

Security Now! #959 - 01-30-24

Stamos on "Microsoft Security"

This week on Security Now!

What changes will the EU's soon-to-be-in-force Digital Markets Act be bringing to Apple's traditional iOS policies? What OS is ransomware unable to infect? What has HP done now with their printer ink policy? How many stolen user database records will fit in 12 terabytes? Can't you just delete that incriminating chat stream? Did Mercedes-Benz leave their doors unlocked? What's the latest on ransom payments rates? And after entertaining some questions from our terrific listeners and a long-awaited announcement from me, we're going to take a look at Alex Stamos' reaction to Microsoft's most recent security incident response.

How to get Hipsters to Obey those "Keep Out" Warnings...



Security News

iOS to allow native Chromium and Firefox engines.

In perfect timing following from last week's discussion of Mozilla's complaints against Apple, Google and Microsoft we have the news, initially reported by The Verge that Apple will be changing their browser engine policy for the European Union. Under their headline [Apple is finally allowing full versions of Chrome and Firefox to run on the iPhone](#) The Verge wrote:

With iOS 17.4, Apple is making a number of huge changes to the way its mobile operating system works in order to comply with new regulations in the EU. One of them is an important product shift: for the first time, Apple is going to allow alternative browser engines to run on iOS — but only for users in the EU.

Since the beginning of the App Store, Apple has allowed lots of browsers but only one browser engine: WebKit. WebKit is the technology that underpins Safari, but it's far from the only engine on the market. Google's Chrome is based on an engine called Blink, which is also part of the overall Chromium project that is used by most other browsers on the market. Edge, Brave, Arc, Opera, and many others all use Chromium and Blink. Mozilla's Firefox runs on its own engine, called Gecko.

On iOS, though, all those browsers have been forced to run on WebKit instead, which means many features and extensions simply don't work anymore. That changes with iOS 17.4 — anyone building a browser, or building an in-app browser for their app, can use a non-WebKit engine if they wish. Each developer will have to be authorized by Apple to switch engines "after meeting specific criteria and committing to a number of ongoing privacy and security mitigations," Apple said in a release announcing the change, at which point they'll get access to features like Passkeys and multiprocessing. Apple's also adding a new choice screen to Safari so that when you first open the browser, you'll be able to choose a different default if you want.

Apple is only doing this because it is required to by the EU's new Digital Markets Act (DMA), which stipulates, among other things, that users should be allowed to uninstall preinstalled apps — including web browsers — that "steer them to the products and services of the gatekeeper." In this case, iOS is the gatekeeper, and WebKit and Safari are Apple's products and services. (The same section of the DMA also means Microsoft has to let people disable Bing web search and uninstall Edge, and it will cause other changes, too.)

Even in its release announcing the new features, Apple makes clear that it is not pleased: "This change is a result of the DMA's requirements, and means that EU users will be confronted with a list of default browsers before they have the opportunity to understand the options available to them," the company says. "The screen also interrupts EU users' experience the first time they open Safari intending to navigate to a webpage." Apple's argument for the App Store has always amounted to: only Apple can provide a good, safe, happy user experience on the iPhone. Regulators don't see it that way. And Apple's furious about it.

Again, these changes are only for iPhone users in the EU. Apple says it allows European users to travel without breaking their browser engine but will make sure only accounts belonging to people who live in the EU get these new engines. Elsewhere in the world, you'll still be getting WebKit Chrome and WebKit. Apple argues, without merit or evidence, that these other engines pose a security and performance risk and that only WebKit is truly optimized and safe for iPhone users.

In the EU, we're likely to see these revamped browsers in the App Store as soon as iOS 17.4 drops in March: Google, for one, has been working on a non-WebKit version of Chrome for at least a year. European users are about to get a serious browser war on their iPhones.

Before The Verge's piece ended, when they suggested that users would see new browsers in March, I was shaking my head since porting an entire browser engine to a new OS is no small task. Putting a browser skin over the WebKit engine, which is what everyone else has done until now, is entirely different from running Chromium's Blink or Firefox's Gecko engines under iOS.

And if this will only be allowed for users having accounts based in the EU, I'm wondering why anyone really cares? The branding skin is all anyone sees. As I mentioned last week, I use Firefox on my iPhone and multiple iPads. To me it looks and acts like Firefox and I appreciate its various features. For example, it's possible for me to login to my Mozilla account and then sync tabs across my various devices. Just now when I picked it up to double check that, it asked me whether I wanted to pick up editing this podcast's Google document.

I presume that compatibility with the EU's Digital Markets Act (DMA) will mean that the sort of link stealing behavior Mozilla was complaining about, which pulls users back to Safari will disappear. If Mozilla is resource constrained I'd prefer to have them keep their focus on the desktop where it matters and ignore this distraction which, after all, only applies to EU territory.

Before I share the big worry that this story prompted in me, I want to share a bit more about this. The day after The Verge's coverage, MacRumors followed up with additional coverage under the headline "[Apple Further Explains iOS 17.4's New Default Browser Prompt in EU](#)" They wrote:

After updating to iOS 17.4, which is currently in beta, iPhone users in the EU will be prompted to choose a default web browser when they first open Safari. In an email today, Apple shared additional details about how this process will work.

Apple said iPhone users in the EU will be presented with a list of the 12 most popular web browsers from their country's local App Store at the time, and noted that the options will be shown in random order for every user.

Apple shared an alphabetical list of the browsers that will currently be shown in every EU country. It is a very long list, so we have elected to highlight browsers that will be shown in France, Germany, Italy, and Spain as examples.

