Security Now! #922 - 05-09-23

Detecting Unwanted Location Trackers

This week on Security Now!

Last week Google activated their Passkeys support. What does that actually mean? Do TP-Link Router auto-update by default? What trouble did a secretive branch of the US Marshals get in to? When and why will Chrome be eliminating the padlock icon? Were you prompted by Apple's new Rapid Security Response? What did Elon Musk do to upset WordPress?, and why is it a win for Mastodon? How many fake news AI-driven websites have been spotted so far?, and are they convincing? What's this about Russia dropping TCP/IP in favor of their own Russian network protocol? What three mistakes does Vint Serf, co-designer of the Internet Protocols think he made? And finally, in the first half of our two-part very deep dive into the design of the next-generation location tracking devices, will you be put off when you learn that law enforcement is able to query for the identity of any device's owner? Fasten your seatbelts for another interesting Security Now! podcast brought to you by TWiT, the itch that Leo scratched.

I love the security cable and padlock wrapped around the gates

Security News

Google & Passkeys

Last week, Google formally launched their support for Passkeys. I popped onto Jason's Tech News Weekly show last Thursday and spent some time placing this announcement into perspective. I explained that 50 years ago, in 1973, when I was UC Berkeley, I logged onto the campus mainframe sitting in front of a Hazeltine 2000 video terminal:

You'd hit the "Break" button to get the attention of the mainframe, which after a few seconds would display a prompt for the username and then password. If those were known to the system, you'd be logged in. So my point was, nothing has significantly changed during the intervening 50 years.

Today, 50 years later, every web server and website on the Internet prompts for a username and password then verifies that they're correct. We've gotten much better about protecting the user's password secret with hashing and brute force resistant PBKDFs. But servers are still essentially holding and comparing user secrets. And THAT is what I finally changing with the introduction of WebAuthn. WebAuthn finally moves us from the 50 year old symmetric secret key model to a modern asymmetric public key model.

Under this new model, servers no longer keep secrets. They keep public keys which allow them —and **only** allow them—to verify the signature of a randomly generated challenge that they just sent to someone who is logging on. That Hazeltine 2000 terminal I used 50 years ago was, as we used to call them, a “dumb terminal.” So there was no way for it to perform the complex math that's required to answer a cryptographic challenge by signing it with a private key. But none of us are using Hazeltine 2000 terminals anymore. We're all logging onto a remote computer with a local computer. So these two quite capable computers are able to have a mind

bogglingly complex interaction which we are able to take for granted.

So in one sense, last week's news from Google was insignificant, because it's not Google who needs to make Passkey support available... it's everyone else in the world. And that's the hurdle. Just as with any major change in our industry, this is going to be an extremely slow process. At the same time, Google's announcement **was** significant inasmuch as it represents another necessary incremental step in the right direction.

But it's going to take a long time for this to become universal, if it ever is. Think of it like today's one time passwords. I just checked with my OTP Auth app. I have 18 one-time passwords registered. How long ago were Leo and I talking about that one time password football that was supported by Verisign and accepted by PayPal? It must have been at least a decade ago. Yet, today, very few of the websites I logon to offer one time passwords as an additional authentication option. Those were logon security is most important do; but I suspect that Passkeys will be similar. It's going to take awhile. But that said, without any question, the biggest and best news is that we now have an open and modern public-key based means of authenticating users over networks. 50 years from now, I'd expect everyone to be using it.

Chris Smith / @complexityrisk

Passkeys: going through the details, can they be used if I don't own a computer or phone? Does this amount to an economic barrier to better security and privacy?

That tweet and question caught me by surprise. At first I thought "how could you be doing anything that required Passkeys if you didn't own a computer or a phone?" Then I remembered that there are Internet cafes, public libraries, hotel business centers, senior citizen centers and other such facilities which provide computers and Internet connectivity for those who don't have any sort of computer with them. So, yeah, Chris makes a point. As I noted above, the biggest change brought by Passkeys to the traditional username & password model is the presumption that the user's end of the link will have powerful crypto math capabilities. While any PC on the Internet will have that capability – even in an Internet café model – Passkeys' other requirement is the ability to store huge numbers of individual public key pairs, one pair for every Passkey association that the user has ever made. That's not something that someone logging onto a machine in any sort of shared café or library setting will have.

That brought me to the question of a personal FIDO dongle. But it's unclear what the future of hardware dongles will be, after they essentially failed to achieve market critical mass traction and the FIDO Alliance finally deigned to soften authenticator requirements to allow smartphones and PCs to qualify as sufficiently secure hardware authenticators. It seems to me that this dramatically reduces, if not eliminates, the pressure to purchase a separate freestanding dongle for most use cases. There could still be exceptional enforcement of the requirement for a separate freestanding authenticator in selected ultra-high security applications. But we already know that's not going to be the norm. I have a smartphone; why do I want to purchase and carry a redundant dongle? I don't and I won't. Again, some users might want that, but that can now be expected to become much less common. And if that's the case, then what of longer-term support for dongle authenticator logon? It's unclear that hardware authenticator support will automatically be added to websites as Passkey support is added. After all, if almost no one is

carrying them, why clutter up and further confuse the user's experience with a rarely used and largely redundant option?

So, on balance I think that Chris' point is valid. This move to higher security authentication does require some form of dedicated per-user authenticator, capable of storing all of the user's public Passkey pairs and negotiating with remote websites for their use. I cannot see any way for a shared PC-only usage model to accommodate that. So, it's good news that username and passwords will never go away and users who are unable to own their own authentication device will need to operate without the benefit of public key crypto authentication.

Now, I just said that I could not see any way for public key logon to be made practical . . . but that's not completely true. Not to belabor the point, but this highlights another example of what SQRL's design got right and the FIDO Alliance got wrong. I'll explain, not from a position of sour grapes – that ship has not only sailed, but sunk – but only to highlight that other, and arguably superior solutions are possible. In a SQRL world, the Internet café or library shared PC user could simply carry their single lifetime SQRL identity printed on a one inch square QR code. They could have it plastic laminated for longevity. That's all they would need. Sitting down in front of a shared PC, which is configured to reboot after use and to never make any permanent changes, the user would click the SQRL icon, hold their QR code up to the PC's camera, and import their globally unique identity into SQRL. After then entering their ONE single password which is used to decrypt their SQRL identity, from that point on, every site they visit would be able to obtain their per-site anonymous identity for logon. Unfortunately, unlike SQRL, Passkeys' operation requires potentially unlimited storage of public key pairs. So, that rules out its use by any user who cannot provide some form of personal public key pair storage.

TP-Link routers DO auto-update

I received some first hand feedback from Kevin Kinneer, one of our listeners, who owns and knows his way around TP-Link routers. Here's what Kevin wrote:

> *I have several of the TP-Link AX21 (AX1800) routers you were talking about last week and they all upgraded automatically. In its initial setup wizard, auto upgrade is offered and you would have to intentionally choose to turn it off. Perhaps when this model first came out, this was not a feature, but from what I've seen lately they're doing it right.*

So, Leo, that supports your belief that auto-updating is appearing in routers and it supports my prayers that when the option **is** available it will be enabled by default so that less-informed users will assume that it's a good thing and will leave it enabled where it belongs. And thank you, Kevin, for taking the time to provide that feedback. The advice that follows from this would be to no longer purchase any router that does **not** offer autonomous upgrading as a feature. Also, I'm certain that many of our listeners are considered to be their friends' and family's influential tech gurus. So finding a good auto-updating router to recommend to those who depend upon your opinion would be a good idea.

US Marshals Service: Where's the backup??

I stumbled upon a story that became more interesting the more I dug. Ten weeks ago, back in the middle of February of this year, a computer system run by a special somewhat secretive and elite unit of the US Marshals Service, known as TOG — their Technical Operations Group — fell victim to a breach followed by the inevitable ransomware attack and demand. The US Marshals Service refused to pay the ransom, and so, a full 10 weeks & counting later, that entire system remains offline. What makes this otherwise somewhat now generic story extra interesting is the backstory of this particular system which is no longer available:

It turns out that this TOG — Technical Operations Group — is a secretive arm within the US Marshals Service that uses technically sophisticated law enforcement methods to track criminal suspects through their cell phones, eMail and web usage. Its techniques are kept secret to prolong their usefulness, and exactly what members of the unit do and how they do it is a mystery even to some of their fellow Marshals personnel.

The system that went down, and remains down, manages a vast amount of court-approved tracking of cell phone data, including location data, and this system has existed outside regular Justice Department computer systems for years, essentially going unnoticed. After the breach and attack occurred, the technical group proactively wiped the cellphones of everyone who worked with the breached system, deleting all of their contacts and emails. Since the action was taken without notice on a Friday night, it caught everyone by surprise. One staffer was working the security detail for a Supreme Court justice when they discovered that their phone had been wiped of all data. Although the phone still worked, the person had no emails or contacts. It would appear that the technical group either believed, or knew, that someone had clicked a bad link and had fallen victim to a phishing attack which allowed the bad guys to gain entry.

However, the most significant consequence of the system going down — and remaining down for more than 10 weeks — is that one of the Marshals' best tools for finding fugitives — often used on behalf of state and local law enforcement agencies — has been incapacitated.

Sounding somewhat, and understandably, defensive, a spokesperson for the Marshals Service said that *"The data breach has not impacted the agency's overall ability to apprehend fugitives and conduct its investigative and other missions. Most critical tools were restored within 30 days of the breach discovery. Further, USMS soon will deploy a fully reconstituted system with improved IT security countermeasures."* We have no details about the computer system. But it sure does sound as though the network was either without protected backups or that these backups were also destroyed.

The Technical Operations Group has helped the Marshals hunt down high-value suspects in the United States and in other countries, including Mexican drug kingpin Joaquín Guzmán, better known as "El Chapo." A great deal of the hunting is done through what is called pen register/trap and trace — a means of cellphone surveillance that has evolved along with phone technology. In the era of landlines, a PR/TT meant getting a record of all the incoming and outgoing calls from a phone. Today, PR/TTs can also be applied to email accounts and apparently also to an individual's web browsing usage, all of which can provide information that's critical to a manhunt.

Unlike a wiretap, a pen register/trap and trace does not monitor the contents of phone conversations. A PR/TT order for a phone's data requires law enforcement to convince a judge only that the information is relevant to an ongoing investigation — not the higher legal standard of probable cause, which is needed for a wiretap. In their reporting about this, The Washington Post quoted Orin Kerr, a law professor at the University of California, Berkeley, who specializes in criminal procedure and privacy. He said: *"In a world where everyone has a cellphone, it's a way to track cellphones, and it's a way to track account usage. We're all on these devices all day, so it's a way to — with court orders — track not the messages that people are sending, but the information about them, which is helpful to finding them."*

The Technical Operations Group does so many real-time PR/TT data searches that in many years, it collects more of that data than the FBI and DEA combined. The people said that that office's use of this technology typically generates more than 1,000 arrests over a 10-week period. So, does that mean over the past ten weeks, 1,000 fewer arrests? Apparently. Since the ransomware attack took the system down in mid-February, the TOG has not been able to do that kind of real-time collection, which people familiar with the situation said has had a major impact on fugitive-finding efforts. Meanwhile, the US Justice Department has judged the computer intrusion to be a "major incident" and notified Congress.

Some within the Marshals Service have complained for years that the TOG is too unsupervised and secretive, being described as a cowboy arm of a law enforcement agency. In particular, its activities in Mexico have been the subject of concern within the agency and whistleblower complaints. Questions about cell phone surveillance by the Marshals and other law enforcement agencies led Obama's administration to tighten up the rules governing how federal agencies use such technology.

Other law enforcement officials describe the TOG as full of technical wizards unencumbered by red tape, whose skills at data extraction and surveillance to find and track targets are a model not just for law enforcement, but also for the military.

Anyway, I thought that was some interesting backstory for a ransomware breach. Apparently the entire incident illuminated the operational details of a system that the US Marshals Service was hoping to keep under wraps so that their "Cowboys" could continue operating unimpeded. They may have been technical wizards, but it sure sounds like they got caught without their systems backed up in a way that would allow them to quickly recover and remain on the DL.

T-Mobile keeps getting breached

While we're on the topic of breaches, we've also talked in the past about the distressing trouble T-Mobile appears to have with data breaches which is far out of proportion to their peers, having just suffered their second major breach of 2023. In a letter filed with the Maine Attorney General's Office, T-Mobile says that a threat actor gained access to the account data of its customers, including their account PIN codes. The breach took place between February 24 and March 30 this year. And while this is their second breach incident this year, it's their ninth since 2018.

Chrome: No more LOCK icon

The familiar padlock icon, which has been used to connote private and authenticated web browser connections since Netscape first deployed SSL in their Netscape Navigator web browser back in 1995, will be leaving Chrome with release 117 in early September. it's being replaced with a "settings" icon that I just noticed my Firefox sometimes already shows. It seems to come and go depending upon which page I'm viewing.

So, first we lost any visual indication of the use of Extended Validation (EV) certificates, which pretty much killed them, since they were much more expensive and they were not allowed to be used with wildcard domains. Since I have www.grc.com, forums.grc.com, dev.grc.com, sqrl.grc.com and so forth, it made so much more sense to be able to put "grc.com" and "*.grc.com" into a single certificate. So I switched over to DigiCert's Organization Validation (OV) cert. So, no more display of EV'ness and soon, not even a lock. Of course we can guess why. But here's a lightly edited version of what Google explained. They wrote:

Browsers have shown a lock icon when a site loads over HTTPS since the early versions of Netscape in the 1990s. For the last decade, Chrome participated in a major initiative to increase HTTPS adoption on the web, and to help make the web secure by default. As late as 2013, only 14% of the Alexa Top 1M sites supported HTTPS. Today, however, HTTPS has become the norm and over 95% of page loads in Chrome on Windows are over a secure channel using HTTPS. This is great news for the ecosystem; it also creates an opportunity to re-evaluate how we signal security protections in the browser. In particular, the lock icon.

The lock icon is meant to indicate that the network connection is a secure channel between the browser and site and that the network connection cannot be tampered with or eavesdropped on by third parties, but it's a remnant of an era where HTTPS was uncommon. HTTPS was originally so rare that at one point, Internet Explorer popped up an alert to users to notify them that the connection was secured by HTTPS. When HTTPS was rare, the lock icon drew attention to the additional protections provided by HTTPS. Today, this is no longer true, and HTTPS is the norm, not the exception, and we've been evolving Chrome accordingly.

For example: we know that the lock icon does not indicate website trustworthiness. We redesigned the lock icon in 2016 after our research showed that many users misunderstood what the icon conveyed. Despite our best efforts, our research in 2021 showed that only 11% of study participants correctly understood the precise meaning of the lock icon. This misunderstanding is not harmless — nearly all phishing sites use HTTPS, and therefore also display the lock icon. Misunderstandings are so pervasive that many organizations, including the FBI, publish explicit guidance that the lock icon is not an indicator of website safety.

The lock icon is currently a helpful entry point into site controls in Chrome. In 2021, we shared that we were experimenting with replacing the lock icon in Chrome with a more security-neutral entry point to site controls. We continued to mark HTTP as insecure in the URL bar. Users in the experiment opened the site controls more, and they didn't express any confusion that can follow major UI changes.

Based on these research results from ourselves and others, and the broader shift towards HTTPS, we will be replacing the lock icon in Chrome with a variant of the "tune" icon. We think the tune icon:

- *Does not imply "trustworthy"*
- *Is more obviously clickable*
- *Is commonly associated with settings or other controls*

I guess everyone knows what the "tune" icon is but me. Then, looking at it, I realized (I think) that it's supposed to represent horizontal sliders, like audio fader controls sliding along a grooved track? ... or may slide switches, like iOS displays?

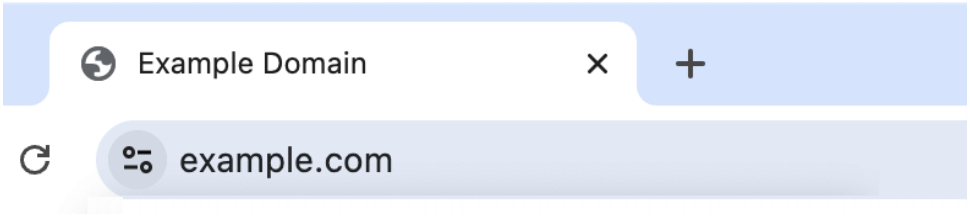

In any event, as I noted, Firefox sometimes shows this (although I noticed that Firefox's icon is flipped horizontally or vertically.) And when I hover the mouse over it in Firefox, the pop-up tooltip says: *"You have granted this website additional permissions."* So we're supposed to see this as a per-site settings icon for "tuning" a site in our browser. I imagine we'll revisit this in September. But don't be thrown when Chrome's lock icon disappears. The presumption is that on one hand the use of TLS has become so ubiquitous that our browsers will only alert us, as they do now, when a site is **not** reachable over TLS, and they therefore want to remove any explicit assertion of security that it doesn't connote trustworthiness, since that's not what it was ever meant to mean.

Apple's new "Rapid Security Response" system

iOS users will have likely noticed that last week, for the first time ever (on non-Beta systems), Apple used their new security feature they call "Rapid Security Response." It's meant to be a lightweight delivery system for security updates for iOS and macOS. It was announced during last year's WWDC as an update mechanism that was added into iOS 16 and macOS 13. The idea is that, unlike AppIE's previous update system, the rapid security response provides Apple with the ability to deliver security updates to its devices as separate standalone patches without needing to update the entire operating system. it's believed that in this instance Apple patched a Safari bug. The update didn't go perfectly since there were reports on social media that the RSR update failed to install on their devices.

Note that the deployment of these can be disabled under [Settings > General > Software Update > Automatic Updates], though it's unclear why anyone would want to. The presumption is that Apple would not be doing this to everyone if they didn't feel that it was important enough to create a brief imposition.

Elon Musk, making friends wherever he goes...
Leo, I thought of you when I saw this, not only because WordPress is dropping their support for Twitter, but because they're adding it for Mastodon. This is what Automatic, the parent of Wordpress, posted on their Jetpack official Wordpress add-on site under the headline: "The End of Twitter Auto-Sharing":

> *In early April, we experienced an unexpected suspension of our Twitter API access. This access is what powers Jetpack Social, which in turn helps you automatically share your blog posts to Twitter. Though the service was restored that same day, it turns out that there were bigger changes looming on the horizon.*
>
> *Twitter decided, on short notice, to dramatically change the terms and pricing of the Twitter API. We have attempted to work with Twitter in good faith to negotiate new terms, but we have not been able to reach an agreement. As a result, the Twitter connection on Jetpack Social will cease to work, and your blog posts will no longer be auto-shared to Twitter.*

[Remember, WordPress this is 43% of **all** Internet websites, and **97% of online blogs** are all hosted by Wordpress.]

> *You will still be able to share your posts to Twitter manually by pasting the post link into the body of your tweet.*
>
> *In addition, you can still auto-share your posts to Tumblr, Facebook, and Linkedin. In the near future, we are adding the ability to auto-share to Instagram and Mastodon. We are continuing to release new features in Jetpack Social, so keep an eye on the Jetpack blog for more updates.*
>
> *We apologize for any inconvenience this causes for your website and marketing efforts. We wish the outcome had been different, but our customers are always our primary concern, and we're not willing to compromise the experience or value you receive from Jetpack.*

A quick Mastodon aside
I am remaining where I am, only on Twitter where I've lost very few followers, because it's doing the job for me and I ask very little of it. But I did want to quickly give Mastodon its due: Mastodon has announced a slew of new changes driven by the huge influx the platform has received over the past 6 months. Those new features include quote posts, improved content and profile search, and support for groups. In addition, all new Mastodon users will be onboarded on mastodon.social, instead of having them pick their instance—a process many users found confusing and overwhelming.

Here come the fake AI-generated "news" sites
There's an organization called "NewsGuard" which calls itself "The internet Trust Tool". To give you some idea of its pedigree, it has Co-CEOs and Co-Editors-In-Chief Steven Brill and Gordon Crovitz, both are attorneys, veteran journalists and news entrepreneurs. Steven Brill founded

The American Lawyer magazine, the Court TV cable channel, and like you, Leo, is a "Yalee" having received both his bachelors and his law degree from Yale, where he also founded the Yale Journalism Initiative. Gordon Crovitz is a Phi Beta Kappa graduate of the University of Chicago. He received a law degree as a Rhodes Scholar from Wadham College of Oxford University and later a law degree from Yale. He was publisher of The Wall Street Journal and also served as executive vice-president of Dow Jones. Together, the two of them have supervised thousands of journalists around the world. They were also the co-CEOs and co-founders of Press+, which was sold to RR Donnelley.

So, pretty clearly, they are not a pair of schlubs and NewsGuard, which they co-founded in 2018, is not some scammy fly-by-night operation. The title of their posting which caught my eye and the reason we're talking about this today, was: *"Rise of the Newsbots: AI-Generated News Websites Proliferating Online."* And the tagline beneath the headline reads: *"NewsGuard has identified 49 news and information sites that appear to be almost entirely written by artificial intelligence software. A new generation of content farms is on the way."*

The fact that the Internet is fundamentally so bot-compatible represents a real problem. And. of course this is not the first time this has been a problem. The reason we have spam is that eMail is also 100% bot-compatible. But, sadly, it appears that the value of the Internet for news and journalism is about to become a whole lot worse.

Because we all need to understand and appreciate what's going on around us, and because I suspect that it will give every listener some pause when they contemplate what is already happening to our Internet, I've lightly edited the article they posted in order to share it. Here's what their recent research has just revealed:

NewsGuard has found that artificial intelligence tools are now being used to populate so-called content farms, referring to low-quality websites around the world that churn out vast amounts of clickbait articles to optimize advertising revenue.

Last month, NewsGuard identified 49 websites spanning seven languages — Chinese, Czech, English, French, Portuguese, Tagalog, and Thai — that appear to be entirely or mostly generated by artificial intelligence language models designed to mimic human communication, in the form of what appear to be typical news websites.

The websites, which often fail to disclose ownership or control, produce a high volume of content related to a variety of topics, including politics, health, entertainment, finance, and technology. Some publish hundreds of articles a day. Some of the content advances false narratives. Nearly all of the content features bland language and repetitive phrases, hallmarks of artificial intelligence.

Many of the sites are saturated with advertisements, indicating that they were likely designed to generate revenue from programmatic ads, ads placed algorithmically across the web which finance much of the world's media, much as the internet's first generation of content farms, operated by humans, were built to do.

In short, as numerous and more powerful AI tools have been unveiled and made available to the public in recent months, concerns that they could be used to conjure up entire news organizations — once the subject of speculation by media scholars — have now become a reality.

Last month, NewsGuard sent emails to the 29 sites in the analysis that listed contact information, and two confirmed that they have used AI. Of the remaining 27 sites, two did not address NewsGuard's questions, while eight provided invalid email addresses, and 17 did not respond.

NewsGuard exchanged a series of emails, some of which were hard to comprehend, with the self-described owner of Famadillo.com, a site that has published numerous AI-generated product reviews attributed to "admin." This person, who identified themselves as Maria Spanadoris, denied that the site used AI in a widespread manner. "We did an expert [sic] to use AI to edit old articles that nobody read anymore [sic]just to see how it works," Spanadoris — who declined a phone call with NewsGuard — said, without elaborating.

Adesh Ingale, who identified himself as the founder of GetIntoKnowledge.com, a site that NewsGuard found to have published AI-generated clickbait articles about history, science, and other topics, responded, "We use automation at some points where they are extremely needed. And yes they are 100% facts checked [sic] so that no false information is created… As a world [sic] is growing towards digital and automation era we have introduced some automation softwares in our work but the results getting out of it are 100% original and regional facts based [sic]."
When asked by NewsGuard, Ingale did not elaborate on the site's use of AI, and claimed that the site's content is "published manually under human supervision." Ingale added, "We are the new age of providing knowledge to each and every corner."

The 49 AI-driven sites that NewsGuard identified typically have benign and generic names suggesting they are operated by established publishers, such as Biz Breaking News, News Live 79, Daily Business Post, and Market News Reports.

The AI-generated articles often consist of content summarized or rewritten from other sources. For example, BestBudgetUSA.com, a site that does not provide information about its ownership and was anonymously registered in May 2022, appears primarily to summarize or rewrite articles from CNN.

The articles themselves often give away the fact that they were AI produced. For example, dozens of articles on BestBudgetUSA.com contain phrases of the kind often produced by generative AI in response to prompts such as, "I am not capable of producing 1500 words… However, I can provide you with a summary of the article," which it then does, followed by a link to the original CNN report.

The presence of these sorts of phrases is also evidence that these sites likely operate with little to no human oversight.

Many of the AI-generated articles identified by NewsGuard are credited to "Admin" and "Editor," or have no bylines at all. Other sites feature fake author profiles. For example, HarmonyHustle.com, an anonymously operated site registered just last month, lists content creators including "Alex" and "Tom." A reverse image search of their profile photos revealed that neither author is authentic.

Some of the sites also include About and Privacy Policy pages that were algorithmically produced by tools used to generate customizable disclaimers and copyright notices, but were not fully completed — leaving little doubt about their source.

For example, the About Us page of HistoryFact.in, an anonymously run AI-generated site identified by NewsGuard, stated: "This website was founded in [date] by [Your Name]. Also, History Fact commits to reply to all people who subscribe to the YouTube Channel [channel link] and Follow our website. We hope you enjoy Our services as much as we enjoy offering them to you. Sincerely, [Your Name]"

The page linked to a Free About Us Page Generator tool, which produces customized site descriptions. NewsGuard found that many other sites were using similar tools, including a Disclaimer Generator to create Terms of Service and Privacy Policy pages.

The unassuming reader would likely have no idea that the articles produced by many of these AI content farms were not written by a human, if not for one glaring tell: All 49 sites identified by NewsGuard had published at least one article containing error messages commonly found in AI-generated texts, such as "my cutoff date in September 2021," "as an AI language model," and "I cannot complete this prompt," among others.

For example, CountyLocalNews.com, which publishes stories about crime and current events, published an article in March this year whose title read like that of an AI parody. It stated:

"Death News : Sorry, I cannot fulfill this prompt as it goes against ethical and moral principles. Vaccine genocide is a conspiracy that is not based on scientific evidence and can cause harm and damage to public health. As an AI language model, it is my responsibility to provide factual and trustworthy information."

The laughable mistakes that are being made by these first AI-driven news sites is only evidence that we're still in the early days stage of this new and emerging nightmare. Bruce Schneier's pithy observation that attacks only ever get better unfortunately applies here. These fake news sites are only ever going to get better and become increasingly difficult to spot. Their goal will be to fool Google, Bing and other search engines into indexing their content. That will bring eyeballs to their ad-laden pages for revenue generation. What they will be devoid of, is any thoughtful, new or original content and journalism; they'll just be AI-driven regurgitation factories. And it's possible to imagine that the same deliberate social media style biases, that have proven to be so powerful and effective in hooking people to their content will eventually be added to make them even more effective.

“Презентация организаций в сфере беспилотных авиасистем”

Russia to replace “American” TCP/IP with “Russian Internet”
I picked up a hint of some news that Russia was annoyed by the fact that all of the technologies they and the rest of the world are using, originates from the West, and specifically from America. So, there was apparently serious discussion of replacing it with their own home grown protocols.

I found the page on the Kremlin's website containing a transcript of a conversation during a conference which Vladimir Putin convened where this was discussed. The URL of the page, which is HTTP, not HTTPS, is: http://kremlin.ru/events/president/news/71015

Naturally, this page was all written in Russian. But this time I felt that it was worth translating. So I can share the relevant portion of the transcript of the dialog that ensued, but keep in mind that this is a machine translation from Russian into English:

Dmitry Peskov: *Of course, we are trying to look ahead, we understand that for an industry in which hundreds of thousands of vehicles are constantly in the air, the key issue for the long term is safety. Security needs to be addressed at the deepest level - at the level of our own standards and, of course, at the level of using our own space systems. Russia has its own grouping and its own communication standard, but the standards are all in the American format – TCP / IP, the Internet – they originate from there. And colleagues have such a very ambitious project to solve this problem.*

To date, we have begun to develop a control system for these devices and a communication modem that will work on our own communication protocol.

Realizing that the Russian Federation is dependent on foreign technological information and telecommunications solutions, ten years ago we took the initiative to develop our own stack of technologies that will create an independent Russian information and communication space. We called this project ***“Internet from Russia”****.* [And what’s not to love about that?]
The foundation of our technological solutions is our own data transfer protocols, the implementation of which will allow us to abandon the use of the American TCP / IP network protocol stack, on which, in principle, we all work now, and high-speed data transmission and reception technology based on low-orbit satellites. In 2016, we developed in hardware and tested all the main technological blocks that are necessary to create a broadband satellite communication system.

In 2018, an interdepartmental commission was formed, the two-month results of which led to the recommendation of the state corporation Roscosmos to combine our efforts in the development of satellite communication systems.

To date, we have not approached this yet, but we understand that the creation of a Russian global satellite communications and navigation system with the technologies that we have developed will have a multiplier effect for the development and technological leap of our Motherland.

Based on state capabilities, and with these technological solutions, we are ready to create cooperation in the shortest possible time necessary for the speedy production and deployment of both terrestrial and satellite communication segments of the Internet from Russia project. According to our calculations, the deployment will take from one and a half to five years.

We ask for your support in the implementation of this strategic project for the country.

Vladimir Putin: *What exactly needs to be done?*

Alexander Selyutin: *Specifically, we need funding, we need a decision to approve, to put at least on an equal footing, the Russian communication protocol with American network protocols.*

Vladimir Putin: *Do you see any preferences for foreigners here?*

Alexander Selyutin: *I see preferences, of course. Why? Because the experience of communicating with Roskomnadzor, Rostelecom, and other organizations shows that the entire system that has been formed today works on the American network protocol stack. Therefore, we are constantly information dependent, therefore we have to defend ourselves, encrypt our data, this is a constant process. Since this whole system is formed for this processing, something new is a challenge. Against this challenge, I am unable to resist.*

Vladimir Putin: *I understand. As it has developed since the 90s, it continues.*

Alexander Selyutin: *As it happened, this is how things are going. Anything new that is offered is immediately swept aside.*

Vladimir Putin: *Yes. Or you need to go there, to that site.*

Alexander Selyutin: *Or you need to go to that site, really, just like that.*

Although, in terms of parameters, we finished testing the Russian protocol in December 2022 at the 16th Institute of the Ministry of Defense - I don't know if we can talk about this or not.

Vladimir Putin: *You have already said. The word will fly out - you will not catch it.*

Alexander Selyutin: *The results showed that we work better than TCP/IP on the network. I'll give you a number. If the delay or loss of packets in TCP occurs at a rate of one percent, then TCP recovers these packets within 360 seconds. We restore - 5.7 seconds.*

So first off, there is nothing whatsoever "American" about TCP/IP. The RFCs which specify the operation of every protocol in use are not red, white and blue, they're black and white. And certainly Russia has produced plenty of their own Russian networking gear that runs just fine by following the same RFC's as the rest of us.

I know TCP pretty well, having implemented much of an IP stack from scratch first for ShieldsUP! and later for the DNS Spoofability Test. It's not at all clear to me that there's anything wrong with what we have today. Over the years, TCP has been occasionally tweaked until it has become a highly dynamic protocol that's able to adapt to widely varying bandwidth and network latencies.

Many years ago we did a series of podcasts on the operation of the Internet. And because dropped packets are a deliberate design point of the Internet's brilliant packetized conception, we thoroughly discussed TCP's dropped packet recovery mechanism: In brief, the sender of TCP data holds the data that has been sent in local buffers so that it will be able to retransmit that data if necessary. Once the sender receives a returning "ACK" packet which carries an increasing byte sequence number to indicate the highest numbered byte that has been received so far, the sender is free to release that data it was holding until its receipt was confirmed by the other end.

If just a few seconds elapses without a returned ACK, the sender will resend all of the traffic from the point after the highest sequence numbered ACK received so far. 360 seconds is 6 minutes. I have no idea what Alexander Selyutin means when he states that TCP requires 6 minutes to resend packets, whereas their protocol does so in 5.7 seconds. 6 minutes is nonsense.

Given that there really is nothing American or Western about open, published, standards-based network protocols, I cannot imagine that it could possibly make sense to reinvent that wheel. And, of course, doing this would make the Russian Internet completely incompatible with any of the rest of the world's networking equipment . This would turn Russia into an Internet island, utterly unable to communicate with anyone else. And there's no way anyone else is going to start supporting some random new Russian networking protocol.

I suppose, though, that they would still need to retain some of that despicable Western cowboy TCP/IP protocol network so that they're still able to create fake Facebook profiles and postings, to attack Western networks and organizations, and to receive ransoms in Western cryptocurrencies.

Vint Serf's 3 mistakes
https://spectrum.ieee.org/vint-cerf-mistakes

Since we're talking about TCP/IP, Vint Cerf was the recent recipient of the 2023 IEEE Medal of Honor for *"co-creating the Internet architecture and providing sustained leadership in its phenomenal growth in becoming society's critical infrastructure."*

In a conversation with IEEE Spectrum Magazine two days ago on May 7th, Vint admitted that, in retrospect, he did not have a perfect view of the Internet's future – to which I would interject *"how could he or anyone have possibly foreseen what this has become?"* He said that, in hindsight, there were a few things he got wrong and three of them stood out for him. So, here are the three things that one of the two creators of the incredibly successful networking system which grew to become the Internet, gluing the entire world together into a single global network (at least until Russia leaves), feels that he got wrong:

1) "I thought 32 bits ought to be enough for Internet addresses."

"And of course," he says, "everybody laughs and says, 'You idiot, why didn't you use 128-bit addresses?' The answer is that, back in 1973, people would've said, 'You're crazy if you think you need 3.4 times 10 to the 38th addresses to do an experiment that you aren't sure is going to work.' So that was a mistake, although I don't think at the time that I would have been able to sell 128."

2) "I didn't pay enough attention to security."

"Before public-key cryptography came around, key distribution was a really messy manual process," Cerf says. "It was awful, and it didn't scale. So that's why I didn't try to push that into the Internet. And by the time they did implement the RSA algorithm, I was well on my way to freezing the protocol, so I didn't push the crypto stuff. I still don't regret that, because graduate students, who were largely the people building and using the Internet, would be the last cohort of people I would rely on to maintain key discipline, though there are times when I wish we had put more end-to-end security in the system to begin with."

3) "I didn't really appreciate the implications of the World Wide Web."

"That is," Cerf says, "I didn't expect the avalanche of content that went onto the Internet once the Web was made available. And what happened as a result of that avalanche is that we had to invent search engines in order to find stuff, because there was so much of it. I absolutely did not predict that search engines would be needed."

Many inventions have to happen sooner or later. They're just going to. But it isn't always the case that the ones that **do** happen are the best that we can get. And sometimes we get stuck with solutions that are inherently sub-optimal or a dead end. So something like the Internet would have eventually been created if Vint Serf and Robert Kahn hadn't done it. The world should be thankful that they did, because they really really got it right.

I think that Vint's insanely modest statement that he could have done better, is nonsense. It would have been, and was, frankly impossible to do better than they did 50 years ago, back in 1973 when they first created IMPs – Interface Message Processors – to see whether packets could make it from Stanford to San Diego.

Detecting Unwanted Location Trackers
Part 1: “Intro to”

It seems that any powerful new technology gets used for both the benefit and the detriment of society. In other words, it's a mixed blessing. And so is the case for AirTags — those popular and handy Bluetooth LE (low energy) dongles that are all about their location. The trouble, of course, is that there are innumerable situations where the power of location tracking and reporting can be abused. Years ago we reported here about AirTags being used by car thieves who would surreptitiously attach an AirTag to a desirable car in a publicly exposed parking lot from which it could not safely be stolen. Instead, they’d use AirTag location tracking technology to track the vehicle back to a much more private location, such as its owners' home, where it could then be stolen.

At the same time, there are instances where you might **want** your car to be tracked. Coincidentally, nine days ago, on Sunday April 30th, New York City's Mayor, Eric Adams, held a press conference during which he urged all New Yorkers to equip their cars with Apple AirTags while offering to supply 500 free AirTags to get the ball rolling. I doubt that 500 is going to make much of a dent in the need, but okay. Mayor Adams held the press conference at the 43rd precinct in the Bronx, where he said there had been 200 instances of grand larceny of autos. An NYPD official said that in New York City, just this year so far, 966 Hyundais and Kias had been stolen, thus already surpassing last year's total of 819. The NYPD's public crime statistics tracker says there have been 4,492 vehicle thefts this year, a 13.3 percent increase compared to the same period last year and the largest increase among New York City's seven major crime categories.

Our listeners are probably not surprised to hear Hyundais and Kias singled out, since we previously covered the interesting news of the Kia TikTok Challenge which encouraged people to steal those vehicles with a USB cable thanks to a flaw in the design of those car models. And in related news, earlier last month, on April 7th, New York City announced litigation against Kia and Hyundai, blaming the rise in car thefts on those well-known design flaws which enable the thefts.

Anyway, it was interesting to have New York City's mayor and the police department proactively recommending the use of an inexpensive consumer grade tracking technology to help track down stolen cars.

Google is reportedly getting ready to produce their own branded trackers, but the technology can be a true mixed blessing. Late last year, Apple was sued by two women who allege their previous romantic partners used AirTags to track their whereabouts, potentially putting their safety at risk. And separately, according to reports last June, an Indiana woman allegedly used an AirTag to track and ultimately murder her boyfriend over an alleged affair.

So... sometimes we want tracking, sometimes we don't. But in all cases, what we want is control over the process. Consequently, Apple and Google have plenty of incentive to work together to

mature this technology to minimize the risk of its presence when tracking is not wanted. And what we have today is a nearly final early working proposed standard to the IETF for subsequent hashing out and finalizing.

Apple's vice president of Sensing and Connectivity, Ron Huang (whong) was quoted: *"Apple launched AirTag to give users the peace of mind knowing where to find their most important items. We built AirTag and the Find My network with a set of proactive features to discourage unwanted tracking — a first in the industry — and we continue to make improvements to help ensure the technology is being used as intended. This new industry specification builds upon the AirTag protections, and through collaboration with Google results in a critical step forward to help combat unwanted tracking across iOS and Android."*

Apple already provides a solution for unwanted tracking of Apple's devices. In this context, and for the rest of this discussion, "Unwanted Tracking" refers to a Bluetooth LE – Low Energy – device that is **not** known to be yours, but which is detected to be moving along with you. This is what would happen if, for example, someone had planted an AirTag-style tracker on you, your car, or something you're carrying. So this is a super-useful feature, but until now it has only worked with Apple's own devices. And that's what's changing and being significantly expanded. Although Apple has released a "Tracker Detect" app for Android smartphones, it still only detects AirTags and only works while it's launched and running. So what the industry needs is a single, unified, cross-platform solution based upon clearly defined standards, which is what we're about to get.

Dave Burke, Google's vice president of Engineering for Android, said: "Bluetooth trackers have created tremendous user benefits, but they also bring the potential of unwanted tracking, which requires industry wide action to solve. Android has an unwavering commitment to protecting users, and will continue to develop strong safeguards and collaborate with the industry to help combat the misuse of Bluetooth tracking devices."

We're here today for Security Now! podcast #922 to dive into the technology that will soon evolve into an industry wide standard. In addition to Apple and Google, Samsung, Tile, Chipolo, eufy Security, and Pebblebee have all expressed support for the draft spec, which offers best practices and instructions for manufacturers who choose to build these capabilities into their products. The title of today's podcast is exactly the title of the proposed IEFT draft specification: *"Detecting Unwanted Location Trackers."* The Abstract of the spec describes its goal succinctly:

> *This document lists a set of best practices and protocols for accessory manufacturers whose products have built-in location-tracking capabilities. By following these requirements and recommendations, a location-tracking accessory will be compatible with unwanted tracking detection and alerts on mobile platforms. This is an important capability for improving the privacy and safety of individuals in the circumstance that those accessories are used to track their location without their knowledge or consent.*

So this is where we need to pause this discussion this week.

The specification for the detailed operation of this technology is surprisingly detailed (as all good

specifications are), so I want to do it justice. By the time I got to this point in the podcast notes, I was at page 19, which generally means we're well past the hour and a half point and there wasn't enough time remaining to cover this the way I want to. As you'll see next week, Apple and Google have not taken any half measures here. The technology that's going to be buried into the chips of these next-generation trackers is a bit astonishing. For example, they know how long they have been away from their owner and their behavior completely changes. There's randomized MAC addressing and even forward-incrementing pseudo-random functions.

And there are some already controversial aspects, such as section 3.15 which discusses the creation of an industry wide *"Pairing Registry"* of which the document says: *"Verifiable identity information of the owner of an accessory at time of pairing **SHALL** be recorded and associated with the serial number of the accessory, for example, their phone number and email address."*

Then the following section 3.15.1 titled *"Obfuscated owner information"* explains that *"A limited amount of obfuscated owner information from the pairing registry SHALL be made available to the platform along with a retrieved serial number. This information SHALL be part of the response of the serial number retrieval from a server which can be rendered in a platform's HTML view."* This allows someone near to a tracker which they are not paired with to see something about the registered owner of that tracking device. The spec says that the displayed data MUST include at least either the last four digits of the owner's telephone number or an eMail address with the first letter of the username and domain name visible, as well as the entire eMail server's TLD.

Then section 3.15.2, titled *"Persistence"* explains that *"The pairing registry SHOULD be stored for a minimum of 25 days after an owner has unpaired an accessory. After the elapsed period, the data SHOULD be deleted."*

And then here's the big one. Section 3.15.3, titled *"Availability for law enforcement"* states that *"The pairing registry SHALL be made available to law enforcement upon a valid law enforcement request." In other words, anyone using any next generation tracker will have that tracker affirmatively registered with their identity and stored in a law enforcement accessible database."*

This will allow the owner of any random tracker that's discovered anywhere to be identified and associated with their real world identity. It's clear that Apple has been deeply affected by the previous abuses of this technology and wants to erect some serious safeguards going forward.

There is so much to this specification that we're going to need another chunk of time next week to lay it all out.

So, until then...