Security Now! #958 - 01-23-24

A Week of News and Listener Views

This week on Security Now!

What mistake did Microsoft make that allowed Russians to access their top executive's eMail? What does the breach of US Health & Human Services teach us? What does Firefox's complaint about Apple, Google & Microsoft mean? Why has the Brave browser just reduced the strength of its anti-fingerprinting measures? Last year CISA started proactively scanning. How'd that go? What new feature of smartphones has become a competitive advantage? And just how Incognito is that mode? Then we'll wrap up the week by looking at some of the best feedback from our listeners, including what's the future of fraudulent media creation?, how should a high school listener of our gets started with computing?, why did a popular Android app suddenly become sketchy?, does Google's Privacy Sandbox allow websites to customize their presentations to their visitors?, how might last week's LG smart washing machine have become infected?, does the Protected Audience API also protect its audience from malvertising?, and why do big ISPs just pull the plug on DDoSed sites rather than attempt to protect them?

You have to wonder how much use that peep hole gets...

Security News

Microsoft's Top Execs' Emails Breached in Sophisticated Russia-Linked APT Attack

Last Friday the 19th, the rest of the world learned that Microsoft's top executives had fallen victim to a Russian state-sponsored password attack which breached their eMail accounts. Here's what Microsoft shared in their Friday blog posting, titled: "Microsoft Actions Following Attack by Nation State Actor Midnight Blizzard". They wrote:

> *The Microsoft security team detected a nation-state attack on our corporate systems on January 12, 2024, and immediately activated our response process to investigate, disrupt malicious activity, mitigate the attack, and deny the threat actor further access. Microsoft has identified the threat actor as **Midnight Blizzard**, the Russian state-sponsored actor also known as **Nobelium**. As part of our ongoing commitment to responsible transparency as recently affirmed in our Secure Future Initiative (SFI), we are sharing this update.*
>
> *Beginning in late November 2023, the threat actor used a password spray attack to compromise a legacy non-production test tenant account and gain a foothold, and then used the account's permissions to access a very small percentage of Microsoft corporate email accounts, including members of our senior leadership team and employees in our cybersecurity, legal, and other functions, and exfiltrated some emails and attached documents. The investigation indicates they were initially targeting email accounts for information related to Midnight Blizzard itself. We are in the process of notifying employees whose email was accessed.*
>
> *The attack was not the result of a vulnerability in Microsoft products or services. To date, there is no evidence that the threat actor had any access to customer environments, production systems, source code, or AI systems. We will notify customers if any action is required.*
>
> *This attack does highlight the continued risk posed to all organizations from well-resourced nation-state threat actors like Midnight Blizzard.*
>
> *As we said late last year when we announced Secure Future Initiative (SFI), given the reality of threat actors that are 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. We will act immediately to apply our current security standards to Microsoft-owned legacy systems and internal business processes, even when these changes might cause disruption to existing business processes.*
>
> *This will likely cause some level of disruption while we adapt to this new reality, but this is a necessary step, and only the first of several we will be taking to embrace this philosophy.*
>
> *We are continuing our investigation and will take additional actions based on the outcomes of this investigation and will continue working with law enforcement and appropriate regulators. We are deeply committed to sharing more information and our learnings, so that the community can benefit from both our experience and observations about the threat actor. We will provide additional details as appropriate.*

So, this reads as a bit of a wake-up call for Microsoft. And it's interesting because there's a lesson here for every large enterprise. It sounds as though some older systems still have older levels of security and that they've been allowed to continue purring along, undisturbed, since

they hadn't been bothering anyone – though they were still online and accepting incoming connections. Presumably, all newer systems are being deployed with stronger password quality minimums, multiple-factor authentication, brute force detection and prevention and all of the additional layers of security that have become modern standard practice.

The lesson here, which Microsoft just learned the hard way, and which I wanted to bring up to our IT-managing listeners, is that the law of the lowest hanging fruit applies to legacy machines that aren't bothering anyone. They aren't bothering anyone until they become the source of ingress into an enterprise's interior.

So, just a note of caution to remember that bad guys won't attack the most secure entry points to an organization, they will attack the weakest. And that might be some machine that still has a password and policies that haven't been considered safe since the turn of the century. And this is really a thing. We all know the lesson LastPass learned by failing to proactively enforce PBKDF iteration counts which were current with password cracking capabilities. Security really is a moving target and older systems won't improve their older security without being revisited.

We would not expect Microsoft not to put the best face on this possible. So we can assume they did. But the news of this breach received some harsh criticism from other quarters. Here's what one respected security reporting group had to say. Starting by speaking about some eMail content, they wrote:

> *Microsoft's disclosure language does not specifically state that this was the **only** stolen information, but it is worth pointing out that Microsoft is currently hosting the Ukrainian government's entire network on its Azure cloud infrastructure.*

That's interesting. I hadn't encountered that little tidbit before. They continue, and write:

> *The breach has drawn quite an avalanche of criticism and ridicule for Microsoft for various and well-deserved reasons.*
>
> *First, Microsoft disclosed the breach late on a Friday night — a well-known scummy tactic to hide the incident from extended media coverage.*
>
> *Second, the breach took place weeks after Microsoft announced, with bells and whistles, its new Secure Future Initiative, a new plan to re-focus the company's engineering efforts to improve the security of its own products. The new initiative was meant to mimic a similar pledge made by Bill Gates in 2002—named Trustworthy Computing—that led to significant changes to Microsoft's security posture and the creation of what we now know as Patch Tuesday.*
>
> *Third, the new breach took place four months after Microsoft disclosed another state-sponsored hack, this one by China's Storm-0558, which also had access to its internal network.*
>
> *Fourth, after promoting multi-factor authentication as the next evolution of online account security, the fact that one of its test accounts got popped via a password spray suggests Microsoft was not high on its own supply.*

The hack is quite bad, but not for most of you reading this. It may not have a material impact on day-to-day Microsoft users, but it has quite the reputational damage on Microsoft's position in the cybersecurity market.