- *France: Aloha, Brave, Chrome, DuckDuckGo, Ecosia, Edge, Firefox, Onion Browser, Opera, Private Browser Deluxe, Qwant, and Safari*
- *Germany: Aloha, Brave, Chrome, DuckDuckGo, Ecosia, Edge, Firefox, Ivanti Web@Work, Onion Browser, Opera, Safari, and You.com AI Search Assistant*
- *Italy: Aloha, Brave, Chrome, DuckDuckGo, Ecosia, Edge, Firefox, Ivanti Web@Work, Onion Browser, Opera, Safari, and You.com AI Search Assistant*

- *Spain: Aloha, Brave, Chrome, DuckDuckGo, Ecosia, Edge, Firefox, Onion Browser, Opera, Safari, Vivaldi, and You.com AI Search Assistant*

There are 23 other countries in the EU that this change applies to though, of course, this no longer includes the UK, which withdrew from the EU in 2020.

It has been possible to change an iPhone's default web browser through the Settings app since iOS 14. Apple has now gone a step further and added the default browser prompt in Safari to comply with new regulations under the EU's Digital Markets Act.

In the EU, iOS 17.4 also allows web browsers to use web engines other than Apple's WebKit.

Apple said iOS 17.4 will be released to the public in March.

So this clarifies a few things. If all of these browsers are currently the top 12 most popular in each EU country's regional Apple App store then they're all currently using simple skins over WebKit since that's all that's been possible until now. That means that users will likely not initially be changing the underlying engine. They will be prompted to proactively pick a skin. And I imagine that a great many will opt for Chrome since that's the browser that dominates the desktop. Then after that, if specific browser vendors see some reason to invest in porting the Chromium or Mozilla engines over to iOS then that might happen.

The one reason I can see for Google to invest in a full Chrome port is that they badly need their Privacy Sandbox API to run everywhere. If tracking is abandoned and eventually outlawed, the Privacy Sandbox's technology will be needed to continue delivering interest-targeted advertising. It's unclear what Apple's plans are for WebKit and whether Apple has any interest in following Google. So if WebKit doesn't support the Privacy Sandbox, true Chrome on iOS will need to.

I mentioned a big worry that this announcement triggered in me: What we see here is Apple capitulating to the demands imposed by a regional legal framework. I suppose they have no choice if they wish to continue operating in the EU. The Verge made it clear that Apple is furious about this... but capitulating they are. And this reminded me of the pending European eIDAS 2.0 legislation which intends to compel the world's web browsers and operating systems to accept, without recourse, any and all root certificates that the EU may choose to require browsers to honor. The EU's Digital Markets Act is about competition and antitrust. Its aim is to water down Apple's vise grip on its traditional heavy-handed business practices. So it's not directly comparable to the eIDAS 2.0 legislation... but I get a sinking feeling about this.

Under the heading "Dodged a Bullet"

Under the heading "Dodged a Bullet" we have the news that the 3rd largest bank in the world, which is China's ICBC, was hit with a ransomware attack which got into and would have compromised their entire network... except for one tiny detail. Believe it or not, in this year of our Lord, 2024, the critical currency trading network used by China's ICBC bank was being run by a Novell Netware server. Yes, in 2024. If it's not broken... Anyway, a Novell Netware server was entirely alien to the ransomware which had no idea how to infect the server or get up to any other mischief. Consequently, the bank just shrugged off the attack, cleaned some modern workstations that had succumbed, and got on with their day in this Year of the Dragon.

HP back in the doghouse over “anti-virus” printer bricking

The news that many of our listeners forwarded to me recently was that HP has once again been bricking their printers when those printers are found to contain 3rd-party Ink. Here's what 9to5mac wrote under the headline [“Third-party ink cartridges brick HP printers after ‘anti-virus’ update”](#):

*HP is pushing over-the-air firmware updates to its printers, bricking them if they are using third-party ink cartridges. But don't worry, it's not a money-grab, says the company – it's just trying to protect you from the well-known risk of **viruses** embedded in ink cartridges ...*

HP has long been known for sketchy practices in its attempt to turn ink purchases into a subscription service. If you cancel a subscription, for example, the company will immediately stop the printer using the ink you've already paid for.

HP's CEO Enrique Lores somehow managed to keep a straight face while explaining to CNBC that the company was only trying to protect users from viruses which might be embedded into aftermarket ink cartridges, saying: “It can create issues [where] the printers stop working because the inks have not been designed to be used in our printers, to then create security issues. We have seen that you can embed viruses in the cartridges, and through the cartridge, go to the printer; from the printer, go to the network.”

ArsTechnica asked several security experts whether this could happen, and they said this is so far out-there, it would have to be a nation-state attack on a specific individual. Three expert replies were:

“Purely from a threat-modeling perspective, I'm skeptical – unless it's a nation-state doing a tailored attack.”

“As someone who works for a different inkjet print company – I'd say it's pretty terrible engine design if you could maliciously craft a cartridge to contain a virus. The amount of information which needs to be stored on the cartridge is fairly small. If the data is not in the format you expect – reject it as invalid. [HP is known to be quite good at this!]”

“I've seen and done some truly wacky hardware stuff in my life, including hiding data in SPD EEPROMs on memory DIMMs (and replacing them with microcontrollers for similar shenanigans), so believe me when I say that his claim is wildly implausible even in a lab setting, let alone in the wild, and let alone at any scale that impacts businesses or individuals rather than selected political actors.”

HP is facing a class action lawsuit for deploying the bricking code without informing printer buyers of its intention to do so. The suit explains:

“This is a class action brought against HP, Inc., for requiring consumers who had purchased certain brands of printers to use only HP-branded replacement ink cartridges, rather than purchasing ink replacements from its competitors. HP accomplished this through firmware updates it distributed electronically to all registered owners of the printers [...] which effectively disabled the printer if the user installed a replacement ink cartridge that was not HP-branded. In the same time period, HP raised prices on the HP-branded replacement ink cartridges. In effect, HP used the software update to create a monopoly in the aftermarket for replacement cartridges, permitting it to raise prices without fear of being undercut by competitors.”

A super-massive leak database was discovered online

It's being called MOAB because it's the Mother Of All Breaches, totally an astounding 12TB of data contained within 26 billion (with a 'B') database records.

The supermassive leak contains data from numerous previous breaches, including data from LinkedIn, Twitter, Weibo, Tencent, and other platforms' user data, making it the largest collection of stolen user data ever discovered. The data includes records from thousands of meticulously compiled and reindexed leaks, breaches, and privately sold databases.

Bob Dyachenko – he's the guy we've mentioned before who appears to specialize in discovering open and exposed data online – was behind this discovery. Although the owner of the database was initially unknown Leak-Lookup, a data breach search engine, said it was the holder of the leaked dataset. The platform posted a message on X, saying the problem behind the leak was a "firewall misconfiguration," which was fixed. Yeah... whoops.

While the leaked dataset contains mostly information from past data breaches, it almost certainly holds new data that has never been published before. For example, the Cybernews data leak checker, which relies on data from all major data leaks, contains information from over 2,500 data breaches with 15 billion records.

But the MOAB contains 26 billion records over 3,800 folders, with each folder corresponding to a separate data breach. While this doesn't mean that the difference between the two automatically translates to previously unpublished data, billions of new records point to a very high probability, the MOAB contains never seen before information.

Researchers believe that the owner of the MOAB has a vested interest in storing large amounts of data and, therefore, could be a malicious actor, data broker, or some service that works with large amounts of data.

The researchers said "The dataset is extremely dangerous as threat actors could leverage the aggregated data for a wide range of attacks, including identity theft, sophisticated phishing schemes, targeted cyberattacks, and unauthorized access to personal and sensitive accounts."

While the team identified over 26 billion records, duplicates are also likely. However, the leaked data contains far more information than just credentials – most of the exposed data is sensitive and, therefore, valuable for malicious actors.

A quick run through the data tree reveals an astoundingly large number of records compiled from previous breaches. The largest number of records, 1.4 billion, comes from Tencent QQ, a Chinese instant messaging app.

However, there are 504M from Weibo, 360M from MySpace, 281M from Twitter, 258M from Deezer, 251M from LinkedIn, 220M from AdultFriendFinder, 153M from Adobe, 143M from Canva, 101M from VK, 86M from Daily Motion, 69M from Dropbox, 41M from Telegram, and on and on.

In addition to data on individuals, the leak also includes records of various government organizations in the US, Brazil, Germany, Philippines, Turkey, and other countries.

If anyone wonders where and how targeted credential stuffing attacks originate, one would need to look no further. The database contains names and addresses, very personal information, password hashes and in-the-clear eMail addresses.

The people who discovered this are understandably hyping it up a bit. This is not to say that it's not a seriously worrisome collection of potentially potent data... but we should keep in mind that it is a collection of data gathered from all previous data breaches. That means that it's aging and is no longer current. No one should **ever** have their security breached. But anyone who is still using "123456" as their single global password – which, fortunately, is quite difficult to do any longer – should not be surprised if their accounts are breached. And, really, no big database is required to do that!

New "Thou shall not delete those chats" rules

Federal investigators are warning companies which are under investigation that they may not and must not delete chats and that they must arrange to preserve conversations that have taken place via business collaboration and ephemeral messaging platforms. In dual coordinated press releases last Friday, the US Department of Justice and the US Federal Trade Commission announced updated language in their preservation letters and specifications —documents they send to companies under federal investigation. The new language updates evidence preservation procedures to cover modern tech stacks such as Slack, Microsoft Teams, and Signal.

Companies that receive subpoenas or other legal notifications must take steps to preserve chat logs and disappearing IM messages, and any who do not will be subject to obstruction of justice charges. The problem, of course, is that being charged with obstruction of justice might be better than revealing what they deliberately chose to delete.

The Deputy Assistant Attorney General of the Justice Department's Antitrust Division said:
"These updates to our legal process will ensure that neither opposing counsel nor their clients can feign ignorance when their clients or companies choose to conduct business through ephemeral messaging."

And this updated guidance comes as the DOJ faced difficulties pursuing its antitrust lawsuits against Google and Amazon. February last year, the DOJ accused Google of lying when it claimed it auto-suspended its chat auto-deletion feature. In addition, the DOJ claimed that for a period of four years, Google trained employees to delete internal chats and move conversations to off-the-record platforms because it anticipated facing antitrust litigation in the near future.

Later, in November last year, the FTC accused Amazon of deleting more than two years worth of internal Signal employee chats after the agency started a multi-state antitrust lawsuit.

I have a representative snippet of the DOJ's evidence-hiding complaint in their antitrust case against Google. The complaint reads:

15 The newly produced Chats reveal a company-wide culture of concealment coming from the
 16 very top, including CEO Sundar Pichai, who is a custodian in this case. In one Chat, Mr. Pichai began
 17 discussing a substantive topic, and then immediately wrote: *“also can we change the setting of this*
 18 *group to history off.”*¹ Then, nine seconds later, Mr. Pichai apparently attempted (unsuccessfully) to
 19 delete this incriminating message. (Byars Decl. Ex. 1, GOOG-PLAY5-000453593.) When asked
 20 under oath about the attempted deletion of the message, Mr. Pichai had no explanation, testifying “I
 21 definitely don’t know” and “I don’t recall.” (*Id.* Ex. 2, Pichai Dep. Tr. 195:7-12.)

22 Like Mr. Pichai, other key Google employees, including those in leadership roles, routinely
 23 opted to move from history-on rooms to history-off Chats to hold sensitive conversations, even though
 24 they knew they were subject to legal holds. Indeed, they did so *even when discussing topics they*
 25 *knew were covered by the litigation holds in order to avoid leaving a record that could be produced*
 26 *in litigation.* As the examples below make clear, Google destroyed innumerable Chats with the intent
 27 to deprive Plaintiffs and other litigants of the use of these documents in litigation.

So the federal government is making it very clear that digital recordings of private conversations may not be deleted from the moment of notification of pending litigation. If executives wish to hold private off-the-record conversations they’re going to need to do it the old fashioned way, face to face in a private setting with no one recording.

Mercedes-Benz source code and MUCH more!

Here’s one that’ll really ruin your day (at least if you’re responsible for software security at Mercedes-Benz: Get this... Mercedes-Benz accidentally exposed a trove of internal data and more by leaving a private key online that gave “unrestricted access” to the company’s source code. And that key was there for more than 90 days before it was discovered and responsibly reported by the a co-founder and the chief technology officer of London-based RedHunt Labs.

What RedHunt discovered during a routine Internet data scan earlier this month was a Mercedes employee’s authentication token sitting in a public GitHub repository. This token served as an alternative to using a password for authenticating to GitHub. As such it would grant anyone full access to Mercedes’s GitHub Enterprise Server which would, in turn allow the download of the company’s entire collection of private source code repositories.

RedHunt said that the GitHub token gave ‘unrestricted’ and ‘unmonitored’ access to the entire source code hosted at the internal GitHub Enterprise Server. The repositories include a large amount of intellectual property – and how! Get this: Connection strings, cloud access keys, blueprints, design documents, [single sign-on] passwords, API Keys, and other critical internal information. RedHunt provided evidence that the exposed repositories contained Microsoft Azure and Amazon Web Services (AWS) keys, a SQL database, and Mercedes source code. It’s not known if any customer data was contained within the repositories.

Last Monday TechCrunch, serving as a middleman for RedHunt disclosed the security issue to Mercedes. On Wednesday, a Mercedes spokesperson confirmed that the company (quote) *"revoked the respective API token and removed the public repository immediately."*

Quote: *"We can confirm that internal source code was published on a public GitHub repository by human error. The security of our organization, products, and services is one of our top priorities. We will continue to analyze this case according to our normal processes. Depending on this, we implement remedial measures."*


Since the exposed key was published last September, it sat there, through October, November, December and most of January this year. What's not known is whether anyone besides RedHunt may have discovered and taken advantage of the exposed key. Mercedes declined to say whether it is aware of any third-party access to the exposed data or whether the company has the technical ability, such as access logs, to determine if there was any improper access to its data repositories. The spokesperson cited unspecified security reasons.

We've previously covered that GitHub has begun proactively scanning repositories for these sorts of inadvertent disclosures. Our software and intellectual property management systems have become so complex and interdependent that they have also become brittle to these sorts of human errors. I'm sure we'll be seeing more of these sorts of mistakes in the future.

Fewer ransoms are being paid

In some good news, the number of ransomware victims who opted to pay ransoms fell to an all-time low at the end of last year. The cybersecurity firm Coveware estimates that only 29% of victims paid ransoms during the 4th quarter of 2023. This is down from the 85% who were choosing to pay back in the 1st quarter of 2019 when the company began tracking the stat.

All Ransomware Payment Resolution Rates



Coveware attributes the fall to improved data backup and recovery strategies in corporate environments and companies getting smarter about not trusting empty promises made by ransomware groups.

Closing the Loop

Conradical / @conradstorz

@SGgrc – Steve, please take a deeper dive into the technology behind verified camera images. My gut reaction is you've overlooked something because public key cryptography should allow the images to be verifiable and unmodifiable.

There are many amazing things that public key crypto can do. And I'm deeply enamored of them. But in the case of the verified camera images, you must ask yourself what could a camera contain that cannot be copied by someone who gets their clutches on such a camera? I contend that anything a camera can know someone can find a way to pry out of that camera to duplicate whatever it knows. And in doing so, duplicate its ability to make a strong assertion of an image's origin. In other words, this entire system depends upon the camera, which is out in public, being able to keep a secret. And everything we know tells us that's almost certainly not possible if someone is sufficiently motivated.

The most common application we have today of public key crypto is the dynamic creation of secure connections to remote servers, where those servers are asserting their identity. Only one thing allows that system to work, which is that those servers are not accessible to others. If they were, the secrets they're protecting could be stolen and others could impersonate them.

That's the difference in the security model of the camera vs a remote server. It's the remoteness of the server that allows it to protect its secrets; the fact that it can only be accessed through a carefully managed TCP connection. The infamous HeartBleed vulnerability demonstrated what would happen if server secrets could be accessed through a side channel. The server's secrets would be compromised.

So it's not that public key crypto doesn't still require secrets. It does. It's just that only one side of the transaction needs to be able to keep something secret. Unfortunately, when a camera is signing the pictures it takes, it's the private key that the camera is using to perform the signing that needs to be kept secret. Building a state of the art hardware security module (HSM) into the camera – which I'm sure it has – will likely make it as difficult as possible to extract the HSM's key. The unanswered question is, will it be difficult enough?

Jg1212G / @Jg1212G

Hi Steve, I was just listening to Security Now and got hooked into the \$15 per week flashlight story. I had to look into it. I found it on the play store and followed the link to their website. <https://simplemobiletools.com/index.html> I thought, very strange, the site says open source and add free. So I clicked on the Github link at the bottom. <https://github.com/SimpleMobileTools/Simple-Flashlight> Sure enough it is open source. So I looked at the developer's page on Github: <https://github.com/tibbi> Wow, his graph shows he was extremely active up until the end of October 2023 then completely stopped. That is so strange. I would really like to know what happened to him. If you hear any news please let us know. I love a good mystery! Thanks, Jason

Ask and you shall receive... our listener "megascrapper" brings an end to the mystery:

megascrapper / @megascrapper

Hi Steve, I'd like to follow up on last week's listener feedback about the absurd subscription prices for a flashlight app. I was made aware of the entire Simple Mobile Tools suite (which includes Simple Flashlight) after watching a video by Brodie Robertson:

<https://www.youtube.com/watch?v=OGr4H7QLVRs>

Unfortunately what happened with Simple Flashlight was exactly what you presumed in your reply to that listener: With very little notice, the owner/primary maintainer of the app sold the entire suite to an Israeli publisher, ZipoApps, which is notorious for the practice of acquiring existing apps and slapping on an outrageously expensive subscription plan.

