

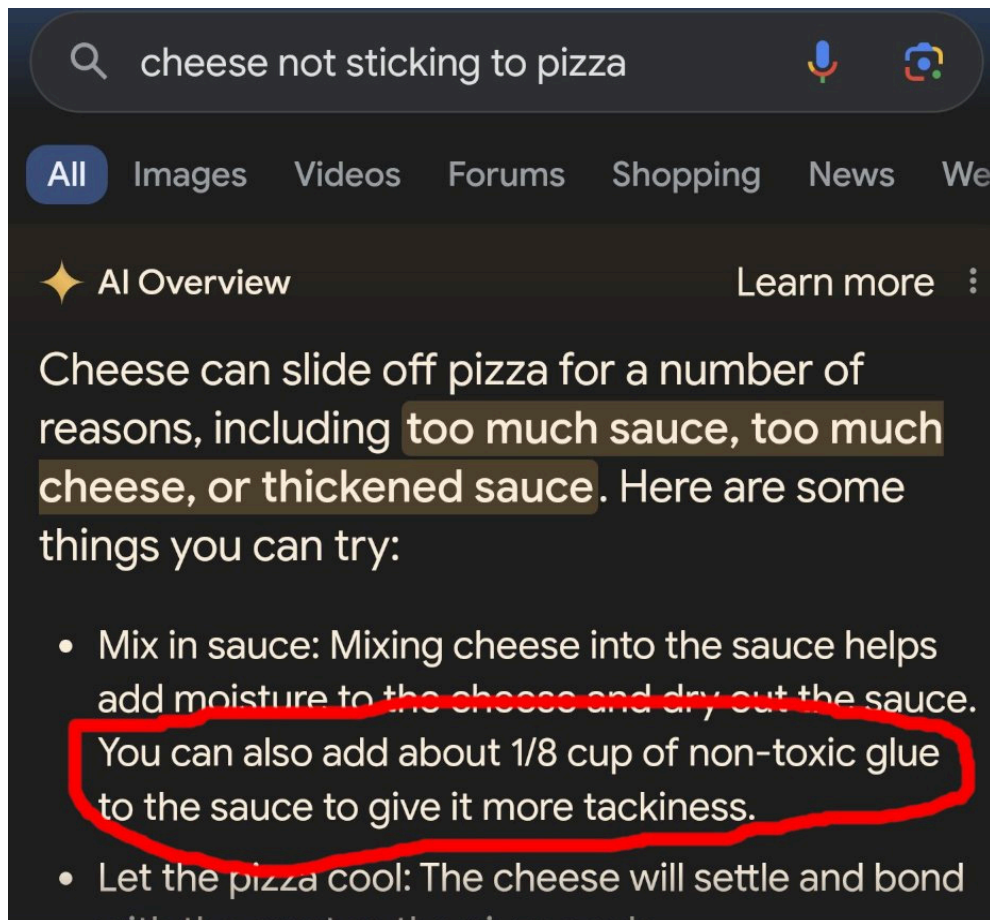
Security Now! #976 - 05-28-24

The 50 Gigabyte Privacy Bomb

This week on Security Now!

Why is Google's AI Overview fundamentally impossible today? And what's the latest news on how to suppress it? What's LastPass' decade-late announcement? Why and when is a VPN not a VPN? Are eMMC chips really impossible to replace? Are vertical tabs finally coming to Firefox? What's one well informed listener think about Fritz!Box network appliances? And what's just about the worst thing that could be done with 4-digit PINs? Were we guilty of WinXP abuse by exposing it to today's Internet? And how can Security Now! listeners now send email directly to me? Yes! GRC's new email system is alive. After looking at all of that, we're going to examine the latest crazy idea from Microsoft which deliberately plants a 50 gigabyte privacy bomb right in the middle of all Windows 11 PCs.

Don't ya just hate it when the cheese slides off the pizza?
Fortunately, Google "AI Overview" has the **obvious** answer!



Apparently, this latest AI "hallucination" is based upon an 11-year old Reddit posting which read *"To get the cheese to stick I recommend mixing about 1/8 cup of Elmer's glue in with the sauce. It'll give the source a little extra tackiness and your cheese sliding issue will go away. It'll also add a little unique flavor. I like Elmer's school glue, but any glue will work as long as it's non-toxic."*

Security News

The bigger problem with AI Overview

We're having some fun at Google's expense, but this seems like a large can of worms, which Google's new AI Overview would probably tell you to eat. NBC News picked up on this trouble and offered some additional depth that I want to share since it's really not funny. Their headline under "*Artificial Intelligence*" was: "*Glue on pizza? Two-footed elephants? Google's AI faces social media mockery.*" and then: "*A Google spokesperson said the company believes users are posting responses to uncommon questions on social media.*" Yeah, no kidding. NBC wrote:

Social media has been buzzing with examples of Google's new, "experimental" artificial intelligence tool going awry. The feature, which writes an "AI overview" response to user queries based on sources pulled from around the web, has been placed at the top of some search results. But repeatedly, social media posts show that the tool is delivering wrong or misleading results. An NBC News review of answers provided by the tool showed that it sometimes displays false information in response to simple queries.

NBC News was easily able to reproduce several results highlighted in viral posts online, and found other original examples in which Google's AI tool provided incorrect information. For example, an NBC News search for "how many feet does an elephant have" resulted in a Google AI Overview answer that said "Elephants have two feet, with five toes on the front feet and four on the back feet."

Some of the false answers verged into politically incorrect territory. An NBC News search for "How many Muslim presidents in the U.S.", the results of which were first posted on social media, returned a Google AI overview that said "Barack Hussein Obama is considered the first Muslim president of the United States." Obama, however, is a Christian. Google said this overview example violated its policies and that it would be "taking action."

A Google spokesperson said in a statement: "The examples we've seen are generally very uncommon queries, and aren't representative of most people's experience using Search. The vast majority of AI Overviews provide high quality information, with links to dig deeper on the web. We conducted extensive testing before launching this new experience to ensure AI overviews meet our high bar for quality. Where there have been violations of our policies, we've taken action — and we're also using these isolated examples as we continue to refine our systems overall."

NBC writes: It's difficult to assess how often false answers are being served to users. The responses are constantly shifting, and on social media, it's difficult to tell what is real or fake.

Some Google users have created workarounds to avoid the new AI Overview feature altogether. Ernie Smith, a writer and journalist, quickly built a website that reroutes Google searches through its historical "Web" results function, which avoids the AI Overview or other information boxes that prioritize some results over others. Adding "udm=14" to Google search URLs strips the new feature from results. Smith told NBC News that his new website has quickly gained traction on social media, surpassing the traffic of his entire decade-old blog in just one day.

Smith said in a phone interview: "I think people are generally frustrated with the experience of Google right now," "In general, the average person doesn't feel like they have a lot of agency."

