Apple to stop storing Siri audio after contractors heard private talks and sex

Enlarge / An Apple Watch Nike+ Series 4.
Getty Images | Wachiwit

reader comments

33 with 26 posters participating

Share this story

Apple will stop storing voice recordings captured by Siri-enabled devices except in cases when users intentionally opt in to the practice. “[B]y default, we will no longer retain audio recordings of Siri interactions,” Apple said in an announcement yesterday. “We will continue to use computer-generated transcripts to help Siri improve.”

The change won’t take effect right away. The default storage of Siri recordings will stop “with a future software release in fall 2019,” Apple said in a support document.

Historically, Apple stored Siri recordings by default and had humans review “less than 0.2 percent” of the audio samples “to measure how well Siri was responding and to improve its reliability,” Apple said.

Apple previously paused human

review

This is the second major change made by Apple since a July report that its contractors who review the recordings for accuracy heard private discussions and even sexual encounters. A month ago, Apple said it was temporarily suspending human reviews of Siri voice recordings and that it would resume human reviews only after giving customers the choice of whether to opt in to the practice.

But until this week, Apple hadn’t promised to stop storing the voice recordings by default.

“We know that customers have been concerned by recent reports of people listening to audio Siri recordings as part of our Siri quality evaluation process–which we call grading,” Apple said in its announcement yesterday. “We heard their concerns, immediately suspended human grading of Siri requests and began a thorough review of our practices and policies. We’ve decided to make some changes to Siri as a result.”

Apple intends to resume human review of recordings after it gives users an option to opt in to the storage and human review. “We hope that many people will choose to help Siri get better, knowing that Apple respects their data and has strong privacy controls in place. Those who choose to participate will be able to opt out at any time,” Apple said.

Apple said it keeps track of Siri data with a “random identifier” consisting of a long string of letters and numbers associated with a single device. This system avoids linking Siri data to “your Apple ID or phone number–a process that we believe is unique among the digital assistants in use today,” Apple said.

“When we store Siri data on our servers, we don’t use it to build a marketing profile and we never sell it to anyone,” Apple also said. “We use Siri data only to improve Siri, and

we are constantly developing technologies to make Siri even more private."

No more contractors reviewing audio

Under its new process, Apple won't have contractors review voice recordings anymore."[W]hen customers opt in, only Apple employees will be allowed to listen to audio samples of the Siri interactions," the company said. "Our team will work to delete any recording which is determined to be an inadvertent trigger of Siri."

It isn't clear whether Apple will immediately delete all voice recordings made before the policy change. We asked Apple for an answer to this question and will update this story if we get a response.

The recordings should all be deleted eventually as part of Apple's normal processes, though. An Apple security guide says that "User voice recordings are saved for a six-month period so that the recognition system can utilize them to better understand the user's voice. After six months, another copy is saved, without its identifier, for use by Apple in improving and developing Siri for up to two years."

The current process for deleting your previous Siri recordings from Apple's servers is laborious because it involves disabling Siri on every one of your Apple devices. As Apple notes, "When you turn Siri and Dictation off, Apple will delete the User Data associated with your Siri identifier, and the learning process will start all over again."

Apple will continue storing computer-generated transcripts even for users who don't opt in to storage of recordings. Like voice recordings, Apple said these "are associated with a random identifier, not your Apple ID."

"If you do not want transcriptions of your Siri audio

recordings to be retained, you can disable Siri and Dictation in Settings," Apple said. The company also says that the computer-generated transcripts "are used in machine learning training to improve Siri, determine common usage patterns, and update language and understanding models," and to help Apple "resolve critical problems that affect Siri reliability."

Other major tech companies have also been grappling with voice assistant recording reviews recently. Microsoft declined to pause reviews and instead changed its privacy policy after backlash. More akin to Apple's decision, Google also recently paused reviews of voice recordings after a contractor leaked 1,000 voice recordings to a media outlet.

Google says that the storing of voice and audio activity is set to "off" by default when people create Google accounts. Google users who have turned this on can disable the saving of voice activity and other types of personal information at Google's activity controls site, where they can also delete past recordings.

https://arstechnica.com/?p=1559927

Macron claims deal to end digital service tax fight with the US

Enlarge / France's President Emmanuel Macron (R) and US President Donald Trump shake hands during a joint press conference in Biarritz, France, on August 26, 2019.

LUDOVIC MARIN/AFP/Getty Images

reader comments

78 with 60 posters participating

Share this story

The United States and France have reached a deal over the taxation of digital services, French President Emmanuel Macron said at a Monday news conference alongside US President Donald Trump. But later in the same conference, Trump pointedly declined to confirm that a deal had been reached—leaving the state of the negotiations unclear.

“We have reached a very good agreement,” Macron said, speaking through a translator.

The two leaders have been feuding since last month, when Trump threatened to slap retaliatory tariffs on French products. “I’ve always said American wine is better than French wine,” Trump tweeted.