But not all hope is lost. The entire suite is open source (GPLv3 licensed) and one of the maintainers already forked it under a project called Fossify, including Simple Flashlight:

<https://github.com/FossifyOrg/Flashlight>

It seems to be still in early development and I can't find the app on Google Play Store, but keep an eye out when it gets released. Thank you very much for your work. Looking forward to 999 and beyond.

So, thank you “megascrapper” for your follow-up on this and for the confirmation that this is the sort of thing that happens with highly popular apps in the Google Play Store. The description for the YouTube video he linked to said: *“I was a fan of the Simple Mobile Tools suite for a really long time and then out of nowhere the developer Tibor Kaputa just sold the entire project and ran away with the bag, luckily not all hope is lost.”* It was certainly Tibor’s right to do whatever he wanted to with his own intellectual property. It’s clear that since the entire project is open source, it was his project’s developer keys that was of actual value because they allowed its purchaser to take over the official popular app and then upgrade it into the existing channel.

Jon Dagle / @jondagle

Hey Steve, In response to the Flashlight app story. (1) to access the flashlight brightness, swipe down from the top right to get Control Center; long-press the flashlight icon. Solved!

I tried it on my iPhone and I was amazed! It has four levels of brightness. I never knew. This is super useful to me since the flashlight defaults to a setting that should be labeled “Visible from Orbit.” All I want to do is read the menu in a darkened restaurant, I’m not trying to signal aliens for pickup. So I immediately set it to its lowest level, which will be much more appropriate in the future... and it won’t blind my fellow diners if I inadvertently pass its laser beacon across their vision!

Before I get to Jon’s second point, I just want to mention something that’s quite annoying: I have this dull sense that there is vastly more available from today’s iPhone than I’m aware of. But how would I ever discover this on my own? I guess I just have to sit around and press on everything to see if anything happens? The original concept of the graphical user interface was that it was discoverable. That’s what was so cool about having nested drop-down menus running along the top of the screen. Unlike the text command interfaces that preceded them, you could

sit down and run the mouse around the screen and find everything that you might need. Today, it's easy to do the basic things with a phone, but it's annoying to imagine just how much more remains hidden behind the need to click your heels together three times. How would you ever know?

Jon's second comment was:

(2) I recently ran across another long-time trusted app that was sold, It's the super excellent Network Toolbox (NET-Toolbox) on iOS. I think it's been mentioned on SN in the past; it has a host of powerful networking tools. But, the long time developer sold the app sometime late in 2023. When I first opened the app after recently resetting all settings, I got the "Network Toolbox wants to track you across websites...." alert.

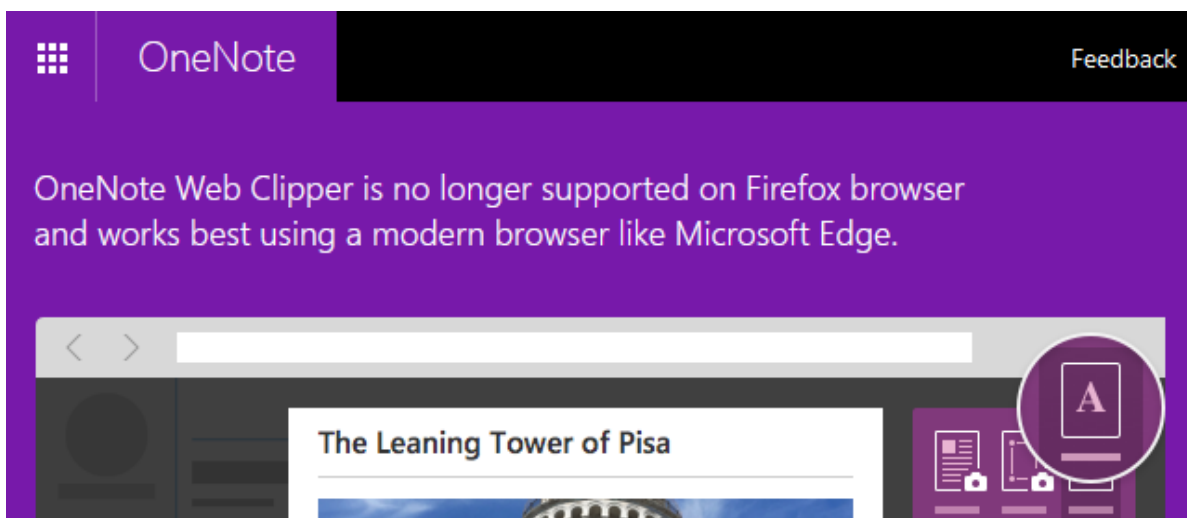
The sale/transfer was all silent as far as I'm aware. Considering the app has a lot of sensitive functions, the trustworthiness the developer is rather important. So beware!

<https://apps.apple.com/us/app/network-toolbox-net-security/id651691453>

All of these stories make me think that perhaps Google and Apple really ought to consider adding proactive notification to apps when their ownership changes hands. I've never participated in such a transfer, so I don't have any clear sense for whether a developer might simply turn over their entire online identity to a third-party purchaser, or where there's some more formal and controlled process for doing so. But if it's knowable to Google or Apple, it would seem useful to add a bit of friction and visibility to this otherwise very slippery and transparent process. This is a lot about trusting the publisher. So if that publisher changes, it seems to me that those being asked to trust someone new should know.

Brian Doyle / @bdoyle159 Brings us another example that Mozilla can add to their growing list of "good luck with that" grievances against the tactics being employed by those who wish to use their own platforms to their competitive advantage. Brian writes:

Hi Steve, I came across this message while looking for a way to save full web pages into OneNote, and had to laugh at Microsoft implying that Firefox is not a 'modern' browser. Thought you might enjoy. Here is the original site <https://www.onenote.com/clipper>



Out of curiosity I brought the same page up under Chrome. Naturally, with Chrome's market share, there was no way that Microsoft was going to snub nearly all of the world's browser users who are, by the way, also not Edge users.

ShipRkt / @Shiprkt (Shipwrecked)

Hello Steve, I hope you don't mind me sending you a message. Could you discuss on a future Security Now episode why "Credit Karma" is storing over 1GB of data on my iPhone. What on earth uses that much data for a credit app? Thank you for your time

First of all, I certainly do not mind receiving messages, which is why I go in search of them every week for the podcast. But neither do I have any idea why the Credit Karma app might be storing over 1GB of data on anyone's iPhone. One thought I had was to wonder whether this 1GB might include the app itself in that total? One of the sad trends we see is applications becoming increasingly and, in fact, obscenely bloated. They evidence no respect whatsoever for the users of these apps. I'm sure that few consumers are even aware of this, which is why there's little cost associated with being so careless with the consumption of other people's storage.

Anyway – old curmudgeon rant off – I would love to put this question to our listeners. I poked around briefly but I didn't find anything about Credit Karma's iOS app resource consumption. So if anything finds anything I'd be glad to share it.

Mark Guy / @SDTwitGuy

Mark Guy, whose Twitter handle is @SDTwitGuy appears to be a fan of the network. He said:

I heard your comment about staying on Windows 7 on the 1/23/24 podcast. My main system is Windows 7 Ultimate. They'll have to pry it out of my cold dead hands. LOL. It's stable, it runs perfectly. I subscribe to OPatch and still get updates for MS Security Essentials, plus I use Malwarebytes Premium. Never had any problems. Plus I know where everything is. I bought a used Windows 10 laptop and I can barely find anything. I also am an avid fan of Windows Media Center. Nothing else comes close to it's functionality. It's how I watch and record TV so I will never update my system.

I'm also a huge Sci-Fi fan and I LOVE that you and Leo talk about your fave Sci-Fi authors, books and series. Thank you!

I wanted to mention two things: First, I know that Mark and I are far from alone among the listeners of this podcast. Just like with that Novell Netware server, it's working so don't mess with it. And yes, at some point I'll rebuild my machine around Windows 10. Since I'm an MSDN developer I could still register a new machine as Windows 7. But I'm not totally insane. I'm typing this into a Win7 Pro workstation mostly because moving to Win10 would take a non-zero amount of time and like Mark and many of our other listeners, why bother when this 64-bit edition of Windows 7 is working just fine.

The second thing I wanted to mention follows from Mark comment: *"I'm also a huge Sci-Fi fan and I LOVE that you and Leo talk about your fave Sci-Fi authors, books and series. Thank you!"*

I've been intending to mention that after investing in about six of those Aeon 14 novels, the ones invariably featuring voluptuous, heavily armed female commandos on their covers, despite the fact that another hundred or so of those remained, I had finally reached my limit.

Following a number of recommendations, I gave the "Expeditionary Force" novels a try, but they just didn't grab me. They're written in a 1st-person narrative style and I kept waiting for something to happen. My trouble might be that they're a bit too realistic. Once you've read much of Peter Hamilton's work you're somewhat cut loose from the need for an excess of reality.

But in the meantime, Ryk Brown, the prodigious author of the Frontiers Saga series, had dropped a few more books in his third of five planned 15-book story arcs. Since we're up to book #10 in arc #3, we've passed the halfway point. I've turned a number of very close friends and family members onto this series and I have been unable to shake them loose. They want nothing to do with anything else. They just want more Ryk Brown. As we know, I've wandered around while waiting for more. I happily consumed the entire Silver Ships series following another recommendation from a listener. And, of course, some of those Aeon14 series series.

I'm bringing all of this up because Ryk Brown's writing style, his deep characterization, his perfect management of a large and growing number of very different and distinct characters, and the fact that you never need to wait long for some action, continues, after 40 books, to be absolutely enjoyable and gratifying. All of the books are available under Amazon's Kindle Unlimited plan and as Audio Books. Through the 19 years of this podcast we've shared our discoveries of many terrific books. For sheer solid entertainment value, I think this series deserves everyone's attention. So I just wanted to be sure that it's on everyone's radar. :)

Dizzle Von Dazzle / @SirDizzleDazzle

Quick question as you are an avid user of Windows 7. How do I continue to use websites that use HSTS. It's a new install on an oldish lenovo IdeaPad All in one. Is there away to update the SSL libraries as none of the update managers for different music production applications I own seem to work either. Keep up the amazing work on SpinRite and here's to episode 999

I'm having no trouble with Win7 and HSTS sites, such as GRC which was one of the earliest to adopt HSTS and permanent registration in Chrome. Under my Win7 setup notes I have a subdirectory named: "Before registering or installing Win7 updates" and that subdirectory contains three specific Microsoft updates: There's an SHA256 Update, a "Servicing Stack Update" and an update "KB3102810". From my notes it appears that you should find those three individual stand alone updates and install them in that order. Then you can successfully bring Windows 7 current and all should be well!

A listener who asked to remain anonymous...

Steve - my company is switching to Bitwarden from LastPass as a result of me raising the issue a year ago, which is a result of your discussion on the podcast. Please keep this anonymous if you mention it on the SN podcast. My question is, can I get a readout from you

on the advisability of adding TOTP codes/secrets into Bitwarden so that it can fill in the field on sites you're logging in to? Personally it gives me a 'Gibsonian response' and feels like all your eggs are in one basket if you do that. What do you think? Regards. Long time listener etc.

We've mentioned this before but it's worth mentioning it again, and I know that Leo concurs since I heard him say the same thing on other podcasts. But I've also had some time to think about this and to perhaps mellow about it a bit. I understand the convenience. But it's also a case of a classic trade off between convenience and maximum security. My truest feeling is that the actual risk from having "all the eggs in one basket" is likely less significant than the benefit that comes from ease of use. So if, for example, it was ever a matter of not registering and user a TOTP one-time-password due to the inconvenience of needing to use a second authentication device — last week Paul Thurrott was explaining that his wife has absolutely zero interest in anything that gets in her way — then yes, it would be better to have Bitwarden able to automatically fill-in the OTP field than to not use time-based multi-factor authentication at all.