Having Russian intelligence services breach your cybersecurity team's email accounts to steal data about themselves four months after the Chinese breached your production systems to steal US government emails is not what this industry calls trustworthy.

Most of that criticism is covered by my earlier observation that Microsoft acknowledges that they're going to need to stir things up a bit and ruffle some feathers in the interest of bring the security of their older systems up to today's standards. Everyone listening to this would be well advised to do the same.

HHS Breached:
A still unknown threat actor stole $7.5 million from the US Department of Health and Human Services in a security breach that took place between March and mid-November of last year. It struck me as interesting that the range is that broad. There's quite a lot of time between March and November. The unknown attackers are believed to have gained access to an HHS system that processes civilian grant payments using spear-phishing. They then proceeded to hijack payments for five grant recipients before being detected. The investigation to identify the perpetrators is still underway.

So our takeaway here is that the human factor remains Security's number one Achilles heel. Having strong outbound security – such as that provided by the AdamNetworks guys who so impressed me last year – and also training training training are probably all that can be done. It's just so easy for a harried worker who has too much going to click a link that they shouldn't.

Firefox vs "The Competition"
Also last Friday, Mozilla posted a complaint to the industry under the heading "Competition". The title of this posting was: "Platform Tilt: Documenting the Uneven Playing Field for an Independent Browser Like Firefox". Here's what Mozilla wrote:

Browsers are the principal gateway connecting people to the open Internet, acting as their agent and shaping their experience. The central role of browsers has long motivated us to build and improve Firefox in order to offer people an independent choice. However, [the centrality of the browser] ~~*this centrality also*~~ *creates a strong incentive for dominant players to control the browser that people use. The right way to win users is to build a better product, but shortcuts can be irresistible — and there's a long history of companies leveraging their control of devices and operating systems to tilt the playing field in favor of their own browser.*

This tilt manifests in a variety of ways. For example: making it harder for a user to download and use a different browser, ignoring or resetting a user's default browser preference, restricting capabilities to the first-party browser, or requiring the use of the first-party browser engine for third-party browsers.
For years, Mozilla has engaged in dialog with platform vendors in an effort to address these issues. With renewed public attention and an evolving regulatory environment, we think it's

time to publish these concerns using the same transparent process and tools we use to develop positions on emerging technical standards. So today we're publishing a new issue tracker where we intend to document the ways in which platforms put Firefox at a disadvantage. We wish to engage with the vendors of those platforms to resolve them.

This tracker captures the issues we experience developing Firefox, but we believe in an even playing field for everyone, not just us. We encourage other browser vendors to publish their concerns in a similar fashion, and welcome the engagement and contributions of other non-browser groups interested in these issues. We're particularly appreciative of the efforts of Open Web Advocacy in articulating the case for a level playing field and for documenting self-preferencing.

People deserve choice, and choice requires the existence of viable alternatives. Alternatives and competition are good for everyone, but they can only flourish if the playing field is fair. It's not today, but it's also not hard to fix if the platform vendors wish to do so.

We call on Apple, Google, and Microsoft to engage with us in this new forum to speedily resolve these concerns.

Many of us prefer to use Firefox as our browser of choice. I have Chrome and Edge but URL clicks are always sent to Firefox. And I have Firefox installed on my various Apple iOS devices. So I dug a bit deeper into this new issue tracking system and it was quickly apparent that Apple had the most strikes against it. At this moment, Mozilla is complaining about:

- *App Store forbids third-party browser engines*
- *Support for third-party multi-process applications on iOS*
- *JIT Support on iOS*
- *Accessibility APIs on iOS*
- *Messages integration on iOS*
- *Importing browser data on iOS*
- *Setting and checking default browser on iOS*
- *Origin-Based Associated Domains dependent features for 3rd-party browser engines*
- *Browser extension support on iOS*
- *Beta testing on iOS*

We know how heavy-handed Apple is. I'm an avid user of Amazon's Kindle readers and also of Amazon's Kindle app on iOS where I use it on iPads and my iPhone. And it is a constant and ridiculous annoyance that Apple refuses to allow Amazon users to purchase books through the Amazon app. It's necessary to use a web browser. Why? Because Apple has iBooks and cannot stand the competition. It's just so petty and it should be beneath Apple... but apparently it isn't.

I'm sure that Apple's reticence to allow Chrome and Firefox and any and all other non-Safari browsers to enjoy the same privileges they have on other platforms is largely about security. As we know, browsers have become the #1 way for evildoers to crawl inside our computers. So I don't blame Apple for that. But given my experience with Amazon Kindle books, I also have no doubt that some of this is just pettiness which, as I said, should be beneath Apple. For what it's worth, though, I'm sure Apple is not singling out Firefox for prejudicial mistreatment. They treat anything that's not Safari as suspect.

Mozilla is also unhappy with their experience over on Google's Android platform. There, they voice three complaints:

- *Importing browser data on Android*
- *Some Android features launch Chrome instead of the user's default browser*
- *Lower quality search result pages in third-party browser engines on Android*

I was curious to look into these three a bit further – especially the last one which we'll get to in a second. What I found was interesting. In detailing their complaint about "*Importing browser data on Android*" Mozilla explained:

Browsing information like history, bookmarked sites, and cookies isn't accessible to third-party browsers on Android. This data is kept within a web browser application's data directory, which isn't directly accessible to third-party browsers, and there's no API or ContentProvider to enable it to be imported. While this is sensitive data, similar import functionality is possible on all major desktop platforms, and Android is able to mediate access to other sensitive data with user consent. Not being able to import data creates significant friction to change from Chrome - a user should be allowed to bring their data with them to another browser.

This seems like a legitimate complaint and a slippery way for Google to give Chrome an anti-competitive edge over any other browser its user might wish to switch to. And the second issue raised: "*Some Android features launch Chrome instead of the user's default browser*" seems even more insidious. Mozilla explains:

Features like Google Search, or Discover, in the pre-installed Google application ignore the user's default browser choice: links to websites outside of the application are always opened in Chrome, regardless of the default browser. This is a widely used application, with additional entry points from built-in features such as the search bar on the home screen and app launcher. Each time it opens a link in Chrome, a user is driven away from their default browser. All built-in applications and affordances that open external links should open them in the user's default browser.

Right. That would really annoy me, and this issue will be quite familiar to anyone who has heard Paul Turrott ranting about Microsoft and Edge doing the same thing. As Mozilla says, every time Chrome is launched when the user has installed Firefox and asked Android to use it, drives the user toward Chrome despite their clearly expressed browser preference.

It was the third item in Mozilla's "Platform Tilt" list of grievances that most caught my eye. They wrote: "*Lower quality search result pages in third-party browser engines on Android*". That seemed like a real antitrust showstopper. Here's what Mozilla explained:

*The web search experience is tightly integrated with a number of built-in features in Android and the experience provided to Firefox is **inferior** compared to the version provided for Chrome.*

As seen in the screenshots, identical search terms show less information and receive a lower quality design in Firefox on Android.

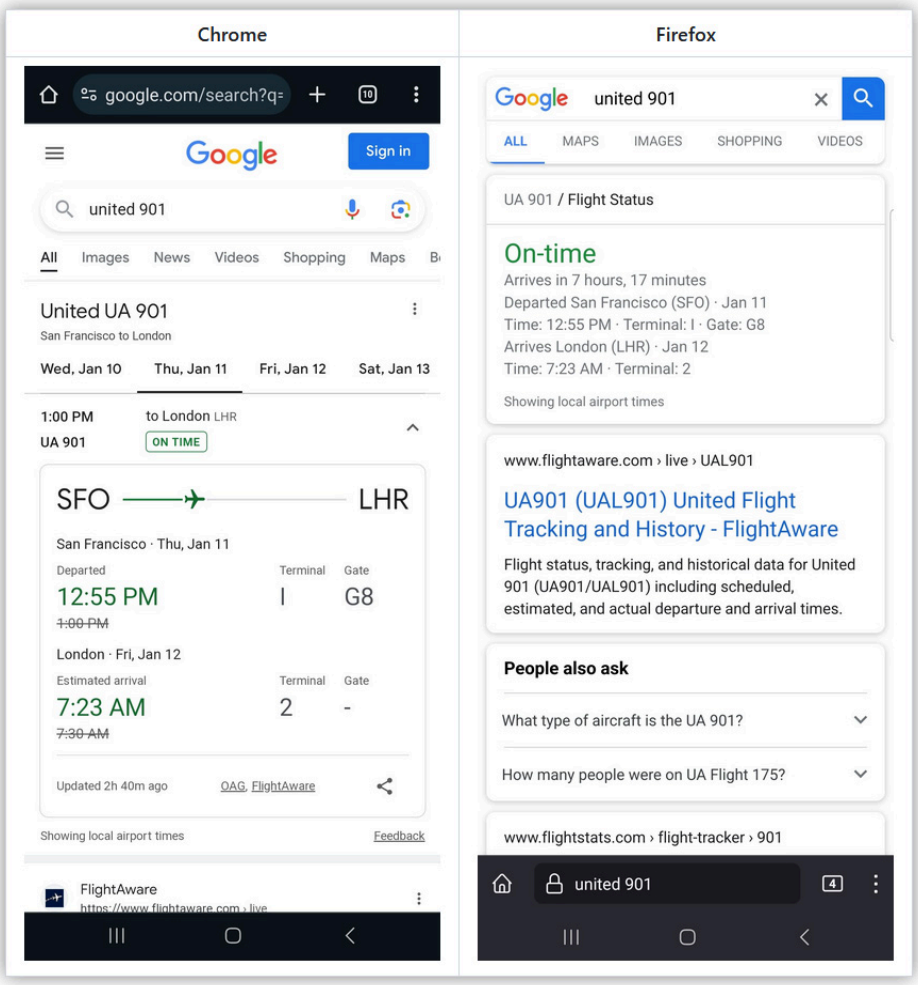

While strictly speaking this is an issue with the Google Search website, given the prominence and integration of search on Android this is a meaningful user experience gap that creates an incentive for users to not choose a third-party browser - especially those implemented with third-party browser engines, like Firefox. There are no technical limitations which would prevent this page from operating in Firefox: an equal experience should be offered.

It turns out that Google's search results are biased against non-Chrome browsers and that if the User-Agent string is changed then Google will provide the same improved experience to Firefox users as Chrome users. User agent dependency is nothing new. And once upon a time Chrome's page results rendering may have necessitated producing different results to differing browsers. But these days this sort of deliberate bias is showing Google's own extreme pettiness.

And speaking of Microsoft and Windows, Microsoft's own incestuous ties to its own web browser actually ***have*** been the subject of antitrust lawsuits. And big ones. Mozilla lists three complaints about Microsoft's and Windows treatment of Firefox:

- *Setting default browser on Windows*
- *Default browser is set to Edge by several Windows flows*
- *Some Windows features launch Edge instead of the user's default browser*

This is starting to sound like a refrain. Under "*Setting default browser on Windows*" Mozilla writes:

Allowing a third-party browser to programmatically set itself as the default is an important platform feature. Without this, even after the user has installed the browser of their choice they must navigate operating system settings and make the choice there as well. This adds friction and creates inertia to continue using Edge, despite the user's obvious preference.

A well-established design pattern is to allow the third-party browser to invoke a system prompt which permits the user to easily confirm or reject the request to set the current browser as the default. This is an intuitive user experience that mirrors similar permissions models used in operating systems, browsers, and web applications. Android and macOS offer such a capability.

Unfortunately, Windows does not support anything like this for third-party browsers: browsers are forced to "deep link" into the Windows settings UI. On Windows 10 this requires several clicks and a double confirmation in the settings UI. On Windows 11 there is a "Set default" button. Neither is sufficient. Windows should instead provide a method for third-party browsers to programmatically request they be set as the default.

To that I'll just say, Yup. This is the traditional way that we've all historically experienced the addition of a 3rd-party browser being installed. The browser notices that it's not currently the system's default URL handler and asks its user whether they would like it to switch them over to using this browser instead. The user says yes please or no thanks and it's done. But no longer. Microsoft, exhibiting the same pettiness we see from Apple and Google, clearly wishes to hold onto the use of Edge every way possible. As I sit in front of Windows 10, I'm periodically reminded of just how much my life could be improved if only I would allow their Edge browser to service my needs. No thank you.

And speak of the devil, here's Mozilla's second complaint:

In general, the Windows 10 and 11 operating systems have persistent messaging that Microsoft Edge is the "recommended" browser for Windows, and offer affordances to change the default browser to Edge. In some cases the wording is misleading, asking a user to adopt "recommended browser settings", which does not obviously suggest a default browser change. This messaging is a moving target, with examples added and removed from Windows over time, often on UI surfaces that appear automatically on update or otherwise, making it difficult to enumerate specific examples.

In all cases these Windows components are able to change the user's default browser directly, and are not forced to use the ms-settings: protocol deep linking that browsers are required to use. Windows should consume the same affordances and APIs that are available to third-party browsers for setting-to-default.

Yep. Just another of the many reasons I'm perched in front of my trusty (and crusty) old Windows 7 system at this very moment. I'm subjected to none of that extraneous crap. And finally...

> *There are at least three prominent Windows features that open URLs in Microsoft Edge and not in the current default browser. The user's default browser choice should be respected when web pages are opened by built-in operating system features.*
>
> *The first is Windows Search, also known as Start Menu Search, and formerly known as Cortana. The UI for this feature is represented by a taskbar search box or search button (depending on user settings), and a search suggestions / results UI that appears when activated and updates as the user types. The suggestions and results UI also appears if the user starts typing when the start menu is open, and by the WIN+S hotkey. All links from this UI, whether they initiate web searches or link directly to articles or results, open in Microsoft Edge regardless of the user's default browser.*
>
> *The second is the new Windows Copilot, currently only available on Windows 11, which appears as a docked window on the right side of the screen. If Copilot produces links in its responses, or offers other links within its rendering area, these links open in Microsoft Edge regardless of the user's default browser.*
>
> *The third are Windows "widgets" which are called "news and interests" on Windows 10, a UI surface area which can be activated by a taskbar button. These show information like news, weather, stocks, and sports scores. On Windows 11 new widgets can be added from 3rd parties. Regardless, all links to a web page from widgets will open in Microsoft Edge regardless of the user's default browser.*

So, in summary, what we have from Mozilla is highlighting and detailing pervasive pettiness on the parts of Apple, Google and Microsoft. Some of Apple's is likely warranted for security reasons, but most of it is just mean spirited. The problem is, each of these major platforms publishes their own web browser. It's what they want their own users to use and they're willing to continually push and prod those users – without showing them any respect – against their express wishes.

It'll be a sad day for the world if we someday lose Firefox.

Brave reduces its anti-fingerprinting protections

Meanwhile, last Thursday, the Braver browser – which is super popular among those who are truly privacy and anti-tracking concerned – notified its users that it would be reducing the strength of its anti-fingerprinting protections. Under the heading "Brave browser simplifies its fingerprinting protections", the Brave team wrote:

> *With desktop and Android version 1.64 in a couple of months (and in today's Nightly release for testing), Brave will sunset Strict fingerprinting protection mode. This does not affect Brave's industry-leading fingerprinting protection capabilities for users. Instead, it will allow us to focus on improving privacy protections in Standard mode and avoid Web compatibility issues.*

Okay now. They say "Brave will sunset Strict fingerprinting protection mode" and then immediately follows that with "This does not affect Brave's industry-leading fingerprinting protection capabilities for users." Hmmmmm. They continue:

> *Brave currently offers two levels of fingerprinting protections. which make it harder for tracking companies to identify you as you browse the Web: Standard and Strict mode. Over time, however, we have observed significant disadvantages of Strict mode:*

Ah.

> - *In order to block fingerprintable APIs, Strict mode frequently causes certain websites to function incorrectly or not at all. This website breakage means that Strict mode has limited utility for most Web users.*
> - *Fewer than half a percent of Brave users are using Strict fingerprinting protection mode, based on our privacy-preserving telemetry data.*
> - *This tiny cohort of users could be more vulnerable to being fingerprinted because they stand out as a result of using Strict mode. Although we have not seen issues around this, it is a valid concern given that users who select Strict fingerprinting protection might have done so because of an elevated concern about tracking.*
> - *Maintaining Strict mode and debugging why some websites are broken on Brave takes our engineers' time away from focusing on default privacy protections that can benefit all of our users.*
>
> *These observations have led us to the conclusion that sunsetting Strict mode in Brave will actually be beneficial to our users' privacy.*
>
> *Brave's Standard fingerprinting protection is already very extensive and the strongest of any major browser. Brave's innovative **farbling** of a number of major fingerprintable Web APIs makes it difficult for fingerprinters to get a reliable unique ID on your browser. Going forward, we will continue to strengthen and expand Brave's Standard fingerprinting protections so that all our users have ever-improving protection against fingerprinters, while maintaining the highest possible level of compatibility with websites.*

Okay, first of all, you did hear me use the term "***farbling***". I have no idea where they came up with that. But, okay... I tracked it down and it's Brave's term for introducing some random jitter noise into the values being returned by the Web APIs that are commonly used for fingerprinting.

Those APIs are: The Canvas API, WebGL, WebGL v2, the WebGL Extensions, the contents of the browser's User Agent header, Web Audio, the browser's Plugins, Hardware Concurrency, the enumeration of system devices – both their ordering and their labels and IDs – and the user's dark mode setting.

Since I was still curious, and knew that our listeners would be, too, I tracked down the difference between Brave's soon-to-be-discontinued "Strict" anti-fingerprinting mode and the mode that all Brave browser users will be left with. Here's how Brave describes the two modes:

Brave has two levels of fingerprinting protections. In the default, "standard" configuration, Brave adds subtle noise to APIs commonly used for fingerprinting. This small amount of noise is enough to make you look different to fingerprinting scripts, without breaking websites, and will provide good protections against web-scale online trackers.

Brave also includes a "strict" option. When set to "strict" mode, Brave only returns random values from APIs commonly used by fingerprinters. This provides a higher level of protection against highly determined attackers, who may attempt statistical and / or targeted attacks to identify users. This mode will also break websites who depend on these features to work correctly.

In other words, the milder "standard" mode uses – dare I say – only moderate **farbling** of API values which do not cause website issues because only some of the least significant bits are being **farbled**. But what strict mode does is to entirely discard the true API values and replace them with fully random values for these API calls that bear no resemblance to reality.

My reading of this is that the original designers of Brave's anti-fingerprinting technology probably got their **farble** turned up too high. They thought: "You know, if a little farbling is good, just think how great it would be if we just farbled the crap out of this!" But, apparently, after gaining more experience with this they learned that some websites became quite upset when they were over- farble. You never want to over-farble, but especially on a school night.

And I can see how the statistical analysis they refer to could theoretically be a problem, since over time the results from a low and safe level of farbling could be averaged out to obtain the true value around which the farbled values are clustered. But on balance I wouldn't worry about that too much. I think Brave is doing the best they can while not causing more trouble than the farbling is worth. So, elimination of Brave's "strict mode" – sounds like a good thing.

CISA's proactive policing results one year later

Coming up on one year ago, in the middle of March 2023, I noted and was quite glad to share that CISA, our already very proactive US Cybersecurity and Infrastructure Security Agency, was launching an even more proactive initiative. They called it the Ransomware Vulnerability Warning Pilot (RVWP) and described it this way. They wrote:

Organizations across all sectors and of all sizes are too frequently impacted by damaging ransomware incidents. Many of these incidents are perpetrated by ransomware threat actors using known vulnerabilities. By urgently fixing these vulnerabilities, organizations can significantly reduce their likelihood of experiencing a ransomware event. In addition, organizations should implement other security controls as described on stopransomware.gov.

However, most organizations may be unaware that a vulnerability used by ransomware threat actors is present on their network. Through the Ransomware Vulnerability Warning Pilot (RVWP), which started on January 30, 2023, CISA is undertaking a new effort to warn critical infrastructure entities that their systems have exposed vulnerabilities that may be exploited by ransomware threat actors

As part of RVWP, CISA leverages existing authorities and technology to proactively identify information systems that contain security vulnerabilities commonly associated with ransomware attacks. Once CISA identifies these affected systems, our regional cybersecurity personnel notify system owners of their security vulnerabilities, thus enabling timely mitigation before damaging intrusions occur.

CISA accomplishes this work by leveraging its existing services, data sources, technologies, and authorities, including CISA's Cyber Hygiene Vulnerability Scanning service and the Administrative Subpoena Authority granted to CISA under Section 2209 of the Homeland Security Act of 2002.

As our listeners know, I'm 100% behind the idea of having the good guys proactively scanning for vulnerabilities. We know that the bad guys are. So, to the good guys my only question would be "what took you so long." Anyway, CISA just published their 2023 Year In Review and it contained some gratifying results from the first year of this pilot program.

During this first year, CISA sent more than 1,200 notifications to US ***and*** international organizations notifying them of early-stage ransomware activity on their networks. CISA also sent 1,700 notifications to organizations that had systems vulnerable to common ransomware entry vectors. In other words, the US is finally proactively scanning the public Internet for vulnerabilities. Together, this totals an average of 8 such notifications sent every day of the year. And it's difficult to imagine that anyone would blow off a notification from this US agency saying that they've found either evidence of an existing network intrusion or an existing public-facing vulnerability. So... bravo CISA!

Longer Life For Samsung Updates

In news of a growing and necessary trend which we've been seeing recently, Samsung's just launched S24 series of Smartphones will be receiving 7 years of software and security updates. That's an increase from the company's previous smartphones which were receiving 5 years of updates. So Samsung joins Google to be the only vendors to offer 7 years of security updates for their Android devices. And the best news of all is that this suggests that the longevity of security support has finally become a recognized competitive advantage. That's nothing less than a big win for consumers.

"This is what we meant all along!"

Remember that we were recently talking about the lawsuit against Google over the consumer's misunderstanding of the protections provided by Chrome's Incognito mode? In response to that, Google is changing the text that appears in Chrome's Incognito Mode browser mode. The new text much more clearly informs its users that their activity **will** continue to be tracked even while they are in the somewhat less than Incognito mode. It now reads:

*"**Others** who use this device won't see your activity, so you can browse more privately. This won't change how data is collected by websites you visit and the services they use, including Google. Downloads, bookmarks, and reading list items will be saved. Learn more"*

Closing the Loop

I wanted to acknowledge that I received everyone's notes that the images contained within the previous two weeks of show notes were not visible for users of Apple's iOS and macOS. It was also a coincidence that I had captioned last week's picture of the week "Please provide an example of irony" – and that the picture was missing and black.

In any event, it appears that the trouble was that I always run the final show notes PDF, which I download from Google Docs where each week's document is authored, through Acrobat's PDF optimizer which very nicely reduces the document's size, which sometimes starts off being many megabytes, down to a few hundred kilobytes. I don't know what may have changed, since it was nothing at this end. But since my Acrobat is v9 with a copyright of 2008, it might be that Apple decided to stop supporting something.

blaž ocepek / @blaz_ocepek

Hello. I started my journey in cyber security 2 years ago. I've learned a lot and still have a lot to learn. Recent news of AI-generated videos got me scared because of the future world my 2 year old son might be growing up in. In your opinion, what would be the best solution for automatic verification of a video, image or audio? I'm thinking of some kind of encryption from the camera like a few episodes ago where photos are signed. But if I remember correctly this is flawed because anyone can buy such a camera, dig the key from the hardware, and sign fake images. I believe this needs to be addressed as soon as possible and not like in 10 years. AI really took off who knows what might be next.

Through the years, as we've observed everything that's happening around us, this podcast has arrived at a number of "rules of the road"; guiding principles that always seem to apply. One of those, that I've occasionally marched out, is quite unsatisfying, though that doesn't render it any less true. And it is: "Not all problems have good solutions." An example we were talking about last week being Internet DDoS attacks. Like it or not, the fundamental design of the Internet has made it inherently vulnerable to spoofed bandwidth flooding and other sorts of attacks.

And I believe that we have the same problem here. When I was growing up, I was fascinated by optical illusions. One that pops to mind is that two parallel lines can be drawn on paper and yep, they look perfectly parallel. But place a series of radial lines exploding outward from a central point and those still perfectly straight parallel lines look curved. No matter how you try to tell yourself that they are not curved, curved is what you see. The trouble is, we're built to believe our senses and our senses can be fooled.

The corollary rule to "Not all problems have good solutions" is that technology cannot solve all of our problems for us. In fact, it's probably a zero-sum, with technology inadvertently creating just as many problems as it solves. Unfortunately, I'm virtually certain that this listener's 2 year old son is simply going to grow up in a very different world than ***we*** have. It's going to be a world where the many things we were able to take for granted as being real depictions of events will simply never be the case in a world for someone born recently.

And of course there's been some of that for us. The phrase "Oh, that image was Photoshopped" has long been a common meme. But until now that's been the exception. The fact that everyone

perceives that we're about to witness a wholesale explosion in the volume of fictitious content masquerading as authentic suggests that, if nothing else, it's going to be a self fulfilling prophecy.

And I don't see this view as pessimistic; I think it's realistic. As they say, being forewarned is to be forearmed. For those of us who have been around for a while this seems like a big change for the worse. We're accustomed to trusting our senses and believing what we see. But it seems all but certain that this is a comfort future generations will simply not have. But then neither will they miss it, since things will have always been that way for them. We old codgers will eventually die off, grumbling "when I was a boy..."

Matthew Burrell / @Matttthhhhew

Quick guidance question please. What 3-4 computer languages should I learn? That as a kid graduates high school and looks to start a business. That's mostly ground up, open source, secure, from server side, (that can do almost it all, etc., backup, database, etc.), to web interface (basically website, login, manage clients, database, other, etc). Is there a good platform to start from like Synology or something? That those languages could be built on top of? Thank you for all your knowledge, guidance, and all that you do.

Well, Matthew, I presume that you're describing yourself here, so you're a young person who is interested in computing technology and wanting to create intellectual property with computers and eventually support yourself. What you need more than anything is knowledge and experience. I told a story many years ago that had a somewhat surprising moral. The story was about my misadventures surrounding my construction of a sonic beam weapon which, being a high schooler at the time, I had named "The portable dog killer." The moral of the story was that all sorts of interesting and unexpected things transpired – but only because I was actively doing things. I was not sitting on my butt playing video games. Okay, so we didn't have video games back then, but there were still plenty of similar ways that my peers managed to burn away the seemingly endless hours of their day not learning anything, not pushing themselves and rarely experiencing anything new. I didn't really have any choice since I loved electronics back then as I love computers and computing today. So if you truly love computers, be active not passive. Turn off the video game that someone else created and start figuring out how to create your own stuff. And more than anything, don't let having no idea what you're doing in the beginning stop you. That's not where you stop... that's where everyone starts.

So pick a language, any language, it really doesn't matter which one. Python is nice, general purpose, easy to get going with, lots of help available online, and it can probably take you anywhere you want to go. Figure out how to get it to print "Hello, world!" and you'll be off and going. Then choose another problem that's not much harder, and solve that one... and so on. And before you know it, you'll be programming.

DBloor / @DaBloor

Hi Steve, I'm a computer forensics instructor up in Canada and have been listening since episode 20ish and love the podcast! Thanks for agreeing to push past 999!

I just wanted to see if you think my hypothesis holds up for a $720 + flashlight app. Did this app potentially get hi-jacked?

I've had this free flashlight app on my Pixel phone for over a year as it allows me to control the brightness since the stock flashlight doesn't have this option. It's been great and simple! I think it used to show a small banner ad occasionally, but nothing intrusive.

Today I tried to use it and I got a pop-up video ad play for about 20 seconds before I could use the light. I thought "Okay maybe they need some money to keep development. How much can the paid version be? I really like this app so why not chip in a bit?"

$15 USD..... PER WEEK! You read that right... a subscription for a flashlight like 4x the price of Netflix! $60/month or $720/year for a FLASHLIGHT!

My hypotheses are:

1. *The app developer got compromised or hi-jacked and someone is trying to scam its users of hundreds or thousands of dollars. Or*
2. *The developer had this in mind all along hoping to get a handful of users subscribe thinking a flashlight app couldn't possible be more than $15, not realizing it's a subscription.*

Either way... wow. Is this what our world has come to? Subscription-based FLASHLIGHTS!?

Anyway keep up the amazing work! I recommend this podcast on all my courses. The app is called "Simple Flashlight", produced by "Simple Mobile Tool" and has 1+Million downloads!

The fact that this app has over one million downloads and that it played a 20-second ad video, and that the app is just a once-free flashlight app, strongly suggests that its original developer, who had acquired a large user base, accepted an offer to sell the app to another party.

We covered this happening many years ago, also in the Android app store. The developer would be conscientious and well meaning, perhaps tired of keeping an app upgraded and current for little or no return. Then someone would come along and offer to buy the app from them outright. The developer of the free app, seeing one last chance to make some money, would take the deal and turn over his developer keys to its new owner.

The new owner, a scam artist, would quickly burden the app with crap designed to make more money than the purchase price of the app, which was likely not much. The scammer would figure that all those million+ users would run it and generate revenue from the ad, and just as our listener suggested, there might be a few, if only a tiny fraction, who might not be paying attention and who would inadvertently subscribe at this inflated rate.

Technically, the app's new owner had done nothing wrong, but neither is this a particularly upstanding way to generate income.

Michael Garrison / @iammikejed

Hey Steve, I'm listening to episode 957 about the Protected Audience API, and I have a question I'm hoping you can help me figure out. I work with small businesses who have no

interest in putting ads on their site, but I'm wondering whether they can still make use of the new ad functionality.

Say a company like TWIT wants to be able to customize what shows they feature on their homepage for new visitors. With the proposed (and abandoned) FLOC proposal, the bitmap of interests stored by the browser were available to the site itself, so if they knew which bit represented interest in space, they could alternatively move the This Week in Space banner higher on the page, or move it below other banners.

With the Topics API, my understanding is that wouldn't have been possible, because the same requester (ad company) would have to have seen the same browser on multiple other sites, obviously leaving first-party site owners in the dark (probably by design).

Now with the new API, I can't find a solid answer yet on whether the site itself would be able to see what categories of interests visitors might have. I assume it won't be available to the site owners, but if you know one way or the other, I'd love confirmation :)

Thank you for your great work on the show! Can't wait to see what email system you come up with for communicating in the future!

My reading and understanding of Google's Privacy Sandbox system agrees with Michael's. One of the significant objections raised against the earlier FLoC – Federated Learning of Cohorts – system was that a first time visitor to a site would be disclosing information about themselves to that site without any previous interaction and many privacy advocates found that to be a big step in the wrong direction. The Privacy Sandbox is vastly more complex than FLoC and it employs that complexity to effectively blind all of the parties so that all information flows into the user's browser and not outward. And when ads are shown their fetch and display frames enforce a new level of inter-frame and page isolation. The objections to FLoC taught us that websites are specifically unable to learn anything about their visitors.

Guillermo García / @gmogarciag

I have a question about the washer machine bot. Is it safe to assume that this malware is configured to infect this specific washing machine? In other words, is someone writing code for this model? How can malware infect various IoT devices? Or is there a unique one for each make and model? Many thanks, and I'm looking forward to being part of 999 and beyond!

I'd suggest that the best answer is a little bit of both. In the generic case of scanning the Internet for potential victims, we know that there have been turnkey IP stacks sold into the embedded device market which were later found to contain critical remote code execution vulnerabilities – often in their fragmented-packet reassembly implementation – since that's an example of something that's very easy to specify and surprisingly difficult to implement securely. So it would be possible to scan the Internet for any IP presence that could be compromised by such an attack that might be common to a wide variety of different IoT devices.

And then we definitely have the frequent and common case – and the washing machine might be part of that, too – of a known vulnerability having been discovered in some Internet connected appliance, after which those devices are explicitly targeted.

And then we have the other frequent case of a patch being made to a widely popular device, and that patch being quickly reverse engineered to start a race to see how many devices can be compromised before each individual device's administrator has applied the patch to prevent just such remote takeover.

In the case of the LG Smart Washer, I was wondering how a remote attacker might have gotten into such a machine in the first place. Any such device would be behind a NAT router that would not admit unknown traffic. It might support UPnP, allowing it to open a port for incoming traffic. But why would a washing machine need to be publicly visible? Another means of compromise might have been entry through some other means, such as a border router vulnerability, for example if its owner had enabled remote web administration. Then, once inside the network, a scan would have likely found everything on the LAN and at that point a known vulnerability in that specific washing machine might have been exploited.

Roger Stenerson / @rogerms

Hi Steve, The Protected Audience API sounds interesting and promising. However, the number one reason I use uBlock Origin and ScriptSafe is to block malvertising. Will the Protected Audience API help in that area? To 999+ and beyond! Thanks for all you do! Best regards.

I'm pretty certain that individual ads, once selected by the web browser, will still be able to run their own scripts within their "fenced frame." So we should not expect any added protection from malvertising. Roger's question caused me to do a bit of digging into the related "Fenced Frames API." The fencing that's created applied the same sort of cookie and other asset stove piping that Firefox implemented quite a while ago. The idea being that rather than all cookies and other asset storage sharing a single large database that's indexed by domain, thus allowing, for example, an advertiser to access information stored under their domain from any website hosting one of their ads... now each first-party website has its own private database containing anything that any 3rd-party might set while at that site. And other sites have the same. So there's no longer any chance for advertisement scripts to share data across sites.

Defensive Computing - Michael Horowitz / @defensivecomput

Steve: re the hacked washing machine: if the router supports outbound firewall rules, the hacked device can be blocked from making any outgoing network connections. Or, depending on the router, perhaps blocked from contacting certain IP addresses or certain domains. Surely pfSense can do this?

Oh, of course pfSense will do this. And such firewall rules can be tied to the machine's fixed Ethernet Mac address in case its IP should ever change. But I'm not the one with the LG Smart Washer on the Net so I'm not the one with the problem. I was more referring to the typical LG Smart Washer owner who would typically have no idea what was going on within their own network. I'm certain that our Security Now! Podcast listeners could readily block this activity. But Michael added something else that was interesting. He added:

I live in an apartment building with a Laundry Room in the basement. Both the washing machines and the dryers report their status to the Internet: running, not running or x minutes left until finished. This was helpful in the pandemic to avoid personal contact.

Huh. That's interesting. Michael says they report their status online, so presumably they're reaching out and communicating with a remote status-gathering server that's been configured to present a webpage which residents of the apartment building who share this common pool of machines are able to login to, to view their current status. So I stand corrected in my previous assertion that connecting washers and dryers to the Internet has little value. I can certainly see how this would provide tremendous value to a large pool of users who are sharing these machines and don't want to travel back and forth needlessly to check on them.

SKYNET / @fairlane32

Hi Steve, Is there really no way for ISP's like Cogent to differentiate between good and bad traffic so when a DDoS occurs they can null route only the bad traffic? Can you explain why an ISP isn't able to do this?

Could they? Perhaps. The best way to describe it is that doing so is beneath them. They simply cannot be bothered. I've had about 30 years of experience with ISPs of all sizes, and I've seen that the smaller the ISP and the closer they are to their subscriber, the more individualized service it's possible to have. But the top tier Internet backbone carriers like Cogent don't need to be bothered with those details, so they aren't. If some traffic is causing their downstream equipment any trouble whatsoever they'll simply drop that traffic as far upstream as possible.

The other change we've seen since the first early attacks is in their blockability. The days of the simple ICMP, TCP/SYN packet or UDP reflection floods have waned a bit. Back when spoofing source IP addresses was important to hide the traffic source, they were used. They still exist, but they've largely been replaced by non-spoofable HTTPS query floods. And those require highly specialized services such as those we often talk about offered by Cloudflare and others, to block. So these attacks can no longer be selectively blocked by simple firewall rules.

SpinRite

Since we appear to be stuck for the moment on the cusp of SpinRite's imminent release, I thought I'd update everyone briefly. As I planned, after last week's podcast, I finished the work on identifying and patching the known-buggy AMI BIOS which handled USB-connected drives. And after some verification and testing, later in the week I posted the next incremental test release. The overall reaction within SpinRite's testing community was jubilation since SpinRite had now lifted what I had been feeling was my previously heavy-handed 137GB clamp on any and all USB access. But it wasn't long before reports of new SpinRite crashing began to be seen. People were running SpinRite on their larger drives plugged into a USB port and SpinRite's own attempt to execute an illegal opcode capture screen began popping up. It turned out that HP and Lenovo and other BIOSes on older machines were also being found to be buggy and they were altering main memory and causing application crashes.

SpinRite's ability to patch the flaw that we found in those AMI BIOSes to allow them to safely work past 137GB was a fluke which just happened to work. In general, altering something as significant as an access size limit would not be simple. And HP BIOSes are in ROM, so they're not even patchable. And even if they had been shadowed into RAM for increased performance and could be patched, there's no reason to imagine that a simple patch would be available as it was for AMI.

So that's where we are at this instant in time. I spent yesterday and today focused upon this podcast. And tomorrow I'll see what more has transpired and decide what to do. But I already no longer feel that my previous policy of limiting SpinRite's use of USB to 137GB was as heavy handed as I thought. It might be that the best course of action will be to allow AMI BIOSes to be patched and used without limit, since that appears to be safe. To identify those, such as older HP BIOSes that are known to be unsafe and prohibit their use beyond 137GB. And to then add a command line option to override SpinRite's otherwise default USB clamp for unknown situations. That would allow a user who wishes to give it a try to do so, but that it will be at their own risk and if SpinRite crashes due to bugs in their old machine's BIOS, it should not come as any surprise.

The ultimate solution will be to get this finished and on to the work on SpinRite 7 which will not have any of these problems.

And of course...

Stay tuned for much more next week!