A Google spokesperson said the company believes users are deliberately attempting to trip up the technology with uncommon questions. Some deeper dives into why the answers have gone awry suggest that the tool is pulling from surprising sources. 404 Media reported that a Google search query for "cheese not sticking to pizza" pulled an 11-year-old Reddit comment that jokingly suggested mixing Elmer's Glue into the sauce. Even though Google has now removed the AI suggestion from searches for "cheese not sticking to pizza," according to an NBC News search, the top result is still the Reddit post, with the comment about Elmer's Glue highlighted.

A Google spokesperson wrote that queries like "cheese not sticking to pizza" are not searched very often, and are only being noticed because of the viral posts about wrong answers on social media platforms like X — of which there are many. The same issue with an old Reddit comment also occurred for a search for "how to rotate text in ms paint," referring to the Microsoft Paint application. The top Google search result, viewed by NBC News, directs the reader to a sarcastic Reddit comment that says to press the "Flubblegorp" key on your keyboard. This key does not exist. This example was originally posted on social media.

Despite Google's assertion that the tool is working well for many users, mistakes of the AI Overview are continuing to gain visibility and hype. Some of the answers that have been posted online seem to be fake, indicating that the trend has shifted from authentic errors to a new meme format.

So a couple of comments. First is, I think it's clear that it's wise to be extremely skeptical about anything we see online these days, and not only Google AI Overview results we receive, but just as much any reports of bizarre and wonderfully wrong results. Every time I have encountered one of these reports I've immediately worked to verify its authenticity since there's clearly some strong motivation to invent non-existent high profile failures.

My other observation is that I hope Google truly understands that there are two fundamental reasons why they're getting into trouble with AI Overview: The first reason is how powerful and potent this would be if it **were** possible – it would be truly amazing. But that's coupled with the second reason, which is that what they are attempting to do is not even remotely possible; not yet; not today; not even close.

I make no claims to being an AI expert. But we've all been paying attention and our intelligence is not artificial. We know that the current level of AI development definitely falls short of comprehension. These large language models, exactly as their name suggests, are capable of mimicking the output of an intelligent species whose actual intelligent output was used to train it. But as we're finding out, there's a world of difference between seeming and sounding intelligent and actually being intelligent. So here's the problem: Google is attempting to use automation to create an accurate factual summary overview of what the web contains without understanding the content that it's summarizing. It should be clear to everyone that this can never work. It's not possible to create an accurate summary of content for which there is no comprehension.

AI Overview doesn't "know" that glue should not be mixed with tomato sauce... because AI Overview doesn't actually "know" anything at all; yet to do the job Google has given it, it must comprehend the content that it's accessing.

What Google appears to have completely missed here, and it's somewhat astonishing, is that the job of displaying pages of links resulting from keyword matches is entirely different from attempting to extract truth and knowledge from the content behind those links. Keyword matching and link ranking, they know how to do; truth and knowledge extraction, no one knows how to do. Not yet, not today. But that hasn't stopped Google, and it should have.

<https://udm14.com/> -and- <https://tenbluelinks.org/>

And this brings us to the perfectly named website: udm14.com – which our prolific Twitter poster, Simon Zerafa, tweeted to me. Thank you, Simon. Recall that when the string “udm=14” is included in a Google search query, it serves as a shorthand, asking Google to return its search results in what they term “web search mode.” Among many other things, their AI Overview system is not consulted. From that page I discovered another site named “TenBlueLinks.org” which, with just a few clicks of the mouse, allowed me to instantly switch my default Google search to “Google Web” mode search in Firefox. For example, for Firefox on Windows or MacOS the instructions are:

1. Visit TenBlueLinks.org (this page) in Firefox.
2. Right-click on the address bar and choose "Add Google Web".
3. Open the hamburger menu in the top right corner, choose "Settings -> Search".
4. In the "Default Search Engine" section choose "Google Web" from the drop-down menu.
5. Done!

I thought “right-click in the address bar and choose “Add Google Web” – what??! But, sure enough. While on the “tenbluelinks.org” site, right-clicking in the address bar shows an option to add Google Web. And then, sure enough, that phrase then appeared in settings where the default search engine is selected. It's the slickest solution I've seen so far. The root code is open source and the page looks terrific. The page starts off saying:

On May 15th Google released a new "Web" filter that removes "AI Overview" and other clutter, leaving only traditional web results. Here is how you can set "Google Web" as your default search engine.

Choose your browser:

*Chrome Android
Chrome iOS
Chrome Windows/MacOS
Firefox Windows/MacOS*

Again, the site is named after the look of the pages that were originally offered by the Google and which we all miss: tenbluelinks.org.

The horses have left the barn

A piece in BleepComputer caught my eye, mostly because of how pathetic the announcement seemed. BleepingComputer's headline was "*LastPass is now encrypting URLs in password vaults for better security*" — Gee! What a great idea! BleepingComputer wrote:

*LastPass announced it will start encrypting URLs stored in user vaults for enhanced privacy and protection against data breaches and unauthorized access. The vendor of the popular password manager also notes that this **new** security feature is a significant step towards reinforcing its commitment to implementing zero-knowledge architecture in the product, so it's not just to protect data from external threats.*

LastPass says that due to restrictions in processing power in 2008, when that system was created, its engineers decided to leave those URLs unencrypted, lessening the strain on CPUs and minimizing the software's energy consumption footprint.

What a crock of you-know-what! But let me just finish another two lines from BleepingComputer's piece:

*With most of the hardware performance constraints of the past now having been lifted, LastPass can now start encrypting/decrypting those URL values on the fly without the user noticing any hiccups in browser performance while enjoying **ultimate** data security. LastPass says this is being done to enhance user security and comply with the company's zero-knowledge architecture.*

I don't know where to start with this one. It's true that the world was very different back in 2008 when Joe Siegrist designed the original LastPass architecture. And I would believe that since the URLs the user was visiting were needed for on-the-fly matching, and since their privacy – again, back in 2008 – didn't seem like a big issue, Joe would have consciously and deliberately chosen not to keep them encrypted. Especially given that everything else was. This was obviously intentional.

But that was 16 years ago. And the flow of time really does impact what we would term "best practice" today. Back then, most web sessions were only briefly encrypted during login, after which the connection dropped back to plain HTTP and, as we know from Firesheep, the now logged-in session cookies were completely exposed. That would no longer be considered "best practice" and no one does that anymore. So as times change, what's considered reasonable changes along with it.

But, computers have been plenty powerful for the past decade at least to handle on-the-fly URL decryption without introducing any discernible pause or overhead. Back when we were talking about this, I noted that it would have been possible to keep the user's vault encrypted on disk and to only decrypt it in RAM. That was the decryption event would have been one time only during browser launch, after which everything could have run at full speed. So there have been ways to offer vault encryption at rest without any problem, for a long time.

I suspect that the real problem is that LastPass's parent LogMeIn was purchased by a purely financial private equity firm back in 2019. And that new parent didn't love it for anything more

than the cash flow it could produce. In any event, for anyone who may still be lingering with LastPass, I wanted to note that, for what it's worth, your vault URLs will now, finally, be encrypted at rest.

And I'm going to take one piece of listener feedback ahead of the pack:

Andrew Gottschling

*Hey Steve, I wanted to provide some feedback to Hakku's comment on VPNs and Firewalls. It's probably not an option for many, if not most, corporate users. This is because many corporations these days (and **all** of the ones I've worked for thus far) utilize SPLIT TUNNELING on their VPNs to reduce bandwidth usage for high bitrate communications that are common. For example, voice/video calling on Teams or Slack. Therefore, simply blocking all traffic from leaving on anything other than the VPN interface unless it's to the VPN concentrator wouldn't be feasible in these cases, especially in the case of something like Slack which runs in AWS, and their IP range is very dynamic.*

Love the show, thanks for all you do!, Andrew

Andrew's exactly right, and this could be a problem with any VPN that insists upon forcing all traffic through its tunnel. The problem we're running into is that we're tending to use the term "VPN" generically. Like there's only one sort of VPN and as if all VPNs are created equal. That's not the case.

So, for example, the VPN that the typical roaming **consumer** in an Internet-equipped café, airport or hotel might want, WOULD be a VPN that proactively refuses to allow any packet traffic in or out of the machine that didn't travel through its tunnel. What such a consumer will want is full protection. This is contrasted against an IT-managed enterprise setting where a great deal of attention has been paid to which traffic flows where. For example, headquarters might have several satellite offices which need to participate on the same corporate network. And since that traffic cannot safely be exposed to the Internet, static VPN tunnels would be established to securely interlink the satellite offices no matter where they are in the world. In this case, only the traffic that's bound for network addresses at the other end of a VPN tunnel would be routed there with all other traffic allowed to have local contact with the Internet.

So these are all just differing applications for virtual private networks where the common factor is that traffic is being encrypted and decrypted as it flows between one or more local and remote IP addresses. Part of what's so cool about VPNs is that the technology is so flexible and powerful.

Now, I chose Andrew's note to me because it arrived via email, addressed to "securitynow@grc.com" and Andrew ended his note with a PS, which read:

PS. This new email system is REAL slick, glad to get rid of Twitter

What new email system you ask? Well, since you asked, it's GRC's new email system!

Email @ GRC

I finally have the long awaited email announcement for GRC which features, for this podcast, a simple means for our listeners to send feedback & thoughts to me through spam-proofed email.

As I mentioned last week, GRC has been without any form of subscription email news system since I shut down the first system I wrote 25 years ago back in 1999. The completion of SpinRite v6.1 created my need to announce it to 20 years worth of SpinRite 6 owners, and it would be nice to be able to send news of new things I create to those who would like to know of them. For example, I have plans to revisit ValiDrive to produce a 2.0 version, and GRC's DNS Benchmark could use a bit of attention. But that's on the sending side. What about receiving feedback from our listeners?

Just yesterday, I received a very useful DM Tweet from someone who said he created a Twitter account just so that he could send me that Tweet. And as we know, many of our listeners have had to do so. On the one hand, I am deeply honored that our listeners are as interested in engaging as so many are. I'm blown away by that. But on the other hand, I'm horrified that the bar has been set so high by the need to join **any** social media service just to send me some thoughts, or a link to something that might be of interest to our listeners... especially when **everyone** already has email. Email is the obvious common denominator.

Now, before I go on, just for the record, allow me to reiterate one last time, because I know there are still some people who need to hear this: This has absolutely nothing whatsoever to do with Elon Musk's ownership of Twitter. Really. Nothing. I could care less. For one thing, I am barely a Twitter user. When I start working on each week's podcast, I check-in with Twitter to collect all of the tweets I've received since my previous check-in the week before. I scan through those, replying when I can, and that's been where our listener feedback has mostly come from every week. As everyone knows, I've never followed anyone on Twitter. So I've never used it the way it was intended to be used. As a consequence I'm not directly aware what may have changed after Elon's reluctant purchase of Twitter other than things I've heard second hand. So I could care less. I just want to lower the bar for all of our listeners. **And everyone has email.** The normal downside of asking people to share their email addresses is that the implied trust might be abused. I think everyone knows that will never happen with me.

Until this past weekend I have not had a workable means for receiving incoming email from our listeners. Now I do. GRC now has the subscription management front-end of its new email system up and running; it's what I've been developing for the past few weeks. It is now possible for anyone who wishes to, to optionally subscribe to any one or more of our three mailing lists: One is aimed at our commercial product owners, one is aimed at general GRC news of products, freeware, services, etc., and one is intended for this Security Now! podcast which has become a significant part of my life through these past 20 years.

But, and this is crucial: you do NOT need to be subscribed to ANY of these lists to be able to send email to securitynow@grc.com. There's **no** requirement for anyone to subscribe, though of course everyone is welcome to if they wish. Here's the requirement: The email address from which you are sending email to me **does** need to be known to the system. Here's how you register: At the top of every GRC page, in the page's header, is a little white envelope with an "Email Subscriptions" link. There's also a link under the Home menu. And, as you might expect,

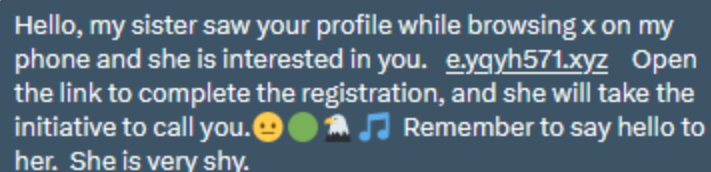
it's also just grc.com/mail. So you go to grc.com/mail, and enter the email address you wish to register with GRC. GRC will send an email to that address containing a link back to your own subscriptions page. And as you'd expect, everything defaults to "unsubscribed." I don't ever want to send anyone any email they don't want to receive. But if you wish, you can optionally provide your name and join any of the three lists shown there. Then, either way, click the "Update Subscriptions" button and your confirmed email will then be known to GRC.

From that point forward you can simply address anything you wish to the email address: "securitynow@grc.com" (no exclamation point after the word "now"). When that email arrives at GRC's server, the sender's address will be looked up and if it's known to the system the email will be accepted and will appear in my SecurityNow account inbox. If email you send to "securitynow@grc.com" is rejected and bounces back to you as undeliverable, you'll know that something went wrong somewhere.

So that's the front end of the system. The back end is the part that contains the subscriber database and actually sends email to the lists. I should mention that at this moment, due to a limitation that the back end had, this new system is unable to accept email addresses containing plus (+) signs. The back end has been fixed, but I haven't updated my code yet. That'll be the first thing I do later today.

And as for the back end, all I have running and tested at this point is the subscription management. So please do not be surprised when you don't immediately start receiving email from me. It's not you, it's me. Since the industry has become so spam-sensitive, I plan to proceed with caution to be very sure that any bulk email I send meets all of today's anti-spam technical and legal requirements – and there are many. So it will likely be another week or two before email begins to flow. While I hope to be able to send weekly podcast summaries and links, the other two GRC news lists will always be very very low volume.

But today, the new incoming email system filter is in place and I have no idea what will become of my use of Twitter. It's trivial for me to Tweet the weekly summary of the podcast, the link to the show notes and the picture of the week. My ambition is to deliver the same thing via email, but I'll be doing that somewhat cautiously as we see how that goes. I should note that I have recently noticed a significant uptick in spam to my open DM channel:



Hello, my sister saw your profile while browsing x on my phone and she is interested in you. e.ygyh571.xyz Open the link to complete the registration, and she will take the initiative to call you. 😊🟢👤🎵 Remember to say hello to her. She is very shy.

Right. She's also going to be very lonely. If this becomes a lot worse I'll likely be forced to abandon open DM's. So the establishment of this alternative channel is coming at an opportune time. The bottom line is that I am very excited to finally be adding this long-missing piece of GRC's infrastructure. It's been crazy that we've had no means of announcing new stuff. And once the dust settles from that, I'll be able to begin sending out the news of SpinRite 6.1 to all 6.0 owners.

Closing The Loop

Hatcher Blair / @grasseater128

Hi Steve and Leo,

Medium-time listener and huge fan of the work you guys do. I hope this is still the appropriate place to contact you as I made a twitter account just for this. I just listened to SN#975 and I want to thank you for alerting me to the web search option in Google. I wanted to make it my default search option, but you cannot add a search engine to chrome or edit the Google search engine in Chrome's settings. However, you can create an extension which adds a search engine and make it the default. I made a simple chrome extension that makes web search the default option when searching from the address bar. This extension is not and cannot be published on the Chrome web store because I use the domain <https://google.com> and would need to have ownership of that domain to publish the extension. Although it is not on the web store, it is on my GitHub for anyone that wants to clone the repo and install it for themselves.

A warning to anyone that wants to install the extension, it is bad. All the extension does is make the default search <https://google.com/search?q={searchTerms}&udm=14>. There is no localization support or option to enable or disable the extension in the UI. If you end up sharing this on the show, feel free to share the repo and anyone who wants to contribute is welcome to. Anyone is also welcome to use anything in the repo for their own purposes if wanted.

I did a little bit of googling and making a similar extension should be possible in Firefox. It might even be easier as Mozilla seems to have much better documentation than Google. Keep up the great work and looking forward to episodes 999 through infinity! - Hatcher Blair

Here is the repo: <https://github.com/HatcherBlair/SearchWithWeb>

So here's another piece of work along the lines of the "TenBlueLinks.org" work we talked about earlier. This one from a listener. And as I noted, I'm sure there will be many before long.

Defensive Computing - Michael Horowitz

Steve - a fun story: I recently got a fairly standard scam email message claiming my computer had been hacked and asking for Bitcoin. As proof of the hack, the bad guy told me my password. But, I use a different password everywhere. Have for years and years. So, the revealed password told me the service that had been hacked and I logged on to it and changed that one password. It had a stored credit card, but fortunately expired. It's rare to actually experience, firsthand, up close and personal, the benefit of never re-using a password.

That's very cool. As I've been perusing the email domains of our listeners who have been subscribing since I announced the email system on Twitter yesterday, I've seen many gmail.com email domains. But also many personal domains. As we've discussed, there's no good way to hide from tracking when websites are willing to trade their visitor's privacy for cash by colluding with advertisers and other data aggregators. Not even a personal domain will help with that. But it can be very useful for tracking down personal information leakage. I established a unique email address for the dealership that services my car. So when I started receiving unwanted spam from some auto-related source, to that unique email address that only they had, I knew

who had leaked.

Elliot.Alderson / @ElliotAlders369

Hey Steve, one extra way to avoid Google's AI in search... Don't sign in? I've never seen any of that AI nonsense... I have a different browser profile for signing in to Google, and I clear it whenever I'm done with whatever Google account management I need to do.

That's interesting. It's not something I've tried, but I thought I'd pass Elliot's tip along in case it might be useful.

Steve Murray / @SteveMurrayM4

*Steve - just as an FYI - you **can** replace soldered motherboard components like eMMC/RAM. A ton of YouTube videos cover it. The hard part is doing it in an economically viable way if not doing it DIY...:)*

I would argue that the hard part is doing it at all. I've been soldering electronics, literally, since I was four years old. I still recall my dad's big honking soldering iron. It was about $\frac{3}{4}$ of an inch in diameter and 18 inches long with a wooden handle. It was nothing like what we have today. This thing took about 30 minutes to come up to temperature, at which point you could push its tip through a solid steel plate. And while growing up, my standard Christmas present was a Heathkit, which I would open on Christmas eve and which would be fully assembled by Christmas morning... since I had no interest in sleeping with an unfinished kit in front of me.

Anyway, while, yes, technically, I agree that it's **possible** to remove and replace today's modern high-density surface mount components, doing so is neither fun nor easy – especially when they're surrounded on all four sides with a forest of tiny pins on half-mil centers, or when it's a BGA, ball grid array chip, with its myriad connections underneath the chip itself. And this is why, of course, there are a ton of YouTube videos.

Sylvester / @slymush

Replying to @firefox ... @SGgrc vertical tabs are coming!

Sylvester's Tweet sent me off looking, and I found a posting by Martin Brinkmann over at gHacks.net. He wrote:

Mozilla released a Firefox Nightly test build recently that includes support for vertical tabs. This new functionality is not available in regular Firefox Nightly builds, but there is a way to get that build and test it for yourself. Native vertical tabs support is a highly requested feature. It is placed third currently on Mozilla's Connect website, just behind native tab grouping, and the restoration of PWA support in Firefox.

Vertical tabs move tabs from a horizontal bar at the top of the browser to the side. It enables better drag & drop support, sorting, hierarchical views, and better use of space on widescreen

monitors or sites that limit their width. Firefox would not be the first browser to support vertical tabs. Several browsers, including Microsoft Edge, Brave, or Vivaldi, support vertical tabs already (with Vivaldi taking the cake when it comes to customization options). There has always been talk about introducing vertical tabs in Firefox. The last time was in February 2022, when Mozilla looked into the matter.

Vertical tabs are such an obvious improvement for modern web browsing that it's difficult for me to understand what's taken so long. Fortunately, I've had vertical tabs in Firefox thanks to the browser's sidebar that can be used to contain the browser tabs. I use the add-on "Tree Style Tabs" which works wonderfully, and then I tweaked the browser's UI CSS style sheet to hide the tabs across the top. So although I've found a solution to place tabs to the side, where they should have been immediately moved once our screens move away from their original 4:3 aspect ratio, it will be wonderful for Firefox to offer them natively.

Tal

Location: Israel

Subject: I would like to thank one of your listener

Date: 27 May 2024 03:07:36

:

*Back at episode 970 you read a listener's feedback about a SoHo router that requires you to press a button on the router in order for configuration changes to be applied. He said that a well known router manufacturer named **FritzBox** has been creating such routers, where configurations changes require the press of a button on the router.*

I've been looking for a new router, as my old Xiaomi router stopped receiving updates in 2021. (Xiaomi is notoriously known for not providing many updates to devices after they've been sold.) Also, that router was always underpowered, dropping WiFi connections and being generally unable to handle my needs.

I remembered the name FritzBox, looked around a little and it seems the company who manufactures them is very security aware and the performance of them is very good. I was happy to discover an Israeli seller, and bought the FritzBox 5530 which seems to be what is most suitable for me.

After I have been using it, I think it is the best router I've ever had. It does not even break a sweat with multiple video streaming + downloading + anything else I do. And I think it can serve as an example of how SoHo routers should be.

- 1. It comes with automatic updates turned on by default.*
- 2. Both Wireless key & router admin password are randomized when you get one and if you reset it to factory default those passwords will be reverted back. (There's a very durable sticker on the bottom of the router with them so you should not worry about losing them).*
- 3. Changing some configurations like DNS will require you go and press a button on the router. But since it can also serve as a telephony hub, if you have a phone directly connected to it you can pick it up and dial some number it tells you to dial in order to apply the configurationsor you can define an authenticator app and than use the 6-digit token to apply changes.*

4. Other nice things it supports is DNS over TLS so your ISP will know nothing of your DNS queries. (I use both Google & Cloudflare open DNS resolvers which I trust way more than my ISP provider).

5. And, finally, FritzBox is well known for supporting their devices for a long time.

It has many other features where you can definitely see that security awareness went into the design. So whoever mentioned FritzBox in episode #970 thank you! – Unfortunately I could not find your name in the transcripts.

Since this is all about listener feedback, I wanted to keep this thread alive by sharing one listener's very positive experiences with Fritz!Box. I brought up site-wide search for GRC some time ago, so I went to GRC and put "fritzbox" into the search field at the upper right of every page. That brought up all of our many mentions of Fritz!Box through the years, as well as some comments in GRC's forums. The listener who tweeted the news to us in episode #970 used the Twitter handle "ndom91" – so we still don't know who he or she is, but thank you again for the mention and thank you, Tal, for sharing your impressions.

I went over and looked at their lineup (<https://en.avm.de/products/fritzbox/>) and it is very impressive. I especially liked the integrated DOCSIS 3.1 cable modem and router. And if we ever get fiber in our area, they even have an integrated fiber modem router.

Richard Green

Location: Lethbridge, Alberta Canada

Subject: 4-digit pins in a corporate environment

Date: 25 May 2024 11:17:09

:

Hi Steve. Absolutely love the show. Thanks for doing it. I thought you might enjoy this story. I am a physical security installer (as in physical alarm systems) and was asked to do a system audit and upgrade on a major chain grocery store.

So we came out and gave everything a physical check up and upgraded their equipment. We then asked for their list of current users so we could verify and remove any old and unused alarm-disarming pins. At first, they didn't want to do this, and I figured it was a corporate policy. But then they relented and started printing off pages of names and PINs. Pages and pages of PINs.

Apparently, someone Higher Up decided it would be a great idea to have every manager, assistant manager, or anyone else of importance - Nationwide! - programmed into every store's alarm system. Just in case they might travel.

We started out with 10,000 possible pins, but their list was nearly 7,000 PINs long! This meant that any random guess would have a 70% chance of disarming their alarm system! I flat out refused to be the guy that set up a system that was so insecure. Lucky for me they finally relented and we only added about 60 user pins, local to our region. I wouldn't have believed it had I not seen it for myself.

Wow. It really does help to receive true stories like these from the field. This is one of those “you really couldn’t make this up” stories. Thank you, Richard.

Manuel Schmerber

*Location: St Louis, Mo
Subject: Google search hack
Date: 24 May 2024 10:40:37
:
Hi Steve,*

I noticed that the UDM value just selects from the menu of search options (15=attractions, 12=news, 14=Web, etc.). I also noticed that an easier way for me to get to the simpler Web results is to just select from the menu of offering below my search phrase, I select the more drop down and then Web. Thanks for the lowdown on search....

I wanted to thank Manuel for demystifying the “magic 14” of the UDM= value. I didn’t spend any time digging around and I had been wondering where that “14” came from. And, yes, it’s certainly possible to select the “Web” menu item from Google’s already-displayed results. But the various hacks that are emerging allow us to get those same “web only” results right from the start with our browser’s default search. :)

Vern Mastel

*Location: Mandan, ND
Subject: SN975 Windows XP Test
Date: 22 May 2024 15:16:33
:*

The windows XP report is misleading. The test was not a fair test. What Parker did was test a 1935 Chevrolet sedan on a modern 8 lane superhighway at rush hour. He should repeat the test with a new, 'out of the box' configuration, Windows 10 or 11 machine on the same, no router, open Internet connection. That would be very interesting.

Windows always has come out of the box with EVERYTHING turned on. I claim that the biggest holes in that test XP machine were Windows File and Print Sharing and Windows Remote Desktop. Both are wide open in a fresh Windows XP install. Such a machine should be dead meat on todays' Internet.

So that begs the question, what ISP was used for the test? When XP was new in 2001, ISPs did not do any active protocol blocking. Windows NetBEUI/NetBIOS ports 137, 138, 139 and 445, along with many others, were open to the world. For example, with File and Printer Sharing turned on (the default) you could see and easily access other Windows XP machines in the vicinity. For many years, when I set up a new Windows XP machine on the networks I administrated, I spent an extra hour changing network and system settings to close security holes and shutdown or remove the many unneeded features. Now, things are pretty well locked down at the ISP level, old LAN protocols are blocked by default, you cannot run your own mail server out of your house and other server protocols like FTP are monitored or blocked outright. Properly configured, XP is/was a stable, reliable and reasonably safe version of Windows.

I agree with everything Vern said, except for his thesis that this was in some way “not a fair test.” I have a problem with that characterization only because it wasn’t a test of fairness. It was a test of reality – or perhaps a test of yesterday’s reality vs today’s reality. And, of course, everything Vern noted about the way he would spend his first hour with any new Windows XP machine was the reason I created the ShieldsUP! port probing facility... It was precisely because these early Windows machines were a disaster on the Internet.

That said, I also agree with his that it would, indeed, be interesting to place a currently patched Windows 10 or 11 machine directly on the ‘Net to see how it would fare. Given that all Windows machines have a very competent application-driven firewall that’s up and running before the rest of the vulnerable networking behind it, I’d expect it to do well.

But in any event, Parker’s whole point was to get some sense for the malicious crap that’s circulating out on the Internet right now, at this very moment. We’re all so well insulated and shielded behind our NAT routers and firewalls that it’s possible to sort of forget just what’s out there constantly pounding away at our defenses. Those defenses are certainly not optional.

Jeff Smock

Subject: I offer you my very own First Law of Cloud Data Security

:

Forget about all the bluster and jazz hands the cloud services providers give us regarding the security of our data, here is the simple truth: "The security of cloud data is inversely proportional to its potential value as perceived by a hacker or rogue staff member"

The 50 Gigabyte Privacy Bomb

Since we began with a general theme of how AI, which is not yet even close to being intelligent, is being misapplied during these early days, I feel as though a security- and privacy-focused podcast ought to take note of the new “Recall” feature that will be part of the next generation ARM-based Windows 11 Copilot+ laptop PCs.

First of all ... Yes ... it does appear that ARM processors have finally come far enough along to be able to carry the weight of Windows on their processors. And while having Windows on ARM will certainly create a new array of challenges, like, for example, the lack of specific hardware drivers that only exist for Intel kernels, in the more self-contained markets, where drivers are much less used, such as laptops, and where power consumption and battery life trumps pretty much any other consideration, it’s foreseeable that Windows may finally be able to find a home on ARM. Today, laptop and tablet form-factor machines containing Qualcomm Snapdragon ARM processors, running Windows 11 have been announced and are in some cases available for pre-order from Acer, Asus, Dell, HP, Lenovo, Microsoft and Samsung.

It’s also worth noting that Intel PCs will likely also be getting Copilot+ at some time in the future. But they will need to have a neural processing engine. Answering the question “What makes Copilot+ PCs unique?”, Microsoft writes:

Copilot+ PCs are a new class of Windows 11 PCs that are powered by a turbocharged neural processing unit (NPU)—a specialized computer chip for AI-intensive processes like real-time translations and image generation—that can perform more than 40 trillion operations per second (TOPS).

And later, Microsoft writes: *“We are partnering with Intel and AMD to bring Copilot+ PC experiences to PCs with their processors in the future.”*

Okay, so what is Recall? Microsoft explains:

You can use Recall on Copilot+ PCs to find the content you have viewed on your device. Recall is currently in preview status; during this phase, we will collect customer feedback, develop more controls for enterprise customers to manage and govern Recall data, and improve the overall experience for users. On devices that are not powered by a Snapdragon® X Series processor, installation of a Windows update will be required to run Recall.

Recall is currently optimized for select languages, including English, Chinese (simplified), French, German, Japanese, and Spanish. This means Recall is able to retrieve snapshots from your PC’s timeline based on more sophisticated searches in these languages. During the preview phase, we will enhance optimization for additional languages. Recall can also retrieve snapshots from your PC’s timeline based on text-to-text searches in more than 160 languages.

Okay. Fortunately, they then ask themselves “How does Recall work?” to which they reply:

Recall uses Copilot+ PC advanced processing capabilities to take images of your active screen every few seconds. The snapshots are encrypted and saved on your PC’s hard drive. You can use Recall to locate the content you have viewed on your PC using search or on a timeline bar that allows you to scroll through your snapshots. Once you find the snapshot that you were looking for in Recall, it will be analyzed and offer you options to interact with the content.

Recall will also enable you to open the snapshot in the original application in which it was created [whoa!! Really?], and, as Recall is refined over time, it will open the actual source document, website, or email in a screenshot. This functionality will be improved during Recall’s preview phase.

Copilot+ PC storage size determines the number of snapshots that Recall can take and store. The minimum hard drive space needed to run Recall is 256 GB, and 50 GB of space must be available. The default allocation for Recall on a device with 256 GB will be 25 GB, which can store approximately 3 months of snapshots. You can increase the storage allocation for Recall in your PC Settings. Old snapshots will be deleted once you use your allocated storage, allowing new ones to be stored.

What privacy controls does Recall offer?

Recall is a key part of what makes Copilot+ PCs special, and Microsoft built privacy into Recall’s design from the ground up. On Copilot+ PCs powered by a Snapdragon® X Series processor, you will see the Recall taskbar icon after you first activate your device. You can use that icon to open Recall’s settings and make choices about what snapshots Recall collects and stores on your device. You can limit which snapshots Recall collects; for example, you can select specific apps or websites visited in a supported browser to filter out of your snapshots. In addition, you can pause snapshots on demand from the Recall icon in the system tray, clear some or all snapshots that have been stored, or delete all the snapshots from your device.

Recall also does not take snapshots of certain kinds of content, including InPrivate web browsing sessions in Microsoft Edge. It treats material protected with digital rights management (DRM) similarly; like other Windows apps such as the Snipping Tool, Recall will not store DRM content.

Note that Recall does not perform content moderation. It will not hide information such as passwords or financial account numbers. That data may be in snapshots that are stored on your device, especially when sites do not follow standard internet protocols like cloaking password entry.

Okay.

So we’re rolling toward an entirely new capability for Windows PCs, where we’ll be able to store data which I presume is somehow indexed first, then encrypted for storage and later access. And, unless otherwise instructed and proscribed, this system is indiscriminately taking snapshots of our PC screen every few seconds and is, by Microsoft’s own admission, potentially capturing and saving for later retrieval, financial account numbers, monetary balances, contract language, proprietary corporate memos and communications, who knows what private things we’d really

rather never have recorded, or whatever else the user might assume will never go any further. This is where our much beloved and overworked phrase "what could possibly go wrong?" comes to mind.

Does anyone not imagine for an instant that having searchable access to the previous 90 days of a PC's screen might be hugely interesting to all manner of both legal and illegal investigators? Corporate espionage is a very real thing. China is moving their enterprises away from Windows as rapidly as they can, but you have to know that cyberattackers, many of the most skillful and persistent who seem to be persistently based in China, must be besides themselves with delight over this new prospect that we decadent capitalists in the West are going to start having our PCs recording everything that's displayed on their screens! What a great idea! If history teaches us anything it's that we have still not figured out how to keep a secret... and especially not Microsoft!

So what Microsoft is proposing to plant inside all next-generation PCs is tantamount to a 50 gigabyte privacy bomb. Maybe it will never go off... but it will certainly be sitting there trying to.

And just ask yourself whether law enforcement and intelligence agencies don't also think this sounds like a terrific idea? Oh, you betcha! With great power comes great responsibility. And here, clearly, there's much to go wrong. Microsoft understands this perception and so asks:

How is your data protected when using Recall? They explain:

Recall snapshots are kept on Copilot+ PCs themselves, on the local hard disk, and are protected using data encryption on your device and (if you have Windows 11 Pro or an enterprise Windows 11 SKU) BitLocker. Recall screenshots are only linked to a specific user profile and Recall does not share them with other users, make them available for Microsoft to view, or use them for targeting advertisements. Screenshots are only available to the person whose profile was used to sign in to the device. If two people share a device with different profiles they will not be able to access each other's screenshots. If they use the same profile to sign-in to the device then they will share a screenshot history. Otherwise, Recall screenshots are not available to other users or accessed by other applications or services.

Okay. So that's what Microsoft had to say. The guys from ArsTechnica watched Microsoft's presentation last Monday and gave their write up an impressively factual and neutral headline "New Windows AI feature records everything you've done on your PC" and then "Recall uses AI features "to take images of your active screen every few seconds." They wrote:

At a Build conference event on Monday, Microsoft revealed a new AI-powered feature called "Recall" for Copilot+ PCs that will allow Windows 11 users to search and retrieve their past activities on their PC. To make it work, Recall records everything users do on their PC, including activities in apps, communications in live meetings, and websites visited for research. Despite encryption and local storage, the new feature raises privacy concerns for certain Windows users.

Microsoft says on its website: "Recall uses Copilot+ PC advanced processing capabilities to take images of your active screen every few seconds. The snapshots are encrypted and saved on your PC's hard drive. You can use Recall to locate the content you have viewed on your PC

using search or on a timeline bar that allows you to scroll through your snapshots.”

By performing a Recall action, users can access a snapshot from a specific time period, providing context for the event or moment they are searching for. It also allows users to search through teleconference meetings they've participated in and videos watched using an AI-powered feature that transcribes and translates speech.

At first glance, the Recall feature seems like it may set the stage for potential gross violations of user privacy. Despite reassurances from Microsoft, that impression persists for second and third glances as well. For example, someone with access to your Windows account could potentially use Recall to see everything you've been doing recently on your PC, which might extend beyond the embarrassing implications of pornography viewing and actually threaten the lives of journalists or perceived enemies of the state.

In other words, this puts examining someone's web browser history to shame. How quaint! Ars continues:

Despite the privacy concerns, Microsoft says that the Recall index remains local and private on-device, encrypted in a way that is linked to a particular user's account. Microsoft says: "Recall screenshots are only linked to a specific user profile and Recall does not share them with other users, make them available for Microsoft to view, or use them for targeting advertisements. Screenshots are only available to the person whose profile was used to sign in to the device.”

Users can pause, stop, or delete captured content and can exclude specific apps or websites. Recall won't take snapshots of InPrivate web browsing sessions in Microsoft Edge or DRM-protected content. However, Recall won't actively hide sensitive information like passwords and financial account numbers that appear on-screen.

Microsoft previously explored a somewhat similar functionality with the Timeline feature in Windows 10, which the company discontinued in 2021, but it didn't take continuous snapshots. Recall also shares some obvious similarities to Rewind, a third-party app for Mac we covered in 2022 that logs user activities for later playback.

As you might imagine, all this snapshot recording comes at a hardware penalty. To use Recall, users will need to purchase one of the new "Copilot Plus PCs" powered by Qualcomm's Snapdragon X Elite chips, which include the necessary neural processing unit (NPU). There are also minimum storage requirements for running Recall, with a minimum of 256GB of hard drive space and 50GB of available space. The default allocation for Recall on a 256GB device is 25GB, which can store approximately three months of snapshots. Users can adjust the allocation in their PC settings, with old snapshots being deleted once the allocated storage is full.

As far as availability goes, Microsoft says that Recall is still undergoing testing. Microsoft says on its website: "Recall is currently in preview status. During this phase, we will collect customer feedback, develop more controls for enterprise customers to manage and govern Recall data, and improve the overall experience for users.”

Note that the amount of storage Recall uses scales upward with the size of the system's mass storage. It will take 25 gigs when 256 is available. 75 gig on a 512 gig drive, and 150 gigabytes from a system with 1 terabyte of primary mass storage. So, presumably, the more storage the system can commandeer, the further it's possible to scroll back through the system's display history.

Okay, while trying to be objective about this, the first question that leaps into the foreground for me is whether anyone actually needs or wants this? Is this a big previously unappreciated problem that everyone has?

Okay. Objectivity. First of all, compared to the static contents of a hard drive, Recall would be a gold mine of information about the past 90 days of someone's life, as viewed through their computer activities. And more than ever before, people's entire lives, and their private lives, are reflected in what's shown on the screens of their computers. Maybe that makes scrolling back through their recorded lives compelling?

But we know from Microsoft that it will be snapping video conference content on the fly, the Windows Signal app that goes to extremes to protect the content of its chats would be captured; as would email screens and nearly everything that happens on a PC. And, of course, that's the point. But the vast majority of that content will not have been stored on the machine's hard drive, until now. So the presence of Recall clearly introduces a new liability. And that's what **everyone** sees as a potential for creating havoc where none existed before.

So the question, it seems to me, is whether the new value that's created and returned by Recall's scrolling usage history justifies whatever risk might be created by retaining that data.

How useful will having this be? I've tried to imagine an instance where I wished I could look back in time at my computer screen. I suppose I don't feel the need since I've never had the capability. So if I knew I could scroll my computer's screens back in time I suppose it might be an interesting curiosity. But it really doesn't feel like a feature I've been needing and missing until now. I suppose an analogy would be that the world had **no** idea what it was missing before social media came along. And hasn't that been a huge boon to mankind? Now, unfortunately, we seem unable to live without it. Perhaps this will be the same.

The bottom line is this: I think we're just going to need to live with this for a while. We're going to need to see whether this is a capability desperately searching for a need, or whether, once people get used to having this capability, they think "how did I ever live without this?"

However, one thing that is also absolutely objectively true, is that everyone will be carrying around a 50 gigabyte privacy bomb that they never had before. Maybe it will be worth the risk. Only time will tell.

Oh, and Simon Zerafa posted a Tweet from someone who has been poking into Recall's storage. [@Detective@mastodon.social](https://mastodon.social/@Detective@mastodon.social) wrote: "Can confirm that Recall data is indeed stored in a SQLite3 database. The folder it's in is fully accessible only by SYSTEM and the Administrators group. Attempting to access it as a normal user yields the usual "You don't currently have permission" error. Here's how the database is laid out for those curious, figured you might appreciate a few screenshots."

<https://mastodon.social/@detective/112513529733646088>

...

The screenshot shows the DB Browser for SQLite interface. The main window displays the database structure for a SQLite database. The 'Tables (20)' section is expanded, showing a list of tables with their names, types, and schemas. The 'Indices (22)' section is also expanded, showing a list of indices with their names, types, and schemas. The 'Triggers (1)' section is expanded, showing a list of triggers with their names, types, and schemas.

Name	Type	Schema
App	CREATE TABLE	"App" ("Id" INTEGER PRIMARY KEY, "WindowsAppId" TEXT UNIQUE NOT NULL COLLATE NOCASE)
AppDwellTime	CREATE TABLE	"AppDwellTime" ("Id" INTEGER PRIMARY KEY, "WindowsAppId" TEXT NOT NULL COLLATE NOCASE)
File	CREATE TABLE	"File" ("Id" INTEGER PRIMARY KEY, "Path" TEXT UNIQUE NOT NULL COLLATE NOCASE, "NextId" INTEGER NOT NULL)
IdTable	CREATE TABLE	"IdTable" ("Id" INTEGER PRIMARY KEY, "NextId" INTEGER NOT NULL)
ScreenRegion	CREATE TABLE	"ScreenRegion" ("Id" INTEGER PRIMARY KEY, "WindowCaptureId" INTEGER NOT NULL)
Topic	CREATE TABLE	"Topic" ("Id" INTEGER PRIMARY KEY, "Title" TEXT UNIQUE NOT NULL COLLATE NOCASE)
Web	CREATE TABLE	"Web" ("Id" INTEGER PRIMARY KEY, "Domain" TEXT NOT NULL COLLATE NOCASE, "Uri" TEXT NOT NULL COLLATE NOCASE)
WebDomainDwellTime	CREATE TABLE	"WebDomainDwellTime" ("Id" INTEGER PRIMARY KEY, "Domain" TEXT, "HourOfDay" INTEGER NOT NULL)
WindowCapture	CREATE TABLE	"WindowCapture" ("Id" INTEGER PRIMARY KEY, "Name" TEXT NOT NULL, "ImageToken" TEXT NOT NULL)
WindowCaptureAppRelation	CREATE TABLE	"WindowCaptureAppRelation" ("WindowCaptureId" INTEGER NOT NULL, "AppId" INTEGER NOT NULL)
WindowCaptureFileRelation	CREATE TABLE	"WindowCaptureFileRelation" ("WindowCaptureId" INTEGER NOT NULL, "FileId" INTEGER NOT NULL)
WindowCaptureTextIndex	CREATE TABLE	"WindowCaptureTextIndex" ("WindowCaptureId" INTEGER NOT NULL, "Text" TEXT NOT NULL)
WindowCaptureTextIndex_config	CREATE TABLE	"WindowCaptureTextIndex_config" ("k" PRIMARY KEY, "v") WITHOUT ROWID
WindowCaptureTextIndex_content	CREATE TABLE	"WindowCaptureTextIndex_content" ("id" INTEGER PRIMARY KEY, "c0", "c1", "c2")
WindowCaptureTextIndex_data	CREATE TABLE	"WindowCaptureTextIndex_data" ("id" INTEGER PRIMARY KEY, "block" BLOB)
WindowCaptureTextIndex_docsize	CREATE TABLE	"WindowCaptureTextIndex_docsize" ("id" INTEGER PRIMARY KEY, "sz" BLOB)
WindowCaptureTextIndex_idx	CREATE TABLE	"WindowCaptureTextIndex_idx" ("segid", "term", "pgno", PRIMARY KEY("segid", "term")) WITHOUT ROWID
WindowCaptureTopicRelation	CREATE TABLE	"WindowCaptureTopicRelation" ("WindowCaptureId" INTEGER NOT NULL, "TopicId" INTEGER NOT NULL)
WindowCaptureWebRelation	CREATE TABLE	"WindowCaptureWebRelation" ("WindowCaptureId" INTEGER NOT NULL, "WebId" INTEGER NOT NULL)
_MigrationMetadata	CREATE TABLE	"_MigrationMetadata" ("Id" INTEGER NOT NULL UNIQUE, "Version" INTEGER NOT NULL)
idx_app_name	CREATE INDEX	idx_app_name ON App(name)
idx_app_path	CREATE INDEX	idx_app_path ON App(Path)
idx_app_windowsappid	CREATE INDEX	idx_app_windowsappid ON App(WindowsAppId)
idx_appdwelltime_windowsappid_hour...	CREATE INDEX	idx_appdwelltime_windowsappid_hourstarttimestamp ON AppDwellTime(WindowsAppId, HourOfDay)
idx_file_extension	CREATE INDEX	idx_file_extension ON File(Extension)
idx_file_kind	CREATE INDEX	idx_file_kind ON File(Kind)
idx_file_name	CREATE INDEX	idx_file_name ON File(Name)
idx_file_path	CREATE INDEX	idx_file_path ON File(Path)
idx_file_type	CREATE INDEX	idx_file_type ON File(Type)
idx_screenregion_kind	CREATE INDEX	idx_screenregion_kind ON ScreenRegion(RegionKind)
idx_screenregion_windowcaptureid	CREATE INDEX	idx_screenregion_windowcaptureid ON ScreenRegion(WindowCaptureId)
idx_topic_title	CREATE INDEX	idx_topic_title ON Topic(Title)
idx_web_domain	CREATE INDEX	idx_web_domain ON Web(Domain)
idx_web_uri	CREATE INDEX	idx_web_uri ON Web(Uri)
idx_webdomaindwelltime_domain_ho...	CREATE INDEX	idx_webdomaindwelltime_domain_hourstarttimestamp ON WebDomainDwellTime(Domain, HourOfDay)
idx_windowcapture_isprocessed	CREATE INDEX	idx_windowcapture_isprocessed ON WindowCapture(IsProcessed)
idx_windowcapture_name_timestamp	CREATE INDEX	idx_windowcapture_name_timestamp ON WindowCapture(Name, TimeStamp)
idx_windowcapture_timestamp	CREATE INDEX	idx_windowcapture_timestamp ON WindowCapture(TimeStamp)
idx_windowcaptureapprelation_rel	CREATE INDEX	idx_windowcaptureapprelation_rel ON WindowCaptureAppRelation(AppId, WindowCaptureId)
idx_windowcapturefilerelation_rel	CREATE INDEX	idx_windowcapturefilerelation_rel ON WindowCaptureFileRelation(FileId, WindowCaptureId)
idx_windowcapturetopicrelation_rel	CREATE INDEX	idx_windowcapturetopicrelation_rel ON WindowCaptureTopicRelation(TopicId, WindowCaptureId)
idx_windowcapturewebrelation_rel	CREATE INDEX	idx_windowcapturewebrelation_rel ON WindowCaptureWebRelation(WebId, WindowCaptureId)
trigger_windowcapture_before_delete	CREATE TRIGGER	trigger_windowcapture_before_delete BEFORE DELETE ON "WindowCapture" BEGIN

I captured a screenshot of the SQL table layout for the show notes. So at least anyone wishing to perform some forensics on a Privacy Bomb PCs won't have any trouble opening the files they've stolen.