I have no problem keeping my OTP tokens in my iPhone and in manually transcribing them. But that's me. So, better to use ANY OTP than none, even if it's being automatically filled-in by the browser. And if given a choice, better to **not** have the browser filling it in, even though the actual danger is realistically very small.

george palfi / @PalfiGeorge

Steve. I am a devoted listener and long term SpinRite owner. Though I wish it worked on Macs. I gave up Windows completely, years ago.

The good news is, I made some changes a few months ago to allow SpinRite v6.1 to run on macs where it can. "Where it can" means on Intel macs where it's possible to boot from a USB or CD. The previous trouble had been with the keyboard, since SpinRite was accessing the keyboard hardware rather than using the BIOS. I changed that so that SpinRite could work with some less PC-compatible Dell machines and we got mac compatibility in the bargain. A number of testers have confirmed their ability to now run SpinRite on their macs.

SpinRite

I'm finally able to announce that after more than 3 years of work, I'm completely satisfied that SpinRite v6.1 is as good as it can be, and that it is finally ready for release. There's nothing left I'm aware of that can be done to further improve SpinRite's functions. I could keep fussing with it forever, adding this or that convenience feature around the edges, but it's already received a large collection of new convenience features and it is, by far, the best SpinRite that's ever been created. It's been proven to work in every environment it's been placed in by more than 818 testers who've registered with our GitLab instance and who have obtained it through my release announcements in GRC's web forums. It's finally ready.

Officially, its code still calls itself release candidate 6, and it makes sense to let it rest for a bit before it's moved to final release 1 since I would prefer not to have to be tweaking the code after it's been released. And there's really no hurry. While the paint is still wet and drying I'll be working on SpinRite's documentation which will all be browsable and explorable online. Since many people prefer to click on a video than read text, I'll be creating video walkthroughs as I did for ReadSpeed, so that someone can get a feel for what SpinRite looks like when it's running.

Next, once the documentation is finished, I'll bring GRC's long awaited eMail facility online to get our promised incoming eMailbag setup to receive incoming mail from this podcast's listeners. So many people have written that they had to login to Twitter just to get a note to me – I get it. That will finally be changing. And I'll create a weekly mailing list for this podcast so that those who would like to receive a weekly summary and link to the show notes will be able to get that, too. I'm sure I'll continue posting on Twitter, but I'm not yet sure whether I'll continue monitoring incoming Tweets and DM's there. We'll play that by ear. I would very much like to consolidate the channels I need to follow and eMail is the most universal medium we all share.

And once all of that is in place, I'll finally begin the process of notifying all 20 years worth of SpinRite's past purchasers. Since I imagine many of those 20 year old eMail addresses are no longer valid, I plan to send announcements starting from the most recent and heading toward the least recent.

Stamos on “Microsoft Security”

Recall that “Midnight Blizzard” is the dramatic renaming Microsoft gave to the Russian state sponsored group, originally known as NOBELIUM, which most recently managed to crawl inside Microsoft’s network to obtain access to data belonging to their upper-most top level executives.

As we covered last week, late Friday night before last, Microsoft slipped out the news that a lesser-protected system had succumbed to the Russians after being sufficiently sprayed with passwords. What Microsoft shared at the time left no one feeling satisfied. So last Thursday the 25th, Microsoft attempted to offer additional useful information. Most observers have still been left wanting. The reading between the lines that we did last week appears to have been correct. At the top of last Thursday’s lengthy update, they wrote:

As stated in the MSRC blog, given the reality of threat actors that are well resourced and funded by nation states, we are shifting the balance we need to strike between security and business risk – the traditional sort of calculus is simply no longer sufficient. For Microsoft, this incident has highlighted the urgent need to move even faster.

If the same team were to deploy the legacy tenant today, mandatory Microsoft policy and workflows would ensure MFA and our active protections are enabled to comply with current policies and guidance, resulting in better protection against these sorts of attacks.

Microsoft was able to identify these attacks in log data by reviewing Exchange Web Services (EWS) activity and using our audit logging features, combined with our extensive knowledge of Midnight Blizzard. In this blog, we provide more details on Midnight Blizzard, our preliminary and ongoing analysis of the techniques they used, and how you may use this information pragmatically to protect, detect, and respond to similar threats in your own environment.

Using the information gained from Microsoft’s investigation into Midnight Blizzard, Microsoft Threat Intelligence has identified that the same actor has been targeting other organizations and, as part of our usual notification processes, we have begun notifying these targeted organizations.

As I noted, many if not most observers of Microsoft’s handling of this incident have come away less than impressed. So I wanted to share the highlights of an important industry-shaping interview Alex Stamos conducted with CNBC last Friday.

To remind everyone, Alex is a computer scientist, having obtained his EECS degree from Berkeley. Today, he’s an adjunct professor and lecturer at Stanford University's Center for International Security and Cooperation. He first popped onto our map when he left Facebook after serving as their chief security officer and then, in 2021, teamed up with ex-CISA director Chris Krebs. Recall that Chris was fired from his position as director of CISA by President Trump after CISA put out a statement declaring that the 2020 US Presidential election had been the most secure election in American history. So Chris and Alex were both free, and they formed the Krebs Stamos Group. That group later became part of SentinelOne where Alex now has the title of “Chief Trust Officer”. He often serves as an expert witness in court and provides expert testimony to Congress.

Okay. So Alex's credentials are well established within the industry and government. The following, which I wanted to share, is what he posted last Friday following Microsoft's updated breach disclosure. The title Alex gave his Linked-In posting was "Microsoft's Dangerous Addiction To Security Revenue". Under that headline, he wrote:

On Monday, CNBC gave me a chance to discuss Microsoft's Friday-night news dump of a new breach by Russian intelligence services, in which I called for more details from Microsoft so that other organizations could defend themselves.

Yesterday, we gained a bit more transparency in the form of a blog post from "Microsoft Security", the commercial security division of Microsoft. [Alex brackets the phrase "Microsoft Security" in air quotes... I'm not sure how he meant that, but it doesn't seem flattering.]

Some reactions:

1) Microsoft buries the lead with this paragraph:

"Using the information gained from Microsoft's investigation into Midnight Blizzard, Microsoft Threat Intelligence has identified that the same actor has been targeting other organizations and, as part of our usual notification processes, we have begun notifying these targeted organizations."

*Translation: Since the techniques outlined in the blog only work on Microsoft-hosted cloud identity and email services, this means that other companies were compromised using the same flaws in Entra (better known as Azure Active Directory) and Microsoft 365. Microsoft's language here plays this up as a big favor they are doing the ecosystem by sharing their "extensive knowledge of Midnight Blizzard" when, in fact, what they are announcing is that this breach has affected multiple tenants of their cloud products. (And in a subsequent update to his original posting Alex notes that Joseph Menn of the Washington Post has several sources indicating that **at least ten companies** were breached and will be disclosing soon.)*

2) Microsoft continues to downplay the attack by abusing the term "legacy".

One of the big open questions from last week was how an attack against a "legacy non-production test tenant" could lead to access to the emails of key Microsoft executives. We get a bit more detail in this paragraph:

"Midnight Blizzard leveraged their initial access to identify and compromise a legacy test OAuth application that had elevated access to the Microsoft corporate environment. The actor created additional malicious OAuth applications. They created a new user account to grant consent in the Microsoft corporate environment to the actor-controlled malicious OAuth applications. The threat actor then used the legacy test OAuth application to grant them the Office 365 Exchange Online full_access_as_app role, which allows access to mailboxes."

I've seen this fundamental problem in multiple investigations, including the one that Microsoft worked so hard to label as the "Solarwinds Incident": AzureAD is overly complex, and lacks a UX that allows for administrators to easily understand the web of security relationships and dependencies that attackers are becoming accustomed to exploiting.

In many organizations, AzureAD is deployed in hybrid mode, which combines the vulnerability of cloud (external password sprays) and on-premise (NTLM, mimikatz) identity technologies in a combination that smart attackers utilize to bounce between domains, escalate privilege and establish persistence.

***Calling this a "legacy" tenant is a dodge;** this system was clearly configured to allow for production access as of a couple of weeks ago, and Microsoft has an obligation to secure their legacy products and tenants just as well as ones provisioned today. It's not clear what they mean by "legacy", but whatever Microsoft's definition, it is likely to be representative of how thousands of their customers are utilizing their products [today].*

Microsoft does, however, offer all of us some solutions...

3) Microsoft is using their own security flaws as an opportunity to upsell.

These sentences in the blog post deserve a nomination to the Cybersecurity Chutzpah Hall of Fame, as Microsoft recommends that potential victims of this attack against their cloud-hosted infrastructure:

- "Detect, investigate, and remediate identity-based attacks using solutions like Microsoft Entra ID Protection.*
- Investigate compromised accounts using Microsoft Purview Audit (Premium).*
- Enforce on-premises Microsoft Entra Password Protection for Microsoft Active Directory Domain Services."*

*[In other words] Microsoft is using this announcement as an opportunity to upsell customers on **their** security products, which are apparently necessary to run their identity and collaboration products safely!*

This is morally indefensible, just as it would be for car companies to charge for seat belts or airplane manufacturers to charge for properly tightened [door] bolts.

*It has become clear over the past few years that Microsoft's addiction to **security product revenue** has seriously warped their product design decisions, where they hold back **completely necessary functionality** for the most expensive license packs or as add-on purchases.*

I'm going to interrupt Alex for a moment to just note that while all of this is highfalutin enterprise stuff, I've long made the same point about Microsoft leveraging the insecurity of their "out of support" operating systems. They blithely offer additional years of **extended** security support for their otherwise "out of support" operating systems to their enterprise customers,

while at the same time starving the end users of those same operating systems of that vital security in a bald effort to force users to move to newer operating systems which they neither need nor want. If the security updates are available anyway, deliberately withholding, as ransom, the patches to your defective operating system – because you can – is morally indefensible and reprehensible.

Referring to Microsoft's two recent posts, Alex continues:

While these two arrogant and circumspect posts do, at least, admit "the urgent need to move even faster" in securing their products, [Alex writes,] I would argue that Microsoft has a much deeper cultural problem to solve as the world's most important IT company.

They need to [discard] ~~throw away~~ this poisonous idea of security as a separate profit center and rededicate themselves to shipping products that are secure-by-default while providing all security features to all customers. I understand the need to charge for log storage or human services, but we should no longer accept the idea that Microsoft's basic enterprise offerings (including those paid for by the US taxpayer) should lack the basic features necessary to protect against likely attacks.

*My current employer competes against some of these products from Microsoft, but if ~~they~~ [Microsoft] did a better job by default ~~then~~ that would ~~actually~~ reduce the need for SentinelOne and other security vendors to provide **basic safety protections**.*

*For all the language about the sophistication of the **SVR** hackers behind this attack, there is nothing here that is outside the norm for ransomware groups attacking Microsoft technologies, **and Microsoft customers of all sizes should be concerned that these techniques will be deployed against them if they do not pay extra for the secure version of Microsoft's cloud products.***

Twenty one years after the Trustworthy Computing memo, it's once again time for some soul searching in Redmond.

I love the system of free enterprise we enjoy in these United States. The profit motive provides strong impetus to innovate and provide value. But the lure of increased profit carries a danger when an executive faces a decision about whether to include a desirable and important feature in the base product, or to charge extra for it.

A crucial feature that's necessary for this system of free enterprise to deliver its maximum value to the public at large, rather than to simply further line the pockets of those executives and their shareholders ... **is competition**. While it's an enviable position to be in, Microsoft is only able to get away with these usury practices because they have no real competition in the markets they dominate. This has been a problem for them in the past, and it may be again in the future.