Later in Monday’s press conference, a French reporter observed that First Lady Melania Trump had been seen drinking French wine. She asked Trump to confirm that the US and France had a deal on digital service taxes.

“I can confirm that the first lady loved your French wine,” Trump said, pointedly ducking the question about the digital services tax.

On paper, France’s 3% tax on digital services revenue isn’t focused on any specific country. But the tax targets large companies—those that have at least €750 million ($830 million) in global revenue—and big global tech companies are disproportionately based in the United States. So, the US argued last month, the tax “will amount to *de facto* discrimination against US companies.”

This weekend’s G7 summit provided an opportunity for the two countries to patch up their relationship. According to Reuters, French foreign minister Bruno Le Maire spent the weekend huddling with US officials to try to work out a deal to defuse tensions.

Now Macron says that a deal has been reached.

“Some digital players pay very little tax,” Macron tweeted on Monday. “This is an injustice that destroys jobs.” He added that he and President Trump “have just agreed to work together on an agreement at the OECD level to modernize international tax rules.”

Macron is hoping that the Organization for Economic Cooperation and Development will establish an international framework for taxing digital services by next year. If that happens, Macron says, France will refund tax money it has collected under its current 3% tax.

But while Macron is clearly enthusiastic about this plan, Trump doesn’t seem to match his enthusiasm. The US president gave no indication during Monday’s press conference that he shares France’s goal to develop a global system for taxing digital services.

https://arstechnica.com/?p=1558231

Australia plans to block sites hosting extremist

content during attacks

Enlarge / Australian Prime Minister Scott Morrison (left) with British Prime Minister Boris Johnson during this weekend’s G7 summit in France. The Australian government announced the new censorship proposal in conjunction with the gathering.

reader comments

206 with 65 posters participating

Share this story

Australia is working on capabilities to quickly block websites hosting terrorist material during crisis events, officials announced on Sunday, according to Reuters.

The governments of Australia and New Zealand have been particularly concerned about the risk of extremist content online after a man killed more than 50 Muslims in a mosque shooting—and streamed a portion of his grisly crime on Facebook.

Now, Australian officials are planning to develop new infrastructure for quickly responding to events like this in the future. According to Reuters, the Australian government plans to “establish a framework to block domains hosting such material. Australia’s eSafety Commissioner would determine on a case-by-case basis what should be censored and was working with industry on arrangements to quickly block access during an attack.”

“A 24/7 Crisis Coordination Centre would be established to monitor the online world for extreme violence or terrorist material,” Reuters reports.

It’s not clear if the Australian government will be able to accomplish all of this using existing legal powers or if new legislation will be needed to make the scheme fully effective.

The Australian initiative underscores that America’s absolutist free-speech tradition is not widely shared among other developed nations. Officials in the EU are pushing for a one-hour takedown rule for online terrorist content. The United Kingdom recently mandated that pornographic sites age-verify their users.

Laws like this would likely be blocked by US courts under the First Amendment. But, of course, other countries don’t have the First Amendment. Most wealthy democracies have some kind of constitutional protections for free speech, but they’re much narrower, leaving substantial room for governments to censor objectionable content.

https://arstechnica.com/?p=1558175

Google defends tracking cookies—some experts aren’t buying it

reader comments

261 with 151 posters participating

Share this story

Google’s Chrome team is feeling pressure from competitors over ad tracking. Apple has long offered industry-leading protection against tracking cookies, while Mozilla recently announced that Firefox will begin blocking tracking cookies by default. Microsoft has been experimenting with tracking protection features in Edge, too.

But Google has a problem: it makes most of its money selling ads. Adopting the same aggressive cookie blocking techniques as its rivals could prevent Google’s customers from targeting ads—potentially hurting Google’s bottom line.

So in a blog post last week, Google outlined an alternative privacy vision—one that restricts some forms of user tracking without blocking the use of tracking cookies any time soon.

“Blocking cookies without another way to deliver relevant ads significantly reduces publishers’ primary means of funding, which jeopardizes the future of the vibrant Web,” Google’s Justin Schuh writes. (Those publishers, of course, include Ars publisher Conde Nast. We use cookies to serve targeted ads because they generate more revenue to support our journalism.)

Google also warns that completely blocking tracking cookies will cause ad networks to resort to browser fingerprinting as an alternative means of tracking users. Under this technique, a site harvests many small pieces of data about a user’s browser—browser version, fonts installed, extensions active, screen size, and so forth—to generate a “fingerprint” that uniquely identifies a particular device.

So Google’s proposal is to declare war on browser fingerprinting while only gradually restricting the use of cookies for ad targeting.

A privacy sandbox?

To prevent fingerprinting, Google says it's working on a new approach called a "privacy budget." Under this approach, the browser would impose a hard cap on the amount of information any site could request from the browser that might reveal a user's identity. If a site exceeded the cap, the browser would either throw an error or it would return deliberately inaccurate or generic information.

But this is only a proposal, not a shipping feature. And it has some obvious challenges. Some API calls might return so much information that they could identify the user all on their own. If a site made one of these calls, the browser would need to warn the user and get explicit approval—which could be annoying for users. And there's a risk that a too-strict privacy budget could break some existing sites even if they're not engaging in user fingerprinting.

The privacy budget is one component of a larger framework Google calls a "privacy sandbox." The goal is to enable advertisers to serve more relevant ads without allowing them to track individual users:

> *We're exploring how to deliver ads to large groups of similar people without letting individually identifying data ever leave your browser – building on the Differential Privacy techniques we've been using in Chrome for nearly 5 years to collect anonymous telemetry information. New technologies like Federated Learning show that it's possible for your browser to avoid revealing that you are a member of a group that likes Beyoncé and sweater vests until it can be sure that group contains thousands of other people.*

Some pro-privacy experts remain skeptical

Google's post was blasted by a pair of Princeton computer scientists who have long advocated for stricter browser privacy protections. They point out that Apple and Mozilla are also working to restrict browser fingerprinting. They argue that it's a non-sequitur to say that the risk of fingerprinting is a reason not to adopt robust restrictions on cookie-based tracking.

The researchers disputed Google's claim that nuking tracking cookies would undermine the economic foundation of the online advertising industry. They point out that after the EU adopted the General Data Protection Regulation, the New York Times discontinued its use of tracking cookies in Europe. The Grey Lady shifted to using contextual and geographic ad targeting—and its ad revenue hasn't suffered as a result.

They also argued that Google is now endorsing ideas that the company dismissed as impractical earlier this decade.

"Privacy preserving ad targeting has been an active research area for over a decade," the pair wrote. They argue that Google was dismissive of alternatives to cookie-based ad tracking during the Do Not Track debate earlier in the decade. "We are glad that Google is now taking this direction more seriously, but a few belated think pieces aren't much progress."

Browser privacy has emerged as an important differentiator for Google's rivals in the browser market. Apple in particular has been running ads in recent months touting the privacy protections offered by the iPhone. These attacks put Google in a difficult position, because Google can't match its rivals' privacy protections without potentially hurting its own lucrative ad business.

https://arstechnica.com/?p=1558091

Book publishers sue Audible to stop new speech-to-text feature

reader comments

12 with 11 posters participating

Share this story

Seven of the nation’s top book publishers sued Amazon subsidiary Audible on Friday, asking federal courts to block the company from releasing a new feature called Audible Captions that’s due out next month. The technology does exactly what it sounds like: display text captions on the screen of your phone or tablet as the corresponding words are read in the audio file.

The publishers argue that this is straight-up copyright infringement. In their view, the law gives them the right to control the distribution of their books in different formats. Audio is a different format from text, they reason, so Audible needs a separate license.

This would be a slam-dunk argument if Audible were generating PDFs of entire books and distributing them to customers alongside the audio files. But what Audible is actually doing

is subtly different—in a way that could provide the company with firm legal ground to stand on.

The caption feature “is not and was never intended to be a book,” Audible explained in an online statement following the lawsuit. “Listeners cannot read at their own pace or flip through pages as they could with a print book or eBook.” Instead, the purpose is to allow “listeners to follow along with a few lines of machine-generated text as they listen to the audio performance.”

“We disagree with the claims that this violates any rights and look forward to working with publishers and members of the professional creative community to help them better understand the educational and accessibility benefits of this innovation,” Audible added.

Audible’s case likely rests on two landmark video recording precedents

Audible hasn’t fully explained how Audible Captions work, but the company has said enough to provide a pretty clear picture. In a video about the service released earlier this month, an Audible executive explained that the technology was “built on publicly available technology through AWS Transcribe.” That’s Amazon’s cloud-based service for automatic text transcription.

So it seems that the Audible app is generating text captions in realtime as the user plays an audio file. The app sends snippets of audio files to an Amazon server and gets back corresponding sections of text, which it then displays on the screen one word at a time. (It’s possible that AWS Transcribe has an offline mode that allows the transcription to happen on-device, but I haven’t found any documentation about this. I’ve asked Audible about this and will update if they respond.)

Audible is likely doing this because it strengthens the company’s argument that it can do this without a license from

publishers.

To see why, it's helpful to review two of the most important copyright decisions of the modern era. The first was the 1984 decision of *Sony v. Universal* that declared the VCR legal. Hollywood argued that the "record" button on a VCR was an invitation for customers to infringe their copyrights. But the Supreme Court disagreed, arguing that copyright's fair use doctrine allowed "time shifting"—recording a show now to play it later.

The courts built on this decision with a 2008 ruling known as *Cartoon Network v. Cablevision*. In that case, a bunch of media companies sued the cable company Cablevision because it was offering customers a "remote DVR." Like a conventional DVR (or a VCR before that), Cablevision's technology allowed customers to record and play back television shows at their convenience. But unlike a conventional DVR, the remote DVR was located in a Cablevision data center, not in the customer's home.

Television content owners argued that Cablevision was infringing their copyrights by making unauthorized copies of their show on a massive scale. Cablevision disagreed, arguing that the copies were being made by *customers*, not by Cablevision. The physical DVR might be owned and maintained by Cablevision, but the customer was deciding which shows to record. And the customer was entitled to do that under the earlier *Sony* ruling. An appeals court ultimately accepted this argument.

The Cablevision ruling provided a legal foundation for cloud-based "storage locker" services that allowed customers to upload, save, and stream (but not share) their music and video collections.

Expect *Cablevision* to be central to Audible's defense

That brings us back to Audible's new transcription technology. Audible doesn't have the legal right to sell text versions of audiobooks to customers' without publishers' permission. But we can expect Audible to argue that it *does* have a right to sell software tools that allow *customers* to do speech-to-text conversion.

Audible's case will likely be strengthened by the fact that its app never creates or saves a permanent, full transcript of an audiobook. Instead, the software only displays a few words on the screen at a time.

If Audible is sending audio files to Amazon's servers for transcription, publishers are likely to argue this means Amazon—not users—are creating the transcripts. But this seems closely analogous to the *Cablevision* case: the conversion is being done by Amazon servers, but only when explicitly requested by users. And each translation is only sent back to the user who requested it.

Of course, outcomes in disputes like this aren't easy to predict. Courts could decide that a conversion from audio to text doesn't have the same fair use protections as the time-shifting activities the courts have blessed in the past. The courts could decide that Amazon plays too active a role in the conversion process to portray itself as a passive supplier of technology like the maker of a VCR. Or they might decide that this whole legal argument is a too-clever-by-half ploy to justify reproducing publishers' copyrighted books. Ultimately, we'll have to wait and see what the courts say.

Interestingly, Amazon faced the same battle in reverse a decade ago when it proposed adding a text-to-speech feature to Kindle. Like Audible today, Amazon was planning to offer the

text-to-speech feature for all e-books without getting approval from book publishers. But after a backlash from publishers, Amazon caved and gave publishers the right to disable the feature. Amazon now says that “most Kindle content and personal documents are eligible for Text-to-Speech.”

https://arstechnica.com/?p=1557181

Hundreds of “banned” goods still for sale on Amazon, report finds

Enlarge / A drone with an Amazon package floats in front of the Amazon logistics center in Leipzig, Germany, 28 October 2014. Amazon did not comment on whether drones will fuel this default one-day speed boost for paying Amazon Prime subscribers’ deliveries.

reader comments

81 with 60 posters participating

Share this story

Amazon is by far the biggest US online retailer. In the past 20 years it has leapt past its origins as a website you could order books from to become, among other things, the everything store—one-stop shopping for all physical and digital goods from A to Z.

The company’s explosive growth is due in part to its sprawling third-party merchant marketplace. Many marketplace merchants are indeed above-board retailers, manufacturers, and resellers. But thousands more sell not only counterfeit items, but also mislabeled, unsafe, recalled, or even banned items that can put consumers—especially children—in serious danger.

The Wall Street Journal identified more than 4,100 such products for sale on Amazon.com during the course of a months-long investigation, and at least 2,000 are toys or medications that fail to include warnings about risks to children.

Among the Journal’s findings: 116 products falsely listed as FDA-approved, including toys, which the agency does not regulate; 80 listings for infant “sleeping wedges” the FDA says can cause suffocation and that Amazon had previously banned; 1,412 electronics listings falsely claiming to be UL-certified; 2,324 toys that failed to include federally mandated choking hazard warnings; and more.

The WSJ commissioned tests of 10 specific children’s products it bought on Amazon, many carrying the enigmatic “Amazon’s Choice” badge. Of those, four failed tests based on federal safety standards, including one that contained excessively high levels of lead.

Balloons also proved to be a sticking point: the WSJ found 4,500 balloon listings that did not include required choking hazard warnings, and the paper notified Amazon about them. Weeks later, WSJ found another 2,200 balloon listings that did not include required warnings. Including all the balloon listings, the WSJ identified 10,870 problematic listings to Amazon, of which 83% were eventually removed or altered.

One mislabeled product the WSJ included in its report proved deadly to its owner. A 23-year-old man in Missouri purchased a motorcycle helmet from Amazon that was at the time listed as certified as meeting US Department of Transportation safety

standards. Later that year, however, he was killed in a crash while riding. A federal investigation later found that the helmet did not meet DOT standards and was recalled. The WSJ, however, found the product still for sale, with an active listing promising compliance, until the WSJ contacted Amazon to inquire about it.

Whack-a-Mole

The WSJ's investigation found 157 products for sale that Amazon has *already* banned from sale on its site. The motorcycle helmet was one of more than 2,300 product listings altered or pulled after the WSJ drew them to Amazon's attention. Yet, within two weeks, the WSJ found that at least 130 of these problem items reappeared, "some sold by the same vendors previously identified by the Journal under different listings."

In short, the Journal writes, Amazon "has increasingly evolved like a flea market," exercising little to no oversight over items sold by third-party merchants unless a specific complaint or media report draws an item to the company's attention. The marketplace setup that causes Amazon to land in hot water seemingly annually for selling some kind of pro-rape, pro-slavery, or pro-Nazi apparel also leads to endemic listings for recalled or harmful goods.

About 60% of Amazon's physical retail sales come from the third-party marketplace, the company has said. A recent quarterly report showed the marketplace generates $11.14 billion in sales for the ecommerce behemoth in just three months. But even consumers wary of shopping from third-party merchants can still easily find themselves purchasing from one. Many of the items in the WSJ's investigation were fulfilled by Amazon: eligible for Prime shipping, from Amazon warehouses, and in Amazon boxes. Two different shoppers told the WSJ they had assumed harmful or mislabeled products they

bought from the site were reviewed and approved in some way by Amazon, as they would be in a big-box store such as Target or Walmart, until contacted by the WSJ.

Counterfeit products, which are often less safe than their "real deal" counterparts, are also a persistent plague to Amazon and its customers. The WSJ did not test counterfeit goods, which show up in searches for every kind of product, from luxury goods to cheap USB cables.

When reached by Ars for comment, an Amazon representative directed us to a company blog post that says, in part, "We provide a number of ways for regulatory agencies, industry organizations, brands, customers, and our customer service teams to report safety issues. When we receive these reports, we move quickly to protect customers, remove unsafe products from our store, and investigate."

Limited recourse

Swimming against the tide of dodgy listings and questionable goods is challenging for even the best-educated consumer, and prevailing in court if you do end up with a damaging product is, at best, a hit-or-miss exercise.

The family of the young man who died in the motorcycle crash sued Amazon, as well as the driver of the vehicle with which he had collided and the third-party merchant that sold the helmet. The merchant was ordered to pay $1.9 million in restitution; Amazon settled for $5,000 without admitting wrongdoing. A company attorney told the WSJ, "Basically, a third party was using Amazon as a bulletin board to advertise the product and sell."

Your ability to sue Amazon if you are injured by a third party's product largely depends on where you live. A federal appeals court in Philadelphia ruled in July that Amazon could be held liable under state law for a defective product that

blinded a Pennsylvania woman.

The Third Circuit court, however, was the first to do so. The Fourth and Sixth Circuit Courts of Appeals ruled in May and June respectively that Amazon is merely a platform, not a "seller," when it comes to state consumer protection law.

https://arstechnica.com/?p=1556937

US phone carriers make empty, unenforceable promises to fight robocalls

reader comments

43 with 38 posters participating

Share this story

Phone companies and attorneys general from all 50 US states are touting a new agreement to fight robocalls, but it won't actually do much to help consumers.

The top wireless carriers and home phone providers promised attorneys general from every state and the District of Columbia that they would offer free robocall blocking and take other steps to fight robocalls. But the agreement imposes no legally binding requirements on phone providers. "Failure to adhere to these principles is not in itself a basis for

liability," a disclaimer on the agreement notes.

Even if breaking the agreement *was* a basis for liability, there would be no deadline to comply. "Adherence to these principles may take time for the voice service providers to plan for and implement," the disclaimer also said, while providing no specific timeline for the carriers to fulfill their promises.

Given that disclaimer, you'd think carriers actually agreed to make some major changes. But the agreement's top promises are things the phone companies are already doing or in the process of rolling out.

Carriers can still charge extra

The promise of "free" robocall blocking can be met by offering a basic level of call blocking for free, even if carriers also charge a monthly fee for more robust call-blocking services. We contacted the top four wireless carriers yesterday—AT&T, Verizon, T-Mobile, and Sprint—and none of them promised any additional free services beyond what they already offer. The agreement between carriers and states doesn't say what level of call blocking must be offered for free, so it's up to the carriers to decide.

The agreement also includes a promise to deploy the new SHAKEN/STIR technology that uses digital certificates to verify that Caller ID numbers aren't being spoofed. But carriers already promised the Federal Communications Commission that they would do this and started the process of deploying SHAKEN/STIR months ago.

The agreement separately includes a promise to offer "network-level call blocking at no charge," which some carriers have been doing for years. The carriers also promised to "analyze and monitor network traffic" and "investigate suspicious calls and calling patterns," which sound like things the carriers

would do in order to fulfill the other promises in the agreement.

In short, it’s a legally unenforceable agreement consisting largely of promises the companies have already met. It’s more of an exercise in public relations than a consumer-protection initiative.

“No substitute for enforceable rights”

Despite its shortcomings, the agreement drew loads of positive media coverage from major news organizations yesterday. Everyone hates robocalls, so carriers and regulators have faced pressure from consumers and politicians from both major parties to do something—anything—to put a stop to them. In this environment, even a legally unenforceable agreement to do things the carriers were already doing was good enough to achieve the carriers’ and regulators’ unstated goal of getting a public relations victory.

“Voluntary agreements are no substitute for enforceable rights,” Harold Feld, senior VP of consumer advocacy group Public Knowledge, told Ars. “This is the sort of smokescreen designed to stop Congress from passing more effective legislation, like the Stopping Bad Robocalls Act passed by the House [and pending in the Senate]. With nothing to force carriers to meet deadlines, nothing to force carriers to upgrade once robocallers figure out how to outwit the new technological measures, and nothing to stop carriers from walking away if it gets too hard or too expensive, I wouldn’t declare ‘Mission Accomplished.'”

Feld does think one aspect of the agreement could lead to real progress, though. Carriers promised to communicate with state attorneys general about robocall trends and cooperate in investigations by making prompt and thorough responses to

traceback requests from law enforcement. Each carrier would do this “by identifying the upstream provider from which the suspected illegal robocall entered its network or by identifying its own customer if the call originated in its network,” the agreement says.

“The real value is the engagement of all state AGs, particularly on enforcement against robocallers,” Feld said. “The emphasis on cooperating with law enforcement on tracebacks is a clear sign that state AGs are eager to go after robocallers, and this agreement makes it more likely that will happen.”

Limits to free blocking

The coalition of state attorneys general was led by North Carolina AG Josh Stein, who issued a press release touting his “leadership in bringing together a bipartisan, public/private coalition of 51 attorneys general and 12 phone companies.”

We asked Stein’s office if there would be any penalties or legal recourse if carriers don’t follow the anti-robocall principles. Stein’s spokesperson replied to us but didn’t answer that question.

Stein’s office partially answered our question about whether the carriers can meet the “free” robocall blocking requirement even if they continue charging for premium blocking services.

“As part of the agreement, each of the companies have agreed to provide base-level blocking for free,” Stein’s spokesperson told us. But Stein’s spokesperson hasn’t answered a follow-up question in which we asked for a definition of “base-level blocking.” The actual agreement between carriers and states never defines or even mentions “base-level blocking.”

The top wireless carriers seem to think they’ve already met the requirement.

“T-Mobile has been the industry leader in fighting against robocalls and fully supports and already meets the State Attorneys General Anti-Robocall Principles for voice service providers,” a T-Mobile spokesperson told us.

T-Mobile offers free robocall blocking but charges $4 a month for “Name ID.” The paid service lets you block calls from specific callers or from entire categories, “including telemarketers, numbers that have been tagged as nuisance calls, political calls, survey calls and more.”

T-Mobile told Ars that it is still charging extra for Name ID.

AT&T’s robocall blocking service lets some robocalls through unless you pay extra. The free version of AT&T’s Call Protect service automatically blocks calls that AT&T categorizes as “fraud.” AT&T Call Protect also provides warnings for suspected “spam” calls but lets them ring your phone unless you buy the premium version for $4 a month.

We asked AT&T if it will stop charging for spam call-blocking, but the company didn’t provide an answer. AT&T did say that the agreement with states “align[s] with much of what AT&T is already doing to protect our customers against illegal and unwanted calls.”

When contacted by Ars, Verizon noted that it offers “both a free and premium product.” Verizon said it will have “additional, related news to announce next week,” but didn’t say if that will include any more free services.

Sprint doesn’t seem to offer any free call blocking, and it charges $3 a month for its “Premium Caller ID” service. Despite that, Sprint made no promises to offer free call blocking when we contacted the company yesterday.

“I’m not in a position to share specifics on this today, but we are committed to continuing to provide our customers with tools that will help them to combat robocalls,” a Sprint

spokesperson told Ars. If Sprint's pending merger with T-Mobile goes through, the company's policies would likely change to meet T-Mobile's.

Better than "thoughts and prayers"

Besides the top four wireless carriers, the agreement with states also includes Bandwidth Inc., CenturyLink, Charter, Comcast, Consolidated, Frontier, US Cellular, Verizon, and Windstream.

Comcast, Charter, and Frontier have all partnered with Nomorobo to offer robocall blocking on home phones at no extra charge. Verizon and AT&T have done the same for their home phone services. Nomorobo is free for VoIP landlines anyway, so it'd be easy for other carriers to do this as well.

The new promise of free robocall blocking tools is only "for smartphone mobile and VoIP residential customers," according to the agreement with states. Customers with older landlines aren't being promised free blocking tools.

Overall, Feld thinks that the states' agreement with carriers is "a step up" from the FCC's recent order that gave carriers additional authority to block robocalls but didn't require carriers to offer anything for free. Feld called the FCC vote a "thoughts and prayers" order.

As for the promises to states, "all of these are things that the carriers have already largely committed to do, and these principles don't change the practical problem of forcing carriers to keep their promises," Feld said. "At best, it creates more political pressure. That's not nothing, but it's not much."

https://arstechnica.com/?p=1557089

Tesla working to resolve dispute with Walmart over solar panel fires

Enlarge / Fire damage on the roof of a Walmart store in Indio, California.

reader comments

76 with 51 posters participating, including story author

Share this story

Walmart and Tesla are actively negotiating to resolve the lawsuit Walmart filed against Tesla earlier this week over defective solar panels, the two companies said in a joint statement sent out on Thursday evening.

“Walmart and Tesla look forward to addressing all issues and re-energizing Tesla solar installations at Walmart stores, once all parties are certain that all concerns have been addressed,” the statement said. The companies say they’re both committed to a “sustainable energy future” as well as safety and efficiency.

Spokespeople for Tesla and Walmart declined to provide any further details about the state of the negotiations, but it’s not hard to guess what happened. The optics of Walmart suing Tesla over multiple fires on its store roofs were not good for Tesla. Tesla wants the public—and potential customers—to know that it’s now working to address Walmart’s concerns.

Walmart filed its Tuesday lawsuit after suffering a series of seven fires since 2012 on stores with solar panels on them. The solar panels were installed by Tesla's solar panel subsidiary, which was a separate company called SolarCity until that business was acquired by Tesla in 2016. Walmart claimed on Tuesday that its investigation revealed "widespread, systemic negligence" in Tesla's installation and maintenance of solar panels. Walmart also charged that Tesla had dragged its feet in addressing the problems Walmart identified or compensating Walmart for the damage.

Walmart had asked the court to order the removal of all Tesla solar panels from its stores, but now the companies say they're working toward a deal to keep the panels in place, repair them, and begin using them again.

https://arstechnica.com/?p=1556929

The FCC has no idea how many people don't have broadband access

reader comments

76 with 57 posters participating

Share this story

A new broadband mapping system is starting to show just how

inaccurate the Federal Communications Commission's connectivity data is.

In Missouri and Virginia, up to 38% of rural homes and businesses that the FCC counts as having broadband access actually do not, the new research found. That's more than 445,000 unconnected homes and businesses that the FCC would call "served" with its current system.

Given that the new research covered just two states with a combined population of 14.6 million (or 4.5% of the 327.2 million people nationwide), it's likely that millions of homes nationwide have been wrongly counted as served by broadband. A full accounting of how the current data exaggerates access could further undercut FCC Chairman Ajit Pai's claims that repealing net neutrality rules and other consumer protection measures have dramatically expanded broadband access. His claims were already unconvincing for other reasons.

The new research was conducted by CostQuest Associates, a consulting firm working for USTelecom, an industry lobby group that represents AT&T, Verizon, CenturyLink, Frontier, and other fiber and DSL broadband providers. USTelecom submitted a summary of the findings to the FCC on Tuesday. The two-state pilot was intended to determine the feasibility of creating a more accurate broadband map for the whole US.

Why the FCC's current data is wrong

The key problem with today's maps is that the FCC's Form 477 data-collection program that requires ISPs to report census-block coverage lets an ISP count an entire census block as served even if it can serve just one home in the block. It has always been known that this approach could undercount the number of unserved homes, but it was never clear exactly how far off the numbers were.

The FCC's latest numbers suggest that 21.3 million Americans

lack access to fixed broadband with speeds of at least 25Mbps down and 3Mbps up. But those numbers are based on the faulty census block data.

“In addition to other important metrics, our pilot shows as many as 38 percent of additional rural locations in Virginia and Missouri are unserved by participating providers in census blocks that would have been reported as ‘served’ in today’s FCC Form 477 reporting approach,” USTelecom and other trade groups representing telecoms and fixed wireless providers told the FCC. “These locations are homes and small businesses hidden from service providers and policymakers simply because of a lack of knowledge fueled by gaps in data—gaps that we can now fill.”

USTelecom argued as recently as October 2017 that the FCC “should not seek to collect broadband deployment data that is more granular than at the census block level, because such a change would be unduly burdensome to providers.” But industry groups and the FCC itself changed their stance after bipartisan complaints about the inaccuracy of US broadband data.

Three weeks ago, the FCC voted to require ISPs to give the FCC geospatial maps of where they provide service instead of merely reporting which census blocks they could offer service in.

The FCC’s new policy should go a long way toward fixing the data problem. But the more accurate data may make it harder for Pai to claim that his deregulatory moves are closing America’s broadband gaps. Pai didn’t create the process for collecting the Form 477 data, but he has used it to claim that his policies created more broadband, even though the data showed broadband deployment was progressing at about the same rate it did during the Obama administration.

More accurate data will also help the FCC determine which

areas should get the most federal funding to expand broadband access. The FCC's Connect America Fund has given billions to ISPs to expand Internet service since its creation in 2011, and Pai plans to continue that with a 10-year, $20.4 billion fund that would pay ISPs to bring broadband to unserved rural areas.

Data was wrong in 48% of rural census blocks

The CostQuest/USTelecom two-state pilot created a map (or "fabric") of virtually all homes and businesses that could be served by broadband if Internet providers built out to them. The pilot also asked ISPs to submit coverage data, and the ISP-submitted data was compared to the statewide maps to determine how many buildings lacked access.

"Creating the fabric revealed that in just two states over 450,000 homes and businesses exist that are counted as 'served' under current 477 reporting that are not receiving service from participating providers," CostQuest wrote. "While not every broadband provider chose to participate in this pilot—so the actual number of unserved may be lower—that still leaves the potential for substantial misrepresentations about service availability."

The pilot also showed that current broadband-availability data is wrong in 48% of rural census blocks and is "in many cases significantly different."

A nationwide version of the CostQuest map and data set could be completed in 12 to 15 months "for between $8.5-$11 million in upfront costs and $3-4 million in annual updates," the summary of key findings said. To create the fabric, the pilot used data sources including tax assessor and parcel attribute data, georeferenced building footprints, and road data, along with statistical analysis and crowdsourced in-person reviews

of parcels to check accuracy.

CostQuest said its data is far more accurate than what standard geocoding tools provide. The company found that 61% of locations in rural areas were geocoded at an incorrect location and that 25% of the locations were off by more than 100 meters. Additionally, 23% of locations had been geocoded to the wrong census block.

In Missouri, CostQuest found that 9% of non-rural locations were unserved and that 36% of rural locations were unserved. In Virginia, 12% of non-rural locations were unserved and 39% of rural locations were unserved.

It’s not clear whether the USTelecom/CostQuest approach will be completed nationwide. The FCC is requiring ISPs to submit geospatial coverage maps, and the FCC plans to create a crowdsourcing system to collect public input on the accuracy of those ISP-submitted maps. But that wouldn’t require the use of CostQuest’s system, and deadlines for ISPs to submit maps have not yet been announced.

https://arstechnica.com/?p=1556249

While one Texas county shook off ransomware, small cities took full punch

Enlarge / They did.
Hemera Technologies/Getty Images

reader comments

72 with 53 posters participating

Share this story

Few details have emerged about the coordinated ransomware attack that struck 22 local governments in Texas last week. But five local governments affected by the attack have been identified.

On August 20, the Texas Department of Information Resources revised its initial report that 23 “entities” had been affected by the ransomware attack, reducing that count by 1. And a Texas DIR spokesperson said in a statement that about a quarter of the local governments affected have been able to at least partially restore normal operations.

That includes Lubbock County, which apparently escaped major disruptions. Lubbock County judge Curtis Parrish told Magic 106.5 Radio that the county’s IT department “was right on top of it… they were able to get that virus isolated, contained and dealt with in a very quick manner so it did not affect any other computers or computer systems here in Lubbock County.”

Other, smaller local governments were not as lucky. Borger, Texas, which has a population of about 13,000, was one of the first to go public. On August 19, Borger city officials announced on Facebook that the ransomware attack had “impacted normal City business and financial operations and services.” A continuity of operations plan had been put into effect to “provide basic and emergency services (Police, Fire, 9-1-1, Animal Control, Water, Wastewater and Solid Waste Collection),” but there was no estimate of how long it would take for full services to be restored.

A spokesperson for the City of Kaufman, Texas, announced on Facebook that Kaufman had also been affected by ransomware. “The City of Kaufman Computer and Technology Services has been

severely affected by an outside source," the spokesperson said in an August 19 post. "At this time, all of our computer and phone systems are down and our ability to access data, process payments, etc. is greatly limited." Kaufman's police and fire departments remain operational, according to the statement.

Today, the LA Times reported that Keene, Texas—a small incorporated city of 6,500 people 40 miles south of Fort Worth—was also hit by the ransomware. While Keene's police and utilities were not affected, the city's payment system for water bills was taken down.

But the hardest hit known so far is the City of Wilmer, a community of about 3,600 in Dallas County. According to a report from CBS' Dallas/Fort Worth affiliate, systems at Wilmer's police department, water department, and public library were affected. City workers reported that when they turned on computers, they were greeted with a blue screen carrying the message, "all your files are encrypted."

While the ransomware used in the attack has not been identified, the details reported thus far do not appear to match with the ransomware used in attacks on local governments in June and July. Those incidents involved the Ryuk malware, which leaves a ransom note in the form of a text file on the victims' computers. However, the new attack does appear to be targeted.

Small cities are a favored target of ransomware operators because they're more likely to pay, according to Chris Hinkley, a senior security researcher with the Threat Resistance Unit (TRU) team at the cloud security provider Armor. "The enormous impact of city's operations, like 911, courts, police and fire, and even non-emergency services creates a huge sense of urgency and anxiety," Hinkley explained. "With the high level of urgency and potentially mission-critical value of the data being held hostage, the threat actors are more likely to get paid, and at a higher

amount, than if they attacked another target.”

Hinkley said that it was not surprising that Texas had been hit hard by ransomware—aside from the 22 attacks in last week’s coordinated ransomware campaign, “there are seven other Texas municipalities which have been victims of ransomware in 2019,” he said. Texas has 1,216 incorporated cities, of which only 35 have more than 100,000 residents. That, Hinkley said, makes for a “very large attack surface.” And smaller cities are less likely to be able to defend themselves adequately against ransomware.

https://arstechnica.com/?p=1555903