Dis-CONTI-nued: The End of Conti?

Description: This week we'll start by following-up on Microsoft's Patch Tuesday Active Directory domain controller mess. We're going to look at several instances of the Clearview AI facial recognition system making news, and at the systems which fell during last week's Vancouver Pwn2Own competition. We cover some welcome news from the U.S. Department of Justice and some disturbing news about a relatively simple and obvious hack against popular Bluetooth-link smart locks. We have some closing-the-loop feedback from our listeners, including a look at what's going on with the Voyager 1 space probe, and another interesting look into the looming impact of quantum crypto. Then we finish by sharing an in-depth examination of the surprisingly deliberately orchestrated shutdown of the Conti ransomware operation.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-872.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-872-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Well, well. What do we have to talk about? We're going to say hello to Clearview AI's facial recognition. They're using it in Ukraine, but is it okay? And then it's the results of Vancouver's Pwn2Own competition. We'll talk about Steve's bizarre theory about what's going wrong with Voyager 1, and then a look at the Conti ransomware operation. It looks like they've gone out of business. But have they? It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 872, recorded Tuesday, May 24th, 2022: Dis-CONTI-nued: The End of Conti?

It's time for Security Now!, the show where we cover your privacy, security, and safety online with the man, the myth, the legend, Steve Gibson. Hello, Steve.

Steve Gibson: Apparently we are a well-caffeinated team this morning.

Leo: Ah, yeah, a little bouncy right now, yes, I am.

Steve: Or this early afternoon, rather, whatever it is. So you agreed with me that the title for the show is a bad pun.

Leo: It's godawful, yes.

Steve: Dis-CONTI-nued.

Leo: Dis-CONTI-nued, yes.

Steve: Yes. The end of CONTI, question mark.

Leo: This has become I think kind of a new MO for these bad guys.

Steve: Yes. Thank you for the perfect foil, segue, Leo, because that's actually what hooked me on the story and what I thought would make it so interesting for our listeners is that there is evidence that this whole Costa Rican debacle was a setup designed to obscure the reformation of CONTI. And what's really...

Leo: Oh, interesting.

Steve: Yes. And remember how it didn't, like something seemed off about it.

Leo: Yeah.

Steve: It was like, what? You know, it just...

Leo: Yeah, exactly.

Steve: It was weird when we talked about it last week, actually for the last couple weeks in various aspects. And it's looking like it was actually, well, okay. I don't want to give it all away here at the beginning because we have lots to talk about. We're going to follow up on Microsoft's Patch Tuesday Active Directory domain controller mess. We're going to look at several instances of the Clearview AI facial recognition system making news, and at the systems which fell during last week's Vancouver Pwn2Own competition.

We cover some welcome news from the U.S. Department of Justice. They're rethinking what it means to be an ethical hacker, which is really good news. And some disturbing news about a relatively simple and obvious hack against popular Bluetooth-linked smart locks. We've got some closing-the-loop feedback. Actually only three pieces, but one of them ends up being, well, actually two of them. Two of the three end up really big, expanding into something because we're going to take a look at what's going on with the Voyager 1 space probe which has just started to act a little wonky.

Leo: You mean Veeger?

Steve: Oh, please don't say "Veeger." That was, god, was that an awful movie. Thank goodness somehow they were forgiven for producing that first atrocity. And then we got Khan, so...

Leo: Khan.

Steve: Yes, Veeger. I have a theory about what might have happened which you're going to get a kick out of.

Leo: Oh, interesting.

Steve: Yeah, well, anyway, I'll save that, too. And another interesting look into the looming impact of quantum crypto. We then are going to finish by sharing an in-depth examination of the surprisingly deliberately orchestrated shutdown of the CONTI ransomware operation. It turns out it was far more well-planned than was known. And essentially the guys that have been watching this have spoiled what CONTI was trying to do.

Leo: Oh, good.

Steve: And it's, yeah, good. And we have a pretty funny Picture of the Week. So I think another great podcast for our listeners. What a shock.

Leo: Spoil the spoilers. All right. All coming up. Picture of the Week time, Mr. Gibson.

Steve: So, okay. So everything about this picture to me looks authentic. It's showing a standup card, sort of, not an ad, like a...

Leo: It's a table card. They put them on the table or a placard they put outside the event, the conference.

Steve: Yes, thank you, yes, for a Microsoft MVP Global Summit. And again, so this was in a conference in a hotel. You can see this godawful hotel carpet underneath it.

Leo: Isn't it always ugly?

Steve: Oh.

Leo: It's always the worst.

Steve: I think it's so that it hides barf and anything else that, you know...

Leo: He's exactly right. Chewing gum.

Steve: Spilled cocktails.

Leo: Exactly.

Steve: Nothing is going to - no spots will show through this horrible print that this thing has. Anyway, so we have this sign. And it was originally boring signage. It probably said, you know, Security Summit Conference 203B or something. Anyway, some person with a sense of humor changed around the normal heading and everything, which looks 100% authentic. They changed the sign of the Microsoft MVP Global Summit announcement to say, or to caution: "Unattended laptops will be upgraded to Windows 11."

Leo: Ooh.

Steve: Which, you know...

Leo: Mean.

Steve: ...is not what anyone wants.

Leo: No.

Steve: So do not leave your laptop unattended, children, or you may find out that your menu can no longer be docked to the left-hand side of the screen. You must have it in the center because we're sure that's better for you. Okay. So speaking of Microsoft, last week we noted that the previous week's Patch Tuesday had been overly eventful, with admins reporting that the month's patch rollup had adversely affected the operation of some of their enterprise's critical authentication infrastructure. The malfunctioning patch was intended to close an actively abused security vulnerability. And while there were some pre-patch workarounds, most admins were choosing to simply roll back their systems by removing the entire May Day patch bundle. So that's where we were last week.

Last week I wondered out loud whether Microsoft would attempt to fix the fix, and hopefully to test it more thoroughly this time, or whether they would prefer to wait until June's Patch Tuesday. One problem with waiting is that June begins next Wednesday, which places June's Patch Tuesday as late in the month as possible, nearly in the middle, on the 14th.

And in the past, Microsoft's patching logic has been, well, calling it "inscrutable" is being kind. Sometimes they wait half a year before patching something. Other times they patch it in a hurry. Fortunately for Active Directory users, the latter was the case here, with Microsoft fixing the trouble in about a week and releasing a mid-cycle emergency patch last Thursday.

However, presumably since it was only domain controller servers that were being adversely affected, that point that Microsoft took great pains to point out throughout all this, it's like, oh, no, it's all fine, don't worry about it, it's only these people, like the admins of the largest enterprises in the world, who have been brought down by this. So as a consequence the emergency update is not being made widely available through the regular Windows Update channel. They're not wanting to like update everybody again because most people aren't affected.

So enterprising users, if you'll pardon the pun, will need to get it from Microsoft's Update catalog by going there directly. But you can also, those users in the enterprises who have Windows Server Update Services or Microsoft's Endpoint Configuration Manager can import the update from the catalog into those services, and then it'll fix their networks. So since CISA was forced to pull the requirement of patching this flaw from their catalog of known exploited vulnerabilities due to May's debacle we can assume that it will be returning to the catalog now that it is possible to actually patch against this, remember, actively exploited flaw in Active Directory.

In any event, it should be obvious to everyone these days that, more than ever, patching has become everything. Patch, patch, patch, patch, patch. You just, I mean, it's better to patch and regret it than it is not to patch at all because you can unpatch.

Leo: That should be on a pillow or something. Good. That's a good motto.

Steve: Well, it's the motto that evolved from this podcast because just look around. Wow. Okay. We've talked about Clearview AI on several occasions. Recall that they're the company which provides what have turned out to be quite controversial facial recognition services to apparently anyone who wants them. The controversy surrounds the perceived privacy invasion created when an army of both visible and invisible cameras might be scanning and cataloging the real-world identities of everyone who crosses through their fields of vision.

Thirty years ago we lacked the ability to automate anything like this at scale because it all comes down to economics. We didn't have inexpensive cameras. We didn't have inexpensive communications bandwidth, or inexpensive and apparently endless amounts of storage or computation. Today we have all of that, essentially for next to nothing. So what could not be done in the recent past is now feasible at a mass scale. Global. During a recent interview, Clearview AI's CEO, a guy by the name of Hoan Ton-That, explained that he intends to have 100 billion images in their database within a year. Currently they have about 20 billion. And 100 billion would represent 13 photos of each person on Earth. Obviously not all equally represented. But probably the more important you are, the more photos you've got in the database.

Okay. So the potential for the abuse of this technology is breathtaking due to its sweeping scope. One of the times we've talked about these guys in the past was when the ACLU, the American Civil Liberties Union, and others brought a lawsuit against Clearview AI in Illinois, thanks to Illinois's very restrictive Biometric Information Privacy Act (BIPA), which became law in 2008. Thanks to BIPA, any entity wishing to collect biometric identifying information and a photograph is that. I knew it was you, Leo, when you sat down, just by looking through the camera.

Leo: Very high intelligence.

Steve: Any entity wishing to collection biometric identifying information for any Illinois citizen must first obtain their explicit consent before doing so, which of course Clearview AI never does. They're scanning online image sources, so consent is never requested nor received. Thus they immediately fell afoul of Illinois's very restrictive BIPA. Earlier this month Clearview AI agreed to a settlement in that lawsuit, which was brought by the ACLU, that would largely prohibit sale of services to private companies which previously weren't prohibited, and people, similarly. And you were just talking recently about how it

turned out that you could get location, like individuals could purchase location information on other individuals.

Leo: Oh, yeah. It's nothing. Data brokers are amazing.

Steve: Yes. And so Clearview AI is clearly a massive cloud-based broker that also provides image matching. So the terms of the suit prohibit the sale of Clearview AI services to private companies and people in the U.S., as well as to law enforcement throughout Illinois, for the next five years. However, Clearview AI's CEO said that the settlement would not change anything on a material level because they could continue to sell to any government clients elsewhere. So they took what they consider to be a little hit, and they'll continue unabated.

At the same time, Clearview AI is increasingly facing pressure from various nations' privacy regulators who are pushing Clearview to remove their data on their citizens from its systems. Both Australia and Canada last year ordered the company to delete information on their residents. And now last week the U.K.'s government announced that they are levying a fine of more than 7.5 million pounds sterling, that's around $9.4 million USD, against Clearview AI, and have ordered it to stop collecting information about U.K. residents and to delete all of the data that it already has from its database.

The U.K.'s Information Commissioner John Edwards said in a press release: "The company not only enables identification of those people, but effectively monitors their behavior and offers it as a commercial service. That is unacceptable. People expect that their personal information will be respected, regardless of where in the world their data is being used. That's why global companies need international enforcement." Edwards added that he would be meeting with European regulators in Brussels later this week to "collaborate to tackle global privacy harms." And I didn't mention, these guys are the biggest now globally. They are the number one facial recognition service provider in the world.

These recent moves by the U.K.'s Information Commissioner's Office follow a joint investigation by the agency and the Office of the Australian Information Commissioner which launched in July of 2020. So that's been going for some time and was completed last November. A provisional notice from the Information Commission's Office to Clearview that month warned the company to stop processing and to delete U.K. resident data, as well as suggesting a substantially larger - at the time they were saying they wanted 17 million pounds sterling. It ended up being 7.5. But still, you know, it'll get your attention.

In reply, again, Clearview's founder and their chief executive, that Hoan Ton-That, said he was "deeply disappointed" that the U.K. data authority "misinterpreted" his company's technology and their intentions. Uh huh. He said: "I would welcome the opportunity to engage in conversation with leaders and lawmakers so that the true value of this technology which has proven so essential to law enforcement can continue to make communities safe," he said.

Okay. So there's one side. There's also been Clearview, although not by name until recently, in the news relative to Russia and Ukraine. There's no question that facial recognition technology has the potential to be quite useful. It is Clearview AI's facial recognition technology that has been allowing the Ukrainian government to identify both their own citizen casualties of war, obviously in their war with Russia, as well as the Russian soldier casualties who've been left behind by their comrades.

Thanks to Clearview AI's knowledge of Russian citizens, Ukraine has been able to mount a potentially powerful political counteroffensive by identifying these Russian war casualties, notifying their families of their demise, and offering to allow them to travel to Ukraine to reclaim their fallen family member. Some 400 Ukrainian investigators are currently using Clearview AI's technology to build much more airtight war crimes cases as a consequence. It really helps to know who you're talking about. And certainly Russia's not providing that information.

The publication The Record recently interviewed Clearview AI's founder and CEO, this Hoan Ton-That, whom I referred to before. I think that the details revealed by that interview are worth sharing. So I've edited it a bit to bring it up to the level of our listeners. So in the interview the question is posed: How did Clearview come to play such a major role in Ukraine?

Hoan replies: When the war started, it was really shocking to me and a lot of members of our team to see especially the video footage of women and children suffering. And it made me think, how can we help? Originally, I would see photos of captured Russian soldiers, and I realized that with that kind of photo quality, our facial recognition technology could be helpful. So I reached out to a lot of people on our advisory board to ask them, do you know anyone in Ukrainian government? And one person, Lee Wolosky, he's on the National Security Council under three presidents, he quickly said yes. We thought that it could be helpful in identifying deceased soldiers and track misinformation.

So how could it be useful to track misinformation is asked. Hoan says: You'd see a lot of things on social media saying this is a captured Russian soldier. But you might see people on the other side saying actually that's a paid actor, and here's the name of the paid actor. So with this technology, the level of accuracy can be used to identify if someone is who they say they are.

So the question is asked: Who in Ukraine has logins to Clearview AI? Hoan: It's in six different agencies, including the National Police of Ukraine. Then the people on the ground would be using it. So we gave a demo. We'd give them training on how to use facial recognition technology responsibly. So part of the training is that they would send photos of unidentified people and run it through the system, and we could show them, "Hey, this is how you verify who it is." So for example, if they have a tattoo that matches online, there's a very high chance it's the same person.

So the questioner: You inserted this technology into a war zone, which presents a lot more problems than having a police department in the United States using it. How are you accounting for that? Hoan responds: You want to be careful to make sure that they really know how to use it properly, and so there are all these scenarios that we want to make sure don't happen. For example, what if someone takes a photo of someone and says, "Hey, I think you're a Russian traitor," and then they detain them. And it's all incorrect based on incorrect information. So we'd never want something like that to happen. As long as these investigations are done by trained investigators, this can be a very useful tool. If it was used by everybody, I think that's when problems happen.

So the interviewer: What's the most surprising use you've seen in Ukraine? Hoan says: War crimes investigations. We were talking to people in the National Police of Ukraine and others where they have video footage from surveillance cameras. Smartphones are more prevalent, so there's a higher chance of something being recorded. Now, with this technology, if it's easy to identify someone, I think people are going to think twice about war crimes. So that was a surprise to me, he said.

Okay, so the interviewer asks: So this is a subscription service, and you say that gives you more control if someone is misusing it. How does Clearview AI work? He said: We vet every person to make sure they're a government official who's using it. There's also

two-factor authentication, so they still have to verify their device before they log in. Once they have an account, there's an administrator for each agency. So, the administrator can see who is conducting searches, and what reason they're conducting the search. There's an intake form that requires a case number and a crime type before conducting a search. So people, when they're on-boarded and they've learned about the software, they know that their searches can be audited because we want to make sure they're using it for the right kind of stuff. Because it's a cloud service, we have the ability to revoke access. If there's any egregious abuse of the technology, you want to make sure that we have the ability to take it away.

The interviewer asks: The IT Army appears to be using it. In a video, the group demonstrated the use of a facial recognition program that appears to resemble Clearview AI. This is a volunteer hacking force, so how is it that the Ukraine IT Army appears to be using Clearview AI? Hoan says: All I can say is that everyone we've onboarded is a government official. We haven't onboarded anyone in the IT Army directly. Everyone we talk to and onboard, we give them proper training on its usage. The speculation that the IT Army is running a Clearview AI search does not match any information we have on this matter. Clearview AI is intended for use in Ukraine by law enforcement and government officials.

So then the questioner: Did maybe somebody give a username and password to somebody who's in the IT Army? Hoan says: It's possible that someone shared a screenshot or shared how it worked. But we want to make sure that whatever the usage of the technology is say it is to identify someone deceased that is done in a way that is positive. The policy of the National Police and all our users is to tell the family members in a humane way.

Question: Had it occurred to you that the IT Army would use this technology to notify families of dead soldiers as a propaganda tool? Hoan says: I talked to some of the officials in the Russian government, and I said, "Look, is this something you knew about? Is that your procedure for doing that?" Then said that's not our official procedure. And they assured me that's not what they want to have happen either. Again, this is war time. Tensions are really high. Those things can happen. Hoan continues: If I thought it would be used in a really bad way, then I don't think I'd give access to them. We think that just getting the information out in a humane way is the most important thing. What we can control as Clearview is giving access to the right people. So for example, we don't give access to the Russians or anything like that, and we make sure Ukraine is trained as appropriately as possible.

Question: Have you revoked any access related to Ukraine because you thought it wasn't being used properly? Hoan replies: No, not at this time. But the administrators of these agencies in Ukraine have the ability to do so. They can go in and audit the searches, remove access to an account, and give access as they deem appropriate. Clearview AI would only revoke access to an agency if there was an egregious amount of abuse in that agency. Until something really escalates to that level, we haven't revoked any access.

The interviewer: The NSO Group is an Israeli company that makes surveillance software that can be remotely implanted in smartphones. It has come under heavy criticism for its tech being used by authoritarian governments to spy on their citizens. With your facial recognition technology, how do you avoid the NSO trap? Hoan replies: I think NSO is a very different kind of technology than what we do. We are searching public information from the Internet. So it is really just like Google. If you can type someone's name and the word "LinkedIn" into Google, and you find their photo, then Clearview can basically do the same thing, but it's search by photo. We also work with law enforcement. NSO is different because it's breaking into phones, very private data. Also, when they sell their software, they don't have the ability to take it back if they sell it to the wrong actor or government; whereas Clearview's software is deployed in the cloud. If we ever find

anything totally egregious or abusive, we have the ability to revoke access. We also want to be aligned with democratic countries, making sure that this is a technology that can be used responsibly and across democratic nations.

And lastly: Can you imagine a scenario is asked, years and years from now, when everyone has this capability, that it would be like VR glasses or built into a phone? Hoan says: I can imagine like augmented reality is interesting where it could be deployed on military bases. So in Afghanistan, they had a situation where they were pulling out at a checkpoint, terrorists could blow up people. They are looking up close at the IDs. To verify someone at a distance in a hands-free way, I think that's a very positive kind of use case.

So it's obviously trivial to invent synthetic positive use cases. It's at least as easy, probably much easier actually, to imagine quite privacy-intrusive use cases. The logic we often hear being applied by those who object to the use of pervasive facial recognition is the assumption of privacy. They suggest that someone walking in a public space can be seen by everyone, so there's no assumption of privacy. But when a meta-tagged photo is taken among friends in a private setting and posted online to celebrate, is it reasonable for those images to be vacuumed up, identified, catalogued, and made available for searches globally? The argument is made that posting the photo online creates implied consent. But is that for any use whatsoever, forever?

So we have another new problem created by astonishingly inexpensive technology. So as a society we're going to have to figure out what's important and what we want and how we want these things to work. And I don't have an answer. It's clear from the way Hoan was answering that...

Leo: He's on the defensive; right?

Steve: Yes, that he's been sensitized to all the ways this technology of his, which I'm sure is generating lots of revenue for Clearview AI, all the ways it could be abused. So.

Leo: Steve? What you looking at?

Steve: I just turned my temperature down a degree.

Leo: Oh, it's a little warm, a little toasty, isn't it. Yeah.

Steve: Okay. So that brings us to Pwn2Own Vancouver 2022.

Leo: Oh, boy.

Steve: Last Wednesday, Thursday, and Friday, the 18th through the 20th, was the 15th anniversary of Pwn2Own held in Vancouver. And this year, rather than enumerate the victories of many brilliant hackers who we don't know, and many of their names which I cannot begin to pronounce correctly so I'd rather not mangle them...

Leo: No, good.

Steve: I'm going to focus upon the products which we do all know, and which fell to their best efforts. On the first day, Wednesday of last week, the various hacker individuals and teams demonstrated an improper configuration against Microsoft Teams. And I didn't put all the dollar amounts in here. But as I recall, that one got them $150,000. So that was some serious improper configuration against Microsoft Teams. There was also an out-of-bounds read and out-of-bounds write, which was used to achieve privilege escalation against Oracle's VirtualBox; a three-bug chain of injection, misconfiguration, and sandbox escape against Microsoft Teams; prototype pollution and improper input validation against Firefox; an out-of-bounds write escalation of privilege on Windows 11.

Two bugs on Ubuntu Desktop, an out-of-bounds write and a use-after-free, were used. There was a zero-click exploit of two bugs, an injection and an arbitrary file write against Microsoft Teams; an out-of-bound write against Apple Safari; and a use-after-free elevation of privilege again against Microsoft Windows; and a use-after-free exploit on Ubuntu Desktop. That was just the first day.

Day two, the hackers and some teams demonstrated two unique bugs, a double-free and an out-of-bound write, with collision on a known sandbox escape on a Tesla Model 3 Infotainment System. They were able to basically take over the Tesla Model 3's infotainment system. There was also a use-after-free bug leading to elevation of privilege on Ubuntu Desktop, an improper access control bug leading to an elevation of privilege on Windows 11, and a use-after-free bug leading to an elevation of privilege on Ubuntu Desktop.

And on the third and much less active day there was an escalation of privilege via an integer overflow on Windows 11, a use-after-free exploit on Ubuntu Desktop, an elevation of privilege via improper access control on Windows 11, and an elevation of privilege via use-after-free on Windows 11.

So the big targets and losers were Microsoft Teams, Windows 11, and Ubuntu Desktop, with VirtualBox, Firefox, Safari, and Tesla's infotainment system each taking one hit apiece. Overall, for this 15th anniversary event, there were a total of 21 exploitation attempts, three of which failed, from a total of 17 contestants, with Trend Micro and their Zero Day Initiative awarding a total of $1,155,000 in Pwn2Own prize money.

So another great competition. And all of the details of the exploits, as always for Pwn2Own, are provided to those whose products fell. And here we are, where are we, it's not going to probably make it into next month's patch round. And Windows 11 isn't even released yet, still in dev channel mode. So, wait, what? No. The current version will be updated probably, well, maybe not by next month, depending upon the nature of these problems and how quickly Microsoft's able to fix it.

Okay. Under the heading of "Sometimes they get it right," we have last Thursday's very welcome news that the U.S. Department of Justice, our DOJ, has revised its policy on how federal prosecutors should charge violations under the Computer Fraud and Abuse Act, creating - get this - an explicit carve-out exemption covering "good faith" security research. So this is huge. What this means is that the U.S. government is altering how vigorously it enforces a central cybercrime law that security researchers, civil liberty advocates, and others have long argued is overly broad. And to that I say amen.

Under the change to enforcement of this CFAA - the Computer Fraud and Abuse Act - which became law back in 1986, the DoJ will amend its charging policy to explicitly discourage going after so-called good faith, or ethical, security researchers. The DOJ's Deputy Attorney General Lisa Monaco said in a statement accompanying the revamped policy: "Computer security research is a key driver of improved cybersecurity. The department has never been interested in prosecuting good faith computer security

research as a crime, and today's announcement promotes cybersecurity by providing clarity for good faith security researchers who root out vulnerabilities for the common good." Wow.

Federal prosecutors who seek to bring charges under CFAA must now first consult with the Computer Crime and Intellectual Property unit inside DOJ's Criminal Division. If that office recommends going forward with charges, prosecutors must then inform Monaco's team and even then may need special permission to proceed. However, it's still not all clear sailing. Ethical researchers who scour for and discover software vulnerabilities could still face prosecution under existing state laws or be sued in civil court.

So this updated guidance comes a little over a year after the Supreme Court ruled in a major CFAA case that the 1986 law does not apply when an authorized user uses data in improper ways. In that case, the court said that a Georgia police officer did not violate the hacking law - maybe some others, but not the CFAA - when he took money from an acquaintance to search a license plate database. In other words, someone paid a police officer to access the license plate database to which the officer himself had authorized access, even though the purpose of that specific instance of access was unwarranted.

The officer was sued under the Computer Fraud and Abuse Act, and the Supreme Court ultimately said "nope." The DoJ explained that the CFAA should only apply in instances when an outside hacker or authorized user - so even an authorized user, but external - actually breaks into a secure portion of an organization's network. So clearly the instance of a police officer using his authorized access for an unauthorized purpose would not be a CFAA violation.

As expected, the news of this updated policy was broadly welcomed among both federal cybersecurity officials and the researcher community. Jen Easterly, who's CISA's director, tweeted: "Huge news. Well done, Team DOJ!" And Chris Vickery, a prominent cyber researcher, tweeted that the new guidance will "hopefully improve the lives of people like me who fear retaliation for trying to do the right thing."

So as I said, sometimes they get it right. We still have a ways to go since local government and private parties who take offense at the discovery of their own cyber failings being revealed will still be able to somewhat unjustifiably exact their revenge. But perhaps a competent defense attorney will be able to point to the changes which have just been made at the federal level as a basis for applying the same reasonable protections more locally. We can hope.

Okay. So you have a fancy car, or a residential or office door lock, which operates as follows: Your smartphone containing the lock's matching app, and essentially or effectively its key, is in your pocket. You approach your car or your home's or office's front door and just touch the lock. The lock's capacitive sensor senses your touch. So it emits a Bluetooth Low Energy, a BLE ping to query for any nearby, in this case, Kwikset Kevo smart lock apps that may be nearby. The app in the phone in your pocket receives the Bluetooth Low Energy ping and responds. So they engage in a super-crypto triple-scoop post-quantum impossible-to-crack handshake negotiation. And now, satisfied that one of its authorized owners is indeed nearby, the lock disengages to allow entry. And all is happy.

But now we have a different scenario. A bad guy team wants to gain entry into that smart lock-protected car or residence or office. So one of them arranges to place a Bluetooth Low Energy relay near the user's phone, perhaps in an adjacent office cubicle, or next to their locker while they're working out, or in the coffee break room which they frequent. Wherever, it doesn't matter.

The other member of the team waits some discreet distance away from the car, the residence, or the office that's about to be breached. When the first team member messages to the second that the unsuspecting user's smartphone is within range of their remote Bluetooth relay, the second member of the team simply walks up to the locked car, residence, or office and touches its lock. Exactly as before, the lock sends out a BLE discovery ping, which the BLE relay forwards to its matching endpoint.

That distant endpoint simply and blindly rebroadcasts the ping it received, which the user's smartphone picks up and acknowledges. It replies, and its reply is similarly forwarded back to the BLE relay which is now positioned near the lock. So once again they engage in the most unbreakable bazillion qubit quantum-entangled crypto that the world has ever hosted, until the lock becomes satisfied that only the user's smartphone can possibly be at the other end of the link. So it disengages its lock to admit the individual whom it presumes is authorized to gain access.

Unfortunately, this is not fiction, except in this example for my overkill use of deeply entangled quantum crypto. The point being, this oh-so-simple Bluetooth Low Energy relay attack entirely defeats any system's crypto, no matter whether it's pre- or post-quantum. And as I said, this is not fiction.

Last September, after successfully executing exactly this scenario in the field, the U.K.'s NCC Group notified several smart lock makers that they had a problem. Their systems were vulnerable to a simple BLE relay attack. And the NCC group notified these lock makers that they would eventually be publishing the news of this, which they did last week. The accompanying security advisory simply explains: "An attacker within BLE signal range of a smartphone or key fob authorized to unlock a Kevo smart lock can conduct a relay attack to unlock the lock over long distances."

And the problem is there is no obvious in-band way of detecting and preventing this abuse. Which is to say that Bluetooth Low Energy does not offer robust endpoint proximity protection. Various out-of-band solutions have been considered, one being for the smartphone to use its own GPS to ascertain whether it's physically near the lock it has been paired with. But concerns over GPS's availability and speed have been a concern. And local GPS jamming and spoofing technology is available.

Another out-of-band solution considered was to have the app monitoring its owner's physical movement and to disable any lock negotiation if the smartphone was motionless immediately before the negotiation began, since that would never be expected in the system's normal use case. But any scenario where the user would be moving, even if many miles away, would defeat such anti-spoof protection. So GPS would seem to be the best, if still imperfect, solution.

But the nature of the problem suggests that the use of Bluetooth Low Energy, convenient though it is, was not the best idea in the first place. If our smartphones were equipped with radios which incorporated reliable time-of-flight measuring capabilities, then it would be possible to obtain spoof-proof data and direct physical proximity assurance. But today's smartphones are not yet that smart.

So again, apparently the manufacturer scrambled around. I looked at the timeline that these guys, the NCC Group, posted in their advisory. And the issue got escalated from the manufacturers and retailers up to the behind-the-scenes designers whose company name you never learn, where this was all actually created. And there was a lot of head scratching going on, and consternation for it to be discovered and would be published that their system was susceptible to this kind of spoofing. Again, obviously not widespread. It would be used in a targeted attack case. But if bad guys knew that this was possible, this BLE relay technology is not difficult to create. And it robustly unlocks a

target at some distance. So probably not such a good idea to have designed a system that could be spoofed like that.

Okay. A couple of closing-the-loop pieces. Two of them, as I said, are bigger than usual. Brian Phillips tweeted: "Hi, Steve. Quick one. Don't know if you've covered this already, but do you have an alternative for" - actually he said Unlock Origin. He then corrected his typo, it was probably an auto correct. He meant Unblock Origin, or sorry, UBlock. "Do you have an alternative for UBlock Origin on Safari?" He said: "Recently got myself a MacBook Pro and now want to use Safari over Chrome." That is to say, instead of Chrome. And Leo, I had suggested Adblock Plus.

Leo: No, there's nothing good on Safari. There really isn't.

Steve: That's what I wanted to ask because I knew you would know. Adblock Plus was the only thing I knew as like sort of equivalent to UBlock Origin.

Leo: Safari changed its plug-in module, and Gorhill did not like it and is not supporting Safari. I would just use Firefox. I use it on all platforms, including the Mac, and it supports UBlock Origin just fine and also gives you [crosstalk].

Steve: Brian, hope you're listening. I'll see if I can find - I think this was an exchange via DM. So I'll forward that to Brian in case he misses it.

Leo: I mean, there are filters. Adblock has always been...

Steve: There are a bazillion of them; right?

Leo: Because they're the ones that have the, what do they call it, the policy for approved ads. And I really don't like to support that.

Steve: Right.

Leo: Yeah. I can't remember, the one I used had three different plugins. But the problem is the iPhone. You're really - there's nothing you can do on the iPhone. So in that case I'd just use NextDNS and block it at the DNS server. Which is pretty effective as well as an adblocker. That blocks everything at the DNS level instead of on the device.

Steve: And he is MacBook Pro; so as you said, Firefox would be an alternative.

Leo: Firefox is great, yeah. I use Firefox on everything, including on my Macs. It's iOS that's problematic because you have to use WebKit.

Steve: Right, right.

Leo: And so then you have to block it at the source instead.

Steve: So Dave Badia asked: "Can you give a shout-out on Security Now! that Voyager 1 is still happily trotting along after so many years?" He says: "And I just read it goes into 'safe mode' if things get out of whack so the programmers on earth can fix it."

Okay. Well, so what's interesting is that Dave sent this last week, either before, or more likely just exactly because things just started to go wonky with Voyager 1. Now, okay, believe it or not, Leo, Voyager 1 has a Twitter account.

Leo: Of course it does.

Steve: Of course it does. Though being 14.5 billion miles away, with light requiring currently 21 hours, 34 minutes, and 38 seconds to traverse the distance, for a matter of convenience, NASA's public relations people have been posting on Voyager's behalf.

Leo: Of course, yes. Good thinking.

Steve: Yes. So last Wednesday Voyager posted on Twitter: "Do you ever feel misunderstood? My team is investigating an issue with my data. