Security Now! #901 - 12-13-22

Apple Encrypts the Cloud

This week on Security Now!

This week we answer the following questions and more: What browser just added native support for passkeys and where are they stored? What service have I recommended that suffered a major multi-day service outage? How can you recognize a totally fake cryptocurrency trading site? Which messaging platform has become cybercrime's favorite, and how would you go about monetizing desirable usernames? What's the latest in TikTok legislative insanity, and is it insane? Which two major companies have been hit with class action lawsuits following security breaches? Was Medibank's leaked data truly useless? And Apple has finally given us the keys to our encrypted data in the cloud, holding none for themselves... or have they?

Security News

Chrome does Passkeys

ArsTechnica's Friday headline read: *"Passkey support rolls out to Chrome stable. With a huge list of caveats, initial Google passkey support is here"*. Ars wrote: *"Google's latest blog says: "With the latest version of Chrome, we're enabling passkeys on Windows 11, macOS, and Android." The Google Password Manager on Android is ready to sync all your passkeys to the cloud, and if you can meet all the hardware requirements and find a supporting service, you can now sign-in to something with a passkey."*

Ars takes some time explaining the stuff we know about passkeys' technology, getting it mostly right in a watered-down sort of way, then they talk about compatibility, writing:

> *Today passkeys essentially require a portable device, even if you are logging into a stationary PC. The expectation is that you'll use a smartphone for this, but you can also use a Macbook or iPad. The first time you set up an account on a new device, you'll need to verify that your authenticating device—your phone—is in close proximity to whatever you're signing in to. This proximity check happens over Bluetooth. All the passkey people are really aggressive about pointing out that sensitive data isn't transferred over Bluetooth—it's just used for a proximity check—but you'll still need to deal with Bluetooth connectivity issues to get started.*
>
> *When you're signing in to an existing account on a new device, you'll also need to pick which device you want to authenticate with (probably also your phone)—if both of these devices are in the same big-tech ecosystem, you'll hopefully see a nice device menu, but if not, you'll have to use a QR code.*
>
> *Second big issue: Did everybody catch that OS listing at the top? Google supports Windows **11** with passkeys—**not** Windows 10—which is going to make this a tough sell. Statcounter has Windows 11 at 16 percent of the total Windows install base, with Windows 10 at 70 percent. So if you happen to make a passkey account, you could only log in on newer Windows computers.*
>
> *Passkeys get stored in each platform's built-in keystore, so that's Keychain on iOS and macOS, the Google Password Manager (or a third-party app) on Android, and "Windows Hello" on Windows 11. Some of these platforms have key syncing across devices, and some do not. So signing in to one Apple device should sync your passkeys' access to other Apple devices via iCloud, and the same goes for Android via a Google account, but not Windows or Linux or Chrome OS. Syncing, by the way, is your escape hatch if you lose your phone. Everything is still backed up to your Google or Apple account.*
>
> *Google's documentation mostly doesn't mention Chrome OS at all, but Google says, "We are working on enabling passkeys on [Chrome for] iOS and Chrome OS." There's also no support for Android apps yet, but Google is also working on it.*

Then Ars wraps up this news of Chrome's emerging support for passkeys by writing:

Now that this is actually up and running on Chrome 108 and a supported OS, you should be able to see the passkey screen under the "autofill" section of the Chrome settings (or try pasting chrome://settings/passkeys into the address bar).

Next up we'll need websites and services to actually support using a passkey instead of a password to sign in. Google Account support would be a good first step—right now you can use a passkey for two-factor authentication with Google, but you can't replace your password yet.

Everyone's go-to example of passkeys is the passkeys.io demo site, which we have a walkthrough of here.

I have Chrome 108 so I put "chrome://settings/passkeys" into Chrome's address bar and was greeted with:

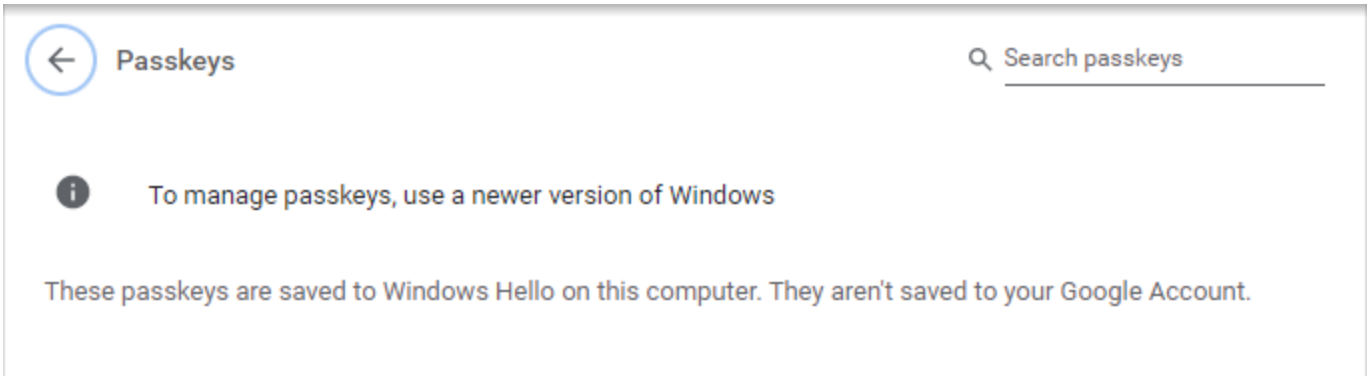

Oh, yeah… like 70% of the world, I'm running Windows 10 on this machine with no plans to change. The notice does say something interesting and encouraging, which is that "These passkeys (if any were listed) are saved to Windows Hello on this computer. They aren't saved to your Google Account. So that's cool. Unlike the browser's saved passwords, the browser is apparently not holding those keys. This potentially means that once Microsoft works out how to safely sync passkeys among Windows machines, assuming that they do, having browsers dynamically obtaining passkeys from the underlying OS might give us browser-agnostic use of passkeys, which would be terrific. And you'd have to imagine that Edge, also built on the Chromium core, would be working toward this quickly.

Overall, though, with these sorts of arbitrary-seeming limitations – lack of synchronization among competing platforms and devices with the attempts to create walled gardens – it's feeling as though passkeys, like passwords, may also become the domain of our existing password managers. We know that any password manager with a pulse has got to be racing as fast as possible toward the goal of supporting passkeys with full cross-client synchronization.

SYNC.COM suffered its first outage

While we're talking about synchronization I thought that I should mention that my favorite cloud synchronization platform, sync.com, had something happen to it. The problems first surfaced last Wednesday, apparently after a scheduled maintenance went wrong... and things didn't appear to be fully restored until just yesterday. I utterly depend upon sync.com and the outage had me revisiting the wisdom of that dependency. I'm not the typical user since I ran a pair of

NAS boxes at each of my two locations, and as we know, there are many alternatives. Reading between the lines of sync's frequent online status updates through the slowly passing days one got the sense that this was causing a major problem for their customers. So, as always, I'm hoping that lessons have been learned.

One thing I learned during this, as I was initially working to troubleshoot what was going on, thinking that perhaps something had become tangled up at my end, was that Sync's OS clients, at least for me on Windows, had **not** been auto-updating. Maybe they don't auto-update. But I've been using sync's service for so many years that my desktop clients had grown quite old.

Since I've talked up sync.com quite a lot here, and until last week my experience with them has been flawless, I wanted to urge everyone listening, who is using sync, to please go to their site to update your clients. In the menu at the top under "Resources" is "Download" which currently offers you the Windows client v2.1.4. If you don't currently have 2.1.4 – and I don't see how you could unless you've been manually keeping current – download and run it. The executable very nicely and seamlessly replaces any older release – even one covered in cobwebs like mine was – and sets itself up without any hassle.

And yes, though I have NAS's at both of my locations and I could easily arrange to self-host some file synchronizing system, or switch to using SyncThing which is a non-cloud peer-to-peer system that I've talked about before, and which I use to keep my wife's large fleet of client laptops remotely synchronized... for the time being at least I'll be remaining with sync.com. They pulled themselves out of the weeds, and this was the first time in many years that they had ever been in the weeds. So they get a pass this time.

Medibank reboot

I titled this piece "Medibank Reboot" because that might literally be what happened this past weekend... We've recently been following the drama with Australia's largest private medical insurer, Medibank. The one that exposed a great deal of the information of nearly 10 million past and current clients. Well, on Friday the Sydney (Australian) Morning Herald had an interesting bit of news. Their coverage began:

> *Private insurer Medibank's app, stores, contact centre and IT systems will go dark this weekend as it overhauls its cybersecurity following the nation's worst data breach in corporate history.*
>
> *From 8.30pm AEDT (Australian Eastern Daylight Time) on Friday, Australia's largest health insurer will* ***shut down its IT systems*** *followed by retail store and customer contact center closures on Saturday to "further strengthen systems and enhance security protections". The company expects normal activity to resume by Sunday at the latest.*
>
> *Microsoft IT security experts from the Asia-Pacific region will travel to Medibank's Melbourne headquarters to assist with the operation* [I guess Microsoft is going to show them where the "Update" button is located.]. *This was said to have been planned over several weeks and will be Medibank's first shutdown of such scale.*

In other words, shut things down, update software and firmware and everything else, then turn everything back on again. The Herald continued:

> *The overhaul is part of a series of maintenance strategies, termed "Operation Safeguard", established after the personal information of up to 10 million current and former Medibank customers was breached in a cyberattack.*
>
> *The data was released on the dark web when Medibank refused to pay a $15 million ransom demanded by the hackers, who police have said were based in Russia. The company said the damaging cyberattack will cost the firm at least $35 million in initial recovery costs, though that is likely to grow as law firms and regulators circle.*
>
> *A Medibank spokesperson said although there had been "no further suspicious activity" detected inside its systems since October, the insurer was carrying out further maintenance to strengthen its online security.*

If we can read between those lines, what might be happening is a complete wide scale coordinated reinstallation of system software. As we've often noted, once a complex system has been compromised you cannot ever really be certain that something isn't still hiding somewhere. The nature and circumstances of what the forensics examination found might have left them with no choice other than to wash everything clean and reinstall. If so, what a nightmare!

> *Since the hack, Medibank has bolstered monitoring, added detection and forensics capability across its system, and scaled up analytical support via specialist third parties. It also recently introduced two-factor authentication* [Oh! Imaging that!] *– where access is granted only after providing a code sent to one's email or SMS* [Oh :(So not good 2-factor authentication.]
>
> *During the shutdown, customers will be unable to access services for Medibank and its discount **ahm** brand online or in person, and instant electronic health claims will be unavailable.*
>
> *The spokesperson said: "We apologize for the inconvenience this operation may cause customers, but this is the next necessary phase of our ongoing work to further safeguard our network."*

And, again, being completely down and offline for as long as two days sure does sound like a major sweeping cleaning out of all mission critical systems.

Totally fake cryptocurrency trading platforms:

It was P.T .Barnum who is credited with the saying "There's a sucker born every minute." I was reminded of that when I came across this bit of news about a cybercrime group that's been named "CryptosLabs." The cybercrime research group known as "Group-IB" identified a new cybercrime operation which they named CryptosLabs. Get this: Since 2018, this CryptosLabs group has operated a network of more than 300 scam websites posing as fin-tech and cryptocurrency trading platforms. Group-IB says the group used search engine ads and social media posts to trick French-speaking users across Europe into investing more than €480 million in these websites – okay? Nearly half a billion Euros – by leading them to believe that they

would get to trade in stocks and crypto assets. But researchers say that once users put money into their accounts, the crooks either asked for more or ignored their customers before shutting down platforms and moving to a new domain. Group-IB said it named this group CryptosLabs after the kit they used to automate the deployment of the fake trading portals, which typically mimicked 40 popular banking, fin-tech, and crypto brands.

So these guys were, for example, pretending to be CoinBase and set up one fake cloned CoinBase site after another. People had heard of CoinBase, their friends were talking about it. Then they encountered an advertisement in a search engine or in online social media, and they thought *"Hey! I just got paid, now's the time!"* They didn't really know what CoinBase's domain was, so when they just clicked the advertisement and went to coinbases.org... and since everything looked quite official they never thought that it might be an illegitimate ad, and site.

People would transfer money in and at some point, once enough had, or someone wanted their money back, the fake domain would be shut down and another would be set up in its place. This is not the Internet that Tim Berners-Lee envisioned back in 1989 when he originated the concept of an Internet full of inter-linked HTML documents that anyone could create and publish on their own. It's not as if crime didn't exist before. It just flowed into a new medium.

There's two interesting pieces of news about Telegram:

Malware on Telegram
The Russian security firm Positive Technologies published a report on Telegram's budding cybercrime ecosystem. According to the company's scans, Telegram has slowly replaced hacking forums and is currently being used for advertising a wide spectrum of hacking services and malware. The sale of remote access trojans, corporate network account credentials, and cash-out services are among some of the most popular topics on the platform.
https://www.ptsecurity.com/ww-en/analytics/cybercriminal-market-in-telegram/

So that's one tidbit. But that's not all...

Telegram
Telegram, which we've just learned is generally becoming the favored hangout of the crime underworld, has decided to further expand their subscriber base by allowing users to sign up without needing one of those pesky SIM cards to anchor their identities. Telegram wrote: *"Today starts a new era of privacy. You can have a Telegram account without a SIM card and log in using blockchain-powered anonymous numbers available on the Fragment platform."*

So I thought, hmmmmmm... What's fragment? So I followed a link <https://fragment.com/> and was told: *"Ooops... This service is not available in the United States"* That's interesting. I wonder why not?

So I thought that Wikipedia might know about Fragment. And perhaps it does. But the word "fragment" is so common that I wasn't able to find it there among all of the other fragments. Googling turned up an abbreviated reference that was more tease than anything else. Google said: *"Buy and Sell Usernames. Secure your name with blockchain in an ecosystem of 700+ million users and assign it as a link for your personal account, ..."* Buy and sell usernames?

What? I know that Kevin Rose might be willing to sell you an icon of a zombie. But what's "Fragment"? So I dug some more and I found some news about fragment over at "crypto.news" where their coverage had the headline *"Telegram now allows users to buy and sell usernames via auction"* and then it goes on to explain...

> *Telegram releases new feature; transforms usernames into digital assets. Popular Cloud-based instant messaging app, Telegram has just launched a new feature to allow users to buy and sell short, recognizable "@ usernames" for personal accounts, public groups, and channels.*
>
> *Telegram has commenced an auction for the best usernames on Fragment, a free collectible trading platform. With this new feature, Telegram usernames have become digital assets that can be secured and sold between parties. According to the innovative platform's unveiling note, ownership of the collectible usernames is secured in the immutable ledger of TON, a fast and scalable blockchain network. Interestingly, the new feature allows owners to add multiple username aliases to their personal accounts, group, or channel. Also, each collectible name can be accessed with its @username on Telegram or outside Telegram using links such as username.t.me and t.me/username.*
>
> *To acquire usernames on Telegram, buyers visit Fragment, search for their desired username and click on auction if that username is still available. Buyers will then be redirected to a page which shows the highest bid along with the bid step and minimum bid.*

So earlier this said that TON was an immutable ledger. Apparently it's also a currency. I went over to https://coinmarketcap.com/currencies/toncoin/ And I learned that a TON has a current value of $2.10 US. It was fluctuating between $2.10 and $2.11 as I was looking at the page. It also has a 24 hour trading volume of $44,628,950. And I saw that there was a TON.org. So I went over there and discovered that TON stands for The Open Network. And from the TON homepage we learn that: *"TON is a fully decentralized layer-1 blockchain designed by Telegram to onboard billions of users. It boasts ultra-fast transactions, tiny fees, easy-to-use apps, and is environmentally friendly."*

Okay... So, let me get this straight. Telegram noticed that they had a lot of users and a popular platform. So they decided that they wanted to monetize the ownership of Telegram usernames. They wanted to create a marketplace which would allow Telegram usernames to be bought and sold. So they created TON, their own cryptocurrency, anchored with their own blockchain. They then established an auctioning system which uses the TON as its exchange currency to allow their users to bid for, purchase and sell Telegram usernames. The rest of the coverage of this, the first part of which I already shared, tells us how this is going. Under the heading: *"Massive instant adoption of new feature; millions of TONs earned in username sales"* We have...

> *Less than 6 hours after the launch, thousands of usernames featuring international brands and celebrity accounts have been put up for sale. Still on auction are @nike, @king, @esport while others such as @auto, @avia, @fifa etc., have been sold for as much as 900,000 TON.*
>
> *Judging by data on the Fragment platform, millions of TONs have been earned by Telegram users from the sales of their short usernames. There is still more to be made as there are still*

lots of usernames currently on auction. An example is the popular shoe brand @nike which has over 300,000 TON bid at the moment. Telegram is affording its users full ownership of their usernames, and they are embracing the idea.

Telegram wrote: *"For the first time in the history of social media, people have full ownership of their usernames. Long-time Telegram users who have been using short usernames they registered early on can now benefit from the platform's growth by selling their usernames in fair, transparent, fully decentralized auctions."*

Wow. It's unclear what this has to do with not needing a SIM card. But apparently it would be possible to purchase a name that you want to use for some number TON, have that transaction locked into a blockchain, after which Telegram will allow you to use that as your immutable identity rather than typing you to a phone's SIM card.

In retrospect it's quite clever. We've seen how desirable short Twitter handles have been. Those with very short handles have needed to be extra vigilant in their protection. So now here's Telegram, another increasingly popular instant messaging platform which has grown into a social media platform. And identity has value. This system monetizes that value.

On the topic of social media...

The TikTok banning continues with Texas joining the ranks. I was going to say that "Texas was the latest...", but that was yesterday so perhaps other states have already followed. But in any event, Texas Governor Greg Abbott has banned the use of TikTok on the devices of state employees and in doing so becomes the 4th Republican-led state to ban TikTok on employee devices following Maryland, South Carolina, and South Dakota. (It was South Dakota that I noted last week being first with their governor Kristi Noem saying that she hoped other states would be following suit. I guess her wish is coming true.) Greg Abbott also ordered state agencies to come up with plans to govern the use of TikTok on state employees' personal devices not owned by the state. Yikes! Abbott wrote in a letter to state agency leaders that there are *"growing threats posed by TikTok"* to the states' sensitive information. *"TikTok harvests vast amounts of data from its users' devices – including when, where, and how they conduct Internet activity – and offers this trove of potentially sensitive information to the Chinese government."*

But wait, there's more...

Indiana's Attorney General brought a pair of lawsuits against TikTok, accusing the company of deceiving users by claiming that their data was protected from the Chinese government and for exposing Indiana children to adult content. The lawsuits claim that the China-based social media giant has deceived and harmed Indiana residents. Indiana's first lawsuit alleges TikTok has marketed its video-sharing platform as safe for teens, even though its algorithm *"serves up abundant content"* depicting drugs, sexual content and other inappropriate themes. The second lawsuit asserts that TikTok has deceived consumers by suggesting their personal information is protected from the Chinese government and Communist Party….

The Attorney General said in a statement: *"The TikTok app is a malicious and menacing threat unleashed on unsuspecting Indiana consumers by a Chinese company that knows full well the harms it inflicts on users. With this pair of lawsuits, we hope to force TikTok to stop its false, deceptive and misleading practices, which violate Indiana law."* The AG wants emergency injunctive relief against the company and is seeking monetary penalties for every time TikTok violated Indiana's Deceptive Consumer Sales Act. **Wow.**

After digesting all of this huffing and puffing I became curious to find a real security company's threat assessment of TikTok. Kaspersky has a very sane and sober writeup which directly addresses these issues and concerns. First, to lay a bit of foundation: *"How does TikTok work?"*

TikTok users sign up with a phone number, email address, or a third-party account like Facebook or Instagram. Once logged in, they can search popular creators, categories, or hashtags to find videos. They can also use their phone contacts or social media connections to find friends who are already using the app.

TikTok makes content discovery a big part of its experience, enabling videos to go viral rapidly. Central to this is its algorithm, which uses AI to make personalized recommendations for viewers. These recommendations appear in the 'For You' feed for each user.

Users can create and watch videos using features such as filters, split screens, green screens, transitions, emojis, stickers, GIFs, and more. Music is a crucial aspect of TikTok – the app has an extensive music library and is integrated with Apple Music, allowing users to add, remix, save, and discover songs.

Users can follow accounts they like and share or give hearts, gifts, and comments on videos they have enjoyed. Coins are used to give virtual gifts on TikTok. There are multiple ways to use and enjoy TikTok – and many users don't create videos but only watch them – but some of the most popular genres are:

- *TikTok Duets – ranging from genuine collaborations to remixes to spoofs.*
- *TikTok Challenges – usually involving a popular song or hashtag, where participants attempt dance moves or their own variation on a theme.*
- *Cringe videos – awkward performances which are embarrassing but funny.*
- *Reaction videos – where users record their reaction to another video.*
- *Many TikTokers with large followings monetize their videos or promote their brand through specialized ads, merchandise, and partnerships.*

Okay. So far, it's your average run of the mill, massively networked, social media short-form video sharing platform. Then Kaspersky directly addresses the question of privacy concerns:

One of the most viral aspects of TikTok has been privacy concerns, with questions like "What data does TikTok collect?" and "Does TikTok steal your information?" regularly circulating online. Like many other social networking platforms such as Facebook, TikTok collects a lot of information about its users. This includes:

- *Every TikTok video watched, and for how long.*
- *The entire contents of every message sent through the app since messages are not encrypted.*

- *The user's country location, internet address and type of device being used.*

With the user's permission, TikTok also captures:

- *Its users' exact location*
- *Their phone's contacts and other social network connections*
- *Their age, phone number, and payment information*

This information can be used to assemble a detailed profile for advertisement targeting – by understanding who its users are, who their friends and family are, what they like and find entertaining, and what they say to their friends.

To use the app, users grant access to their microphone and camera. If they create videos, the app captures close-ups of their face.
Potentially, this provides biometric data which could be used in conjunction with other images which exist online. TikTok also uses technical measures to encode its activity. This means that some of what it does is hidden from external researchers. TikTok says this is to disrupt hackers and other malicious actors.

There has been extensive media coverage of TikTok privacy concerns. However, most social media platforms worldwide collect, use, analyze and ultimately profit from users' personal data. TikTok argues that it collects less data than platforms such as Facebook or Google since it does not track user activity across devices.

Okay. So far this reads exactly like any other of the many social media platforms. Certainly our own FaceBook headquartered here in the U.S. is no better.

To get a balanced look at the government's side, NPR provided some coverage from less than a month ago about the FBI's raising of concerns over what TikTok might be capable of doing. NPR's headline was *"The FBI alleges TikTok poses national security concerns"*:

The head of the FBI says the bureau has "national security concerns" about the U.S. operations of TikTok, warning that the Chinese government could potentially use the popular video-sharing app to influence American users or control their devices.

FBI director Christopher Wray told a House Homeland Security Committee hearing about worldwide threats on Tuesday that the FBI has "a number of concerns" just days after Republican lawmakers introduced a bill that would ban the app nationwide.

Wray said: "They include the possibility that the Chinese government could use it to control data collection on millions of users or control the recommendation algorithm, which could be used for influence operations if they so chose, or to control software on millions of devices, which gives it an opportunity to potentially technically compromise personal devices."

[NPR reminds us that] *TikTok, which hit 1 billion monthly active users in September 2021, is owned by the Chinese company ByteDance. Chinese national security laws can compel foreign and domestic firms operating within the country to share their data with the government upon request, and there are concerns about China's ruling Communist Party using this broad authority to gather sensitive intellectual property, proprietary commercial secrets and personal data.*

TikTok has long said that it stores U.S. user data within the U.S. and does not comply with Chinese government content moderation requirements. But the company has come under increasing scrutiny in recent months, and in July it acknowledged that non-U.S. employees did in fact have access to U.S. user data.

Citing leaked meeting audio, Buzzfeed News reported in June that China-based ByteDance employees have repeatedly accessed non-public data (like phone numbers and birthdays) of U.S. TikTok users. Separately, Forbes reported in October that ByteDance planned to use TikTok "to monitor the personal location of some specific American citizens," which the company denied.

I was curious to learn what information Forbes had, since in the past they've shown a tendency to sensationalize and exaggerate. In their October coverage, which did indeed make those allegations, Forbes wrote: *"Forbes is not disclosing the nature and purpose of the planned surveillance referenced in the materials in order to protect sources."* – Nor have they disclosed any of those materials or provided any evidence to support their story. But it certainly generated some clicks and some quotes, which may have been their goal. Finishing NPR's coverage...

Wray said at the hearing, that Chinese law essentially requires companies to "do whatever the government wants them to in terms of sharing information or serving as a tool of the Chinese government." – "And so that's plenty of reason by itself to be extremely concerned."

The FBI has in the last few years been shifting its focus to China. In July, Wray said China was the "biggest long-term threat to our economic and national security" and accused Beijing of having interfered in recent elections.

So all that the FBI's director said in the middle of last month was that this or that was **possible**. He said that installing the TikTok app created *"an opportunity to potentially technically compromise personal devices."*

I wanted to look around to make sure that there wasn't some smoking gun somewhere that we hadn't heard about regarding TikTok. I couldn't find one. At this point it appears to be a handful of state governors crusading against the Red communist threat. I'm not suggesting that none of this is possible. Our global community has collectively created the technology to make it **all** possible. But so far this looks like a politically driven overreaction for grandstanding.

And while we're on the subject of nonsense...

The LastPass class action lawsuit

Filed two days after LastPass' November 30th disclosure: Password management company LastPass has been hit with a class action lawsuit after experiencing two data breaches in the past three months. The plaintiff in the case *"Debt Cleanse Group Legal Services LLC"* filed the class action complaint against LastPass US LP on Dec. 2 in a Massachusetts federal court, alleging negligence and breach of contract. The Chicago-based debt relief firm said it used LastPass to manage its passwords. However, it says LastPass was negligent in its duties, evidenced by the fact it has experienced two data breaches in the past three months.

The class action lawsuit alleges that LastPass was negligent with data security, stating that LastPass used ineffective data security measures to protect its customers' information. The lawsuit states: *"There is a strong probability that entire batches of stolen information have been dumped on the black market, or are yet to be dumped on the black market, meaning plaintiff and class members may be at an increased risk of fraud and identity theft for many years into the future."*

DebtCleanse responded to news of the breach by changing all of its employees' passwords for accounts that used LastPass, which took considerable resources, it says. It seeks to represent all LastPass users whose information was accessed in the breach. It seeks certification of the class action, damages, fees, costs and a jury trial.

So, the geniuses over at DebtCleanse freaked out and totally unnecessarily changed all of their passwords for everything. They noticed that LastPass made an announcement of a secondary breach following on from the first one. But they somehow failed to heed the statement – which was also right there, both times – that no user information or passwords were at any time at risk or exposed by either of those two breaches.

I know that Leo and I feel similarly about class action lawsuits. From time to time I'm awarded something like 92 cents from something I purchased once that faded too quickly if it was left out in the sun – or whatever. Somewhere some lawyers made some money. The individuals in the class on whose behalf the action was taken... netted 92 cents.

Rackspace had a big embarrassing problem

A tiny piece of Rackspace's overall cloud hosting business — hosted Microsoft Exchange e-mail services — was hit eleven days ago by ransomware. Things have not been going well for them ever since. An article published yesterday in IT World Canada summed things up and offered some interesting perspective on the cloud industry overall. Here's what they wrote:

On December 2nd, Rackspace experienced an outage for its Hosted Exchange environment. The company blamed a "security incident." A status update issued by the company noted, "We proactively shut down the environment to avoid any further issues while we continue work to restore service." [That really stretches the phrase: "putting a good face on the problem." Were they deeply embarrassed by this? Probably. Should they just have said "we got hit with ransomware right from the start? Yeah, probably."] Anyway, back to IT World Canada:

One week later, the outage continues, ***and*** *the company has confirmed that it* ***is*** *due to a ransomware attack. Rackspace has not indicated how much data might be lost, whether it will pay the ransom, or when the managed exchange service will resume. This is the only information from the section of the website dealing with the press.*

However, in an announcement on its investors page, the company notes that the hosted exchange business accounts for less than ***one per cent*** *of the company's revenue* [which, by the way, is a hefty $3 billion per year!] *and reassures investors that the company has cyber insurance.* [Oh... you're gonna be needing it, honey. We'll get to that in a minute.]

But the attempt to reassure investors may not be working. In an article on December 10th – so last Saturday – *investment blog MarketWatch criticized the company for being "frustratingly closed mouthed" about the incident, and noted that the company's stock price had declined.*

The article notes, "Since the incident came to light, Rackspace shares have tumbled by a third. This is a relatively small part of the company's business, only about US$30 million a year in revenue. But the bad blood that Rackspace is generating will leave a lasting stain."

The stinging critique of the company's communication is significant, but another quote from the article raises an issue that could extend beyond Rackspace to the entire cloud industry. The writer notes, "While I remain a big believer in cloud computing, the Rackspace attack is an urgent reminder of the risks in relying on the cloud for mission-critical applications if your provider isn't keeping up with software patches and paying attention to security risks."

The use of cloud computing, even for mission critical applications, has grown rapidly for years, but that growth has accelerated in the past year and is predicted to further accelerate in the next 24 months.

[I so clearly recall the looks I received from the IT guys who attended DigiCert's customer advisory board meeting six years ago. I made some offhand comment about the rack of gear I had over at Level 3 which brought all discussion to a halt. I said "what???" and one of them replied, apparently for the entire group: "Nobody does hardware anymore." Oh. I do. I like hardware... with lots of blinky lights.]

Senior management has bought into cloud in a big way. But could investor nervousness from the Rackspace outage have an impact on attitudes in the boardroom? When a service that gives you one per cent of your revenue leads to a drop of 30 per cent in your share price, cloud proponents may, to quote Ricky Ricardo, have some "splainin' to do."

And that was just the first shoe to drop. Here comes the second one...

Yes, Rackspace is now facing not one, and not two, but at least three class action lawsuits.

Reports are that Rackspace will need to defend themselves in (so far) at least three different class-action lawsuits related to this recent ransomware attack from which they are still recovering and may never be able to fully recover from. The attack left countless companies without access to their cloud-hosted email servers. In an interview last week, Rackspace suggested they might not be able to recover all of their customers data, which they referred to as "legacy data." The company also appears to have given up on hosting Exchange email servers in its cloud and said it was migrating all its existing customers to Microsoft 365 which, according to documents Rackspace filed with the SEC will cost the company $30 million – thus just 1% of their $3 billion annual revenue. But who cares? It's certainly not worth the headache. I imagine that hosting Exchange Server was more of a means to an end. A way of establishing a relationship with an enterprise and then over time, moving more of their non-eMail business into the cloud. How's that working out?

The trend appears to be that **everyone** is getting very tired of the consequences of these never ending attacks. Governments are going on the offensive, saying that they're no longer going to wait to be attacked, other governments are calling their own employees criminally negligent and bringing them up on criminal charges, and the class action ambulance chasers now only take a day or two to file their lawsuits. Why is all of this happening all of a sudden? Because criminal organizations, apparently some with state level sponsorship and protection, have realized that the cryptocurrency boom, coupled with the presence of exchanges to local fiat currencies, provides them with a means of being anonymously paid. That provided the motivation. Clever hackers provided the means, and the final piece, endless opportunity, was readily supplied by the industry's historically lax cybersecurity.

So today, everyone is furious, besides themselves, pointing fingers and suing. But who's still never being held to account? By a perversion inherent in the system we've built, the suppliers of the buggy, brittle and breakable systems are never to blame. Stay tuned. That's going to be next. Juries are made up of end users and they are not going to be sympathetic after everything this industry has been putting them through.

And speaking of losing patience...

Another country goes on the offensive.
On December 11th, it was reported that Japan intends to establish a legal framework that will allow strengthening of defense measures in cyberspace, including the ability to attack preemptively. The Japanese government intends to change the rules so that it can start tracking potential attackers and breaking into systems as soon as there is a potential threat. Current regulations make it extremely difficult to apply such measures unless there is an emergency situation that requires the mobilization of defense forces after a military attack. The plan is reflected in a proposal to amend the National Security Strategy that has just been submitted to the governing coalition. The draft amendment is expected to be approved by the cabinet before the end of this month. So, we're clearly seeing a changing of attitudes across the board.

Closing The Loop

The following was sent to me via Twitter DM and I'm not sure this individual would want his name published so I'm not sharing it. But get this, he's the Australian cybersecurity person who was tasked with finding and downloading the Medibank data from the Dark Web for analysis.

He sent the following:

Hi Steve, I just wanted to drop you a line about the Medibank data. You mentioned in Security Now 900 that you didn't know what use that data would be if it wasn't formatted properly. I'd like to offer one way it could be useful in the wrong hands.

I work in cyber security for a law enforcement government department in Australia and I was tasked with finding and downloading the Medibank (and Optus) data.

I absolutely agree, the data held within the extracted data was not formatted and was of limited use in its raw format (all CSV files with different structures presumably depending on the database it was extracted from). It's true that the data is of limited use as the medical

information about the patients was all in code format which meant nothing without the application to match the code to the treatment the patient received.

However, my job was to search through the data to see what law enforcement officer's details were leaked. Can you imagine how bad it would be if an individual used their government issued email address (which has their law enforcement division within the domain of the email address) to sign up with Medibank? Suddenly you would be able to match a name and home address of a person, and verify it was a person working for the law enforcement division thanks to the email address. In the hands of the wrong people (such as criminal gangs) this information could put these law enforcement officers in serious danger. Sadly, data of this nature was leaked and our organization now has to help protect these individuals. I'm sure you've already considered this scenario with leaked data (really any leaked data could be used this way) but considering the high profile nature of this leak, we're having to take it very seriously. Thank you for such a wonderful podcast.

And thank **you** for the very interesting note from the field! It's abundantly clear that this podcast has aggregated quite an amazing group of listeners. I am continually humbled and flattered by the people who take the time to listen to these ramblings. Thank you, all.

Paul Jolley / @j0113y

I know you've spoken of UTF-8 in past episodes, but was left asking myself what could possibly be the legitimate purpose for some characters after I encountered U+200B recently.

I appreciate English isn't the only language used in the world so understand why this extended character set is useful but (as previously discussed on the podcast) can also see how it can be abused, typo-squatting in domain names for example by using characters that look like their ASCII equivalents.

Getting back to my recent experience, I received an email where the sender name would not sort alphabetically in my mail client. I didn't think much of it at first but then decided to investigate. It turns out they had inserted a "Zero Width Space" character before the first letter of their name, making it appear at the top of my sorted list. The 3 bytes could be seen in a hex editor as E2 80 8B but would not visibly appear in Notepad.

I've never encountered this kind of whitespace before and thought it would be worth mentioning because I expect something like this will only be used for mischief. Try it yourself by renaming a folder in Windows and putting this special character at the beginning to make your folder appear at the top of a sorted list without any visible whitespace, or paste it in Notepad a few times and see the file size increase when you save the file but there's nothing to select or highlight on screen. [And Paul concludes...] *The world doesn't need this!*

Nope. And the lesson here is that for every useful thing we create there are clever people who will subvert that use into abuse.

Jona / @jachness

Hi Steve, long time SN listener. Regarding the last LastPass breach, the zero knowledge architecture keeps passwords safe but what if the base source code of the browser client gets compromised in a way that it exfiltrates the passwords once the user decrypts the site's passwords?

Yes. That is the danger. Client side compromise where the passwords must be in the clear. If Joe Siegrist who originally wrote this stuff was still at LastPass I could ask him whether the entire blog of encrypted data is decrypted at once or whether it's granulated. I can't see any reason why the individual pieces could not be individually encrypted and thus individually decrypted in order to keep everything else always safe.

And remember that this remains a problem that passkeys doesn't cure. Passkeys moves us closer toward the right solution. And it's worthwhile. But clients still need to synchronize and store a large and growing number of private keys which they **must** keep absolutely secret. Not to harp on this, but a huge advantage of SQRL, in fact the reason I originally got off on that tangent, was that SQRL only uses one single secret – ever – and it derives each of its private keys from that one secret coupled with the authentication domain. That makes key management so much more practical. It eliminates synchronization issues and it makes it possible to protect the key in hardware. It's far more difficult to strongly protect a growing, amorphous and ever changing collection of keys in hardware.

Anyway, I imagine we'll get there someday and if future researchers start poking around looking for systems that have solved the related challenges, SQRL will be there waiting in the wings.

And, while we're on the subject of my projects...

SpinRite

SpinRite is currently at alpha release 5 with many further improvements already implemented for its next alpha release. Everything is going very well. Thanks to a truly gratifying level of engagement – we now have exactly 500 registered users in GitLab – we've been uncovering mysteries around the edges that I've been working to solve, and also things that I didn't anticipate.

For example, in 6.1 it's possible to select a drive from the command line by model number. But Western Digital's model numbers contain spaces, and spaces are used as a delimiter to separate command line entities. The answer is simple, right? Allow tokens to be quoted to have their contents parsed literally. When I wrote the command line parser I forgot to add support for quoting literals. Someone booted SpinRite from a CD and it didn't notify them that their request for logging to the CD could not be honored. So, those sorts of things. Things I want to take the time to get correct now, so that I don't have to correct them later.

And we do have a couple of mysteries. One guy has two HP All-In-One PCs where, if he warm boots from Win10, SpinRite locks up when it starts running. But if he cold boots the machine everything works fine. Now, we might say "Okay, so don't warm boot." And that might be the

final answer, but only if there's nothing that can be done. I don't yet understand **why** that's happening and whether it might be a symptom of something more important. So, I found one of those machines on eBay for $47 and it's in my car's trunk right now. I'll figure out what's going on and I'll answer the question that's best expressed using one of my favorite SpinRite development abbreviations which now comes up frequently: WTF?

Miscellany

I also have one piece of miscellany to share. The ZimaBoard fanless single board computer that I stumbled upon and told everyone about, has apparently been a big hit. I keep seeing mentions of it in postings among those who are testing SpinRite and also in Twitter messages. So, given how long it takes to obtain one after ordering, and the fact that shipping from China is not free, I was surprised and delighted to learn a few hours ago that Amazon has all of them in stock and available for astonishing same day arrival, at least to my location. As an Prime subscriber, the cost is $10 higher than ordering the ZimaBoard directly from its source, so $129 versus $119. But then there's no shipping and you can have today. I know that not everyone bothers being a Prime subscriber, but it's there nevertheless. So I just wanted to share that bit of happy news.

Apple Encrypts the Cloud

Last Wednesday, Apple posted the news which generated the most feedback to me of anything that's happened recently. So I knew we needed to address that today. Before that, I was planning to talk about the technology behind a clever new generic bypass of web application firewalls — nerdy but cool. Don't worry, we'll get to that next week. Today, is Apple's news.

So Apple's headline read: *"Apple advances user security with powerful new data protections"* the subhead which introduces three new terms, reads: *"**iMessage Contact Key Verification**, **Security Keys for Apple ID**, and **Advanced Data Protection for iCloud** provide users with important new tools to protect their most sensitive data and communications"*

So we have three new things. Apple explains that **iMessage Contact Key Verification** allows users to verify they are communicating only with whom they intend. We'll see what that means in a minute. Then there's **Security Keys for Apple ID** where users have the choice to require a physical security key to sign in to their Apple ID account. So that sounds interesting. And with **Advanced Data Protection for iCloud** users will reportedly finally obtain what Apple says is end-to-end encryption to provide Apple's highest level of cloud data security where users will have the choice to enhance the protection of important iCloud data, including iCloud Backup, Photos, Notes, and more.

Before we go any further I should note that none of this is available yet. The enhanced iCloud encryption is currently available to beta testers and Apple plans to roll it out to everyone in the US by the end of the year, which is just weeks away. I was in the iOS 16 beta program since I wanted early access to passkeys. But later, as iOS kept evolving, I became annoyed with the continual beta updates and restarts, so I resigned from the beta program. It doesn't look like

any of us in the US will need to wait long. And Apple said that Advanced Data Protection for iCloud would be rolling out across the rest of the world starting early next year. And that's also the case for the new iMessage Contact Key Verification and Security Keys for Apple ID. Both are slated for deployment in 2023.

Let's look at each of those three security improvements.

The good news is that the most information is available about the thing most people are most excited about, which is "not even Apple can get in" true user encryption of user data stored in the cloud. Thus the title of this podcast. So, we'll get to that part last. The bad news is that neither of the other two security improvements is explained in detail, but perhaps there isn't much detail to be had.

Of iMessage Contact Key Verification Apple only said:

> *Conversations between users who have enabled iMessage Contact Key Verification receive automatic alerts if an exceptionally advanced adversary, such as a state-sponsored attacker, were ever to succeed in breaching cloud servers and inserting their own device to eavesdrop on these encrypted communications. And for even higher security, iMessage Contact Key Verification users can compare a Contact Verification Code in person, on FaceTime, or through another secure call.*

And that's all we know presumably until it's rolled up early next year. They provide a photo showing a little Emergency exclamation triangle icon underneath which is written: *"An uncognized device may have been added to {account's name} account.* ***Options...****"* and we don't know what one's options might be.

So, this covers one of the two theoretical ways that we know iMessage can be subverted. Since iMessage conversations are broadcast to all of an account's registered and logged-in devices, if an adversary were to somehow get their own additional device added to a user's account, that additional device would automatically be privy to all of the attacked accounts messaging. So this seems like a useful feature.

But since Apple handles all key management for its users, the iMessage attack we've always wondered and worried about is law enforcement serving Apple with a wiretap subpoena which would compel Apple to surreptitiously add an additional key into an iMessage conversation. Since iMessage can already handle multi-party messaging, adding a hidden key seems eminently plausible. And it's unclear under what grounds Apple could refuse a legal order for something they could probably do. But that's all speculation for which there's no evidence. The issue arises, however, anytime any 3rd party manages keying on behalf of its users. It's convenient but less secure.

The other piece of this announcement that we'll apparently need to wait to see, is this means for somehow comparing and confirming what Apple calls a Contact Verification Code in person, on FaceTime or through another secure call. This CVC is presumably (and hopefully) a hash of the user's public key which matches their private key. Note that this is exactly what my favorite messaging agent, Threema, has understood was important and has always done from the start.

So that's all we know for now about iMessage Contact Key Verification. I'm sure we'll talk about it again next year once it's rolled out.

Next up is "Security Keys for Apple ID". Apple explains it by writing:

> *Apple introduced two-factor authentication for Apple ID in 2015. Today, with more than 95 percent of active iCloud accounts using this protection, it is the most widely used two-factor account security system in the world that we're aware of.*
>
> I should clarify that what they're talking about here is just the use of some other device to receive a 6-digit code to authenticate a user.
>
> *Now with Security Keys, users will have the choice to make use of third-party hardware security keys to enhance this protection. This feature is designed for users who, often due to their public profile, face concerted threats to their online accounts, such as celebrities, journalists, and members of government. For users who opt in, Security Keys strengthens Apple's two-factor authentication by requiring a hardware security key as one of the two factors. This takes our two-factor authentication even further, preventing even an advanced attacker from obtaining a user's second factor in a phishing scam.*

So, hardware authentication dongles. That's good. I imagine that Stina and the team over at Yubico is happy to see 3rd-party hardware dongles becoming far more mainstream. This will likely make a huge difference in this market. For less than $50 you can get an external key which must be present anytime you would normally need to provide your account's passcode to unlock your iPhone or iPad. That's going to be popular.

Okay, finally, Encrypting the Cloud

Most of the content of this announcement page is glitz and self-congratulatory hype. The good news is that there's a second page available containing much more detail. But since it leaves out some of the broad strokes we'll cover them here, first.

Ivan Krstić who is Apple's head of Security Engineering and Architecture is quoted here saying: *"Advanced Data Protection is Apple's highest level of cloud data security, giving users the choice to protect the vast majority of their most sensitive iCloud data with end-to-end encryption so that it can only be decrypted on their trusted devices."* For users who opt in, Advanced Data Protection keeps most iCloud data protected even in the case of a data breach in the cloud.

So Apple is naturally not saying that this keeps **their** hands off any of their users' data, which is of course the main reason everyone wants this, but rather to protect their users in the event of an iCloud security breach. Fine, whatever, just let us have it. They explain, and here's the interesting bit:

> *iCloud already protects 14 sensitive data categories using end-to-end encryption by default, including passwords in iCloud Keychain and Health data.* [Meaning that currently Apple cannot get at that.] *For users who enable Advanced Data Protection, the total number of data*

categories protected using end-to-end encryption rises [from 14] *to 23, including iCloud Backup, Notes, and Photos. The only major iCloud data categories that are not covered are iCloud Mail, Contacts, and Calendar because of the need to interoperate with the global email, contacts, and calendar systems.*

But iCloud Backup, that's the biggie, right? Remember that back in 2016 after the San Bernadino shootings, it was Syed Farook's phone whose iCloud Backups, had they been available, and had the FBI not changed the password on the account, could have been decrypted to provide the FBI with the vital evidence that they say they needed at the time.

Anyway, just to wrap up this overview, Apple is expending some significant effort to spin this additional encryption as user protection in the event of a breach, presumably to deflect some of the government's and law enforcement's annoyance over more stuff being encrypted. Apple went out of their way to write:

Enhanced security for users' data in the cloud is more urgently needed than ever before, as demonstrated in a new summary of data breach research, "The Rising Threat to Consumer Data in the Cloud," published today. Experts say the total number of data breaches more than tripled between 2013 and 2021, exposing 1.1 billion personal records across the globe in 2021 alone. Increasingly, companies across the technology industry are addressing this growing threat by implementing end-to-end encryption in their offerings.

Okay, so that was the broad strokes, what more is known about this at this time? This is all good stuff so I'm just going to share it:

Advanced Data Protection for iCloud *is an optional setting that offers Apple's highest level of cloud data security. When a user turns on Advanced Data Protection (henceforth ADP), their trusted devices retain sole access to the encryption keys for the majority of their iCloud data, thereby protecting it with end-to-end encryption. For users who turn on ADP, the total number of data categories protected using end-to-end encryption rises from 14 to 23 and includes iCloud Backup, Photos, Notes and more.*

Conceptually, ADPis simple: All CloudKit Service keys that were generated on device and later uploaded to the available-after-authentication iCloud Hardware Security Modules (HSMs) in Apple data centers ***are deleted from those HSMs*** *and instead kept entirely within the account's iCloud Keychain protection domain. They are handled like the existing end-to-end encrypted service keys, which means Apple can no longer read or access these keys.*

ADP also automatically protects CloudKit fields that third-party developers choose to mark as encrypted, and all CloudKit assets.

When the user turns on Advanced Data Protection, their trusted device performs two actions:

First, it communicates the user's intent to turn on ADP to their other devices that participate in end-to-end-encryption. It does so by writing a new value, signed by device-local keys, into its iCloud Keychain device metadata. Apple servers cannot remove or modify this attestation while it gets synchronized with the user's other devices.

*Second, the device initiates the removal of the **available-after-authentication** service keys from Apple data centers. As these keys are protected by iCloud HSMs, this deletion is immediate, permanent, and irrevocable. After the keys are deleted, Apple can no longer access any of the data protected by the user's service keys. At this time, the device begins an asynchronous key rotation operation, which creates a new service key for each service whose key was previously available to Apple servers. If the key rotation fails, due to network interruption or any other error, the device retries the key rotation until it's successful.*

[This is very cool and very serious. Apple is saying, since we once had those keys we want all of those keys to be rotated out of service so that even though we've already deleted the keys you're now using new keys that we've never seen.]

After the service key rotation is successful, new data written to the service cannot be decrypted with the old service key. It's protected with the new key which is controlled solely by the user's trusted devices, and was never available to Apple.

[Notice what's just been quietly acknowledged here: This is has nothing to do with breach protection. This is all about Apple strongly selling the truth that they no longer have access to their users' iCloud device backups, photos and notes.]

When a user first turns on ADP, web access to their data at iCloud.com is automatically turned off. This is because iCloud web servers no longer have access to the keys required to decrypt and display the user's data. The user can choose to turn on web access again, and use the participation of their trusted device to access their encrypted iCloud data on the web.

After turning on web access, the user must authorize the web sign-in on one of their trusted devices each time they visit iCloud.com. The authorization "arms" the device for web access. For the next hour, this device accepts requests from specific Apple servers to upload individual service keys, but only those corresponding to an allow list of services normally accessible on iCloud.com. In other words, even after the user authorizes a web sign-in, a server request is unable to induce the user's device to upload service keys for data that isn't intended to be viewed on iCloud.com, (such as Health data or passwords in iCloud Keychain). Apple servers request only the service keys needed to decrypt the specific data that the user is requesting to access on the web. Every time a service key is uploaded, it is encrypted using an ephemeral key bound to the web session that the user authorized, and a notification is displayed on the user's device, showing the iCloud service whose data is temporarily being made available to Apple servers.

ADP and iCloud.com web access settings can be modified only by the user. These values are stored in the user's iCloud Keychain device metadata and can only be changed from one of the user's trusted devices. Apple servers can't modify these settings on behalf of the user, nor can they roll them back to a previous configuration.

[Again, this is not about protection from data breaches. It's all obviously about Apple going well out of their way to demonstrate exactly how they no longer have access.]

In most cases, when users share content to collaborate with each other—for example, with shared Notes, shared Reminders, shared folders in iCloud Drive, or iCloud Shared Photo Library—and all the users have ADP turned on, Apple servers are used only to establish sharing but don't have access to the encryption keys for the shared data. The content remains end-to-end encrypted and accessible only on participants' trusted devices. For each sharing operation, a title and representative thumbnail may be stored by Apple with standard data

protection to show a preview to the receiving users.

Selecting the "anyone with a link" option when enabling collaboration will make the content available to Apple servers under standard data protection, as the servers need to be able to provide access to anyone who opens the URL.

iWork collaboration and the Shared Albums feature in Photos don't support Advanced Data Protection. When users collaborate on an iWork document, or open an iWork document from a shared folder in iCloud Drive, the encryption keys for the document are securely uploaded to iWork servers in Apple data centers. This is because real-time collaboration in iWork requires server-side mediation to coordinate document changes between participants. Photos added to Shared Albums are stored with standard data protection, as the feature permits albums to be publicly shared on the web.

The user can turn off ADP at any time. If they decide to do so:

1. The user's device first records their new choice in iCloud Keychain participation metadata, and this setting is securely synchronized to all their devices.

2. The user's device securely uploads the service keys for all ***available-after-authentication*** *services to the iCloud HSMs in Apple data centers. This never includes keys for services that are end-to-end encrypted under standard data protection, such as iCloud Keychain and Health.*

The device uploads both the original service keys, generated before ADP had been turned on, and the new service keys that were generated after the user turned on the feature. This makes all data in these services accessible after authentication and returns the account to standard data protection, where Apple can once again help the user recover most of their data should they lose access to their account.

The requirements to turn on Advanced Data Protection for iCloud include the following:

- *The user's account must support end-to-end encryption. End-to-end encryption requires two-factor authentication for their Apple ID and a passcode or password set on their trusted devices.*

- *Devices where the user is signed in with their Apple ID must be updated to iOS 16.2, iPadOS 16.2, macOS 13.1, tvOS 16.2, watchOS 9.2, and the latest version of iCloud for Windows. This requirement prevents a previous version of iOS, iPadOS, macOS, tvOS, or watchOS from mishandling the newly-created service keys by re-uploading them to the available-after- authentication HSMs in a misguided attempt to repair the account state.*

- *The user must set up at least one alternative recovery method—one or more recovery contacts or a recovery key—which they can use to recover their iCloud data if they lose access to their account.*

If the recovery methods fail, such as if the recovery contact's information is out of date, or the user forgets them, ***Apple cannot help recover the user's end-to-end encrypted iCloud data.***

ADP for iCloud can be turned on only for Apple IDs. Managed Apple IDs and child accounts (varies by country or region) aren't supported.

I am 100% very impressed. We heard what you need hear which was that if you lose your device or forget your password and are unable to authenticate, and none of the emergency recovery methods we made you set up when you turned on Advanced Data Protection available after we made you jump through all those hoops and acknowledge that your first born might be up for sacrifice**...** there is truly nothing we can do to help. You turned it on. We made you jump through hoops. You acknowledged the risks. And now it's all on you.

Yay. It sure took us a long time to get here. But here we are at long last.

EVERYTHING that Apple is storing for us is encrypted in our devices before it ever leaves them. They're storing keys in keychains, but they cannot decrypt those keychains because that's the one ultimate key that's being held by the user. So Apple has finally been willing to move the rest of the keys that they were holding into the keychain to which only its users have access.