed that the probe had too little power available, it would automatically turn off systems not essential for keeping the spacecraft flying in order to keep power flowing to the critical systems. But the probes have already turned off all nonessential systems except for the science instruments. So the fault protection system turned off the X-band transmitter and turned on the S-band transmitter, which uses less power.

The mission is working with extremely small power margins on both Voyager probes. Powered by heat from decaying plutonium that is converted into electricity, the spacecraft lose about 4 watts of power each year. About five years ago — some 41 years after the Voyager spacecraft launched — the team began turning off any remaining systems not critical to keeping the probes flying, including heaters for some of the science instruments. To the mission team's surprise, all of those instruments continued to operate despite reaching temperatures lower than what they had been tested at.

The team has computer models designed to predict how much power various systems, such as heaters and instruments, are expected to use. But a variety of factors contribute to uncertainty in those models, including the age of the components and the fact that hardware doesn't always behave as expected.

With power levels being measured to fractions of a watt, the team also adjusted how both probes monitor voltage. But earlier this year, the declining power supply required the team to turn off a science instrument on Voyager 2. The mission shut off multiple instruments on Voyager 1 in 1990 to conserve energy, but those instruments were no longer in use after the probe flew past Jupiter and Saturn. Of the 10 science instruments on each spacecraft, four are now being used to study the particles, plasma, and magnetic fields in interstellar space.

Voyagers 1 and 2 have been flying for more than 47 years and are the only two spacecraft to operate in interstellar space. Their advanced age has meant an increase in the frequency and complexity of technical issues and new challenges for the mission engineering team.

That article said: *"The X-band transmitter had been shut off by the spacecraft's fault protection system when engineers activated a heater on the spacecraft."* What it didn't tell us is **why** the JPL engineers turned on that heater. And there's even more fascinating information about that.

Our listener **Jeff Root in San Diego** supplied the link to a story in The Register, titled *"'Best job at JPL': What it's like to be an engineer on the Voyager project."* This was posted two days later on the U.S.'s Thanksgiving Thursday. And it, too, is chock full of interesting science and engineering insight. The Register wrote:

The Voyager probes have entered a new phase of operations. As recent events have shown, keeping the venerable spacecraft running is challenging as the end of their mission nears.

As with much of the Voyager team nowadays, Kareem Badaruddin, a 30-year veteran of NASA's Jet Propulsion Laboratory (JPL), divides his time between the twin Voyager spacecraft and other flight projects. He describes himself as a supervisor of chief engineers but leaped at the chance to fill the role on the Voyager project.

Suzanne Dodd, JPL Director for the Interplanetary Network Directorate, is the Project Manager for the Voyager Interstellar Mission.

Badaruddin told The Register: "She knew that the project was sort of entering a new phase where there was likely to be a lot of technical problems – and so chief engineers, that's what they do. They solve problems for different flight projects."

Dodd needed that support for Voyager. Badaruddin would typically have found someone from his group, but he said: "I was just so excited about Voyager, I said, you know, look no further, right? I'm the person for the job. I'm your engineer. You know, please pick me."

So Badaruddin has spent the past two years on the Voyager project. After decades of relatively routine operation, following plans laid out earlier in the mission when the team was much larger, the twin Voyager spacecraft have begun presenting more technical challenges to overcome as the vehicles age and power dwindles.

The latest problem occurred when engineers warmed up part of the spacecraft, hoping that some degraded circuits might be "healed" by an annealing process. Badaruddin explained that "There's these junction field effect transistors (JFETs) in a particular circuit that have become degraded through radiation. We don't have much protection from radiation in an interstellar medium because we're outside the heliosphere where a lot of that stuff gets blocked. So we've got this degradation in these electronic parts, and it's been proven that they can heal themselves if you get them warm enough, long enough. And so we knew we had some power margin, and we were hopeful that we had enough power margin to operate this heater ... and as it turned out, we didn't. It was a risk we took to try to ameliorate a problem that we have with our electronics. So now the problem is still there, and we realize that we can't solve it this way. And so we're going to have to come up with another creative solution."

The problem was that more power was demanded than the system could supply. A voltage regulator might have smoothed things out, but the Voyagers no longer have that luxury. Instead, engineers took a calculated risk and ran afoul of the then-innovative software onboard the spacecraft. The under-voltage routine of the fault protection software shuts down loads on the power supply, but since the Voyager team had already shut down anything that is not

essential, there isn't much left.

Badaruddin explained: "So the under-voltage response doesn't do much except turn off the X-band transmitter and turn on the S-band transmitter. And that's because the S-band transmitter uses less power, making it the last safety net to save you."

And save the mission it did. While the S-band is great for operations near Earth, such as the Moon, it is almost useless at the distance of the Voyager spacecraft. However, by detecting the faint carrier signal of the S-band transmission, the team was able to pinpoint that the problem had been the act of turning on the heater, even without X-band telemetry from the spacecraft.

The challenge for engineers isn't just the time it takes to get a command to the Voyagers and receive a response, but also checking and rechecking every command that gets sent to the spacecraft. The waiting is apparently not as frustrating as we might think. Badaruddin said: "This is the rhythm we work in; we've grown accustomed to it … it used to be a very small time delay and it's gradually grown longer and longer through the years."

With duplicate physical hardware long gone, the team now works with an array of simulators. Badaruddin said: "We have a very clear understanding of the hardware. We know exactly what the circuitry is, what the computers are, and where the software runs."

And as for the software? ... It's complicated. There have been so many tweaks and changes over the years that working out the exact revision of every part of Voyager's code has become tricky. Badaruddin said: "It's usually easier to just get a memory readout from the spacecraft to find out what's on there."

The challenge for the Voyager team is that the spacecraft are nearing the half-century mark, as is the documentation. <quote> "We have documents that were type-written in the 70s that describe the software, but there are revisions ... and so building the simulators, we feel really good about the hardware ... but we feel a little less good about understanding exactly what each instruction does." The latest bit of recoding occurred with the failure of one of Voyager 1's integrated circuits, which manifested itself as meaningless data last year.

We talked about this on the podcast at the time. Badaruddin reminds us:

"The basic problem was figuring out what was wrong with no information. We could see a carrier signal; we knew we were transmitting in the X-band ... we knew we could command the spacecraft because we could tweak that signal slightly with commands. So we knew the spacecraft was listening to us, and we knew the spacecraft was pointing at Earth because otherwise, we wouldn't get a signal at all."

The engineers went further down the fault tree, and eventually managed to get a minimum program to the spacecraft to give a memory readout. That readout could be compared to one retrieved when the spacecraft was healthy. 256 words were corrupted, indicating a specific integrated circuit. Code was then written to relocate instructions around that failed area.

"The problem there is the code was very compact. There was no free space that we could take advantage of. So we had to sacrifice something."

That something was one of Voyager 1's higher data rate modes, used during planetary flybys.

Okay, so now back to the present:

The current challenge involves dealing with the probes' thrusters. Silicon from bladders inside the fuel tanks has begun to leach into hydrazine propellant. Since silicon doesn't ignite like hydrazine, a tiny amount gets deposited in the thrusters and slowly builds up in the thruster capillaries. Badaruddin uses the analogy of clogging arteries. Eventually, the blockage will prevent the spacecraft from firing its thrusters to keep it pointed at Earth.

However, the pitch and yaw thrusters, each of which have three branches, are clogging at different rates. The current software works on the basis that branch 1, 2, or 3 will be used. But could it be operated in mixed mode, where branch 2 is used for the pitch thruster, but branch 3 is used for the yaw?

Badaruddin notes: "So that's a creative solution. It would be very complicated ... this would be another software modification in interstellar space." And getting it right the first time is not just nice to have; it's almost essential. By the time the results of a command come back from the Voyager spacecraft, it might be impossible to deal with the fallout of a failure.

The Voyager spacecraft are unlikely to survive another decade. The power will eventually dwindle to the point where operations will be impossible. High data rates, which is to say 1.4 kilobits per second, will only be supported by the current Deep Space Network (DSN) until 2027 or 2028. After that, some more creativity will be needed to operate Voyager 1's digital tape recorder.

Badaruddin speculates that shutting off another heater (the Bay One heater) used for the computers would free up power for the recorder, according to the thermal model, but it'll be a delicate balancing act. And, of course, the recent annealing attempt demonstrated the limitations of modeling and simulations on Earth.

So, does Badaruddin have a favorite out of the two spacecraft? He replies: "Well, Voyager 2 is the one that's been flying the longest, and Voyager 1 is the one that's furthest from Earth. So they both have a claim to fame." To use another analogy: "They're essentially twins ... they're basically the same person, but they live different lives, and they have different medical problems and different experiences."

Badaruddin hopes to stick with the mission until the final transmission from the spacecraft.

"I love Voyager. I love this work. I love what I'm doing. It's so cool. It just feels like I've got the best job at JPL."

I just checked on the Voyager 1 mission status, which is what gave me the title for today's podcast. That intrepid little spacecraft is now so far away that light (and radio signals) take more than 23 hours to travel in each direction. It's nearly one entire light-day distant. Yet Voyager 1 is managing to keep itself pointed at our Earth across all that distance and we still have working bi-directional communication with it. This entire endeavor has been an astonishing example of incredible engineering. The original design was flexible enough that – and software controlled enough – that even though it was designed in the 1970's and launched on September 5th, 1977 – all well before the Internet and all of the technology we now take for granted, this machine has endured and has exceeded everyone's expectations many times over.

This story does make one principle absolutely clear: No pure hardware solution could have ever done this. No pure hardware solution would still be alive, functioning and communicating after 47 years of space flight. Nor even could any fixed firmware hybrid hardware / software solution.

The reason is that **none** of what has transpired since Voyager 1's original mission was redefined and extended, after it continued to perform so brilliantly, could have been anticipated by NASA's brilliant engineers in the mid 70's. The sole key to Voyager 1's success today is that to an extremely large degree the original designers of the spacecraft put the machine's hardware under software control. The reason they did that back in the 70's was different from the reason they are now glad they did that. They created a deeply software-based control system back then because software doesn't weigh anything and the spacecraft didn't have an ounce to spare. So the engineers of the 70's put their faith in software and that faith, and the inherent dynamic redesign flexibility it enabled, has given the spacecraft a far longer life than it could have ever otherwise enjoyed.

All of that said, yesterday's and today's software is ultimately at the mercy of hardware. If the attitude control systems' capillaries become clogged with leached and deposited silicon, the spacecraft's ability to maneuver and keep itself pointing at the Earth will eventually be lost. At some point in the not too distant future it will still be alive but we'll have lost contact with one another.

What an amazing accomplishment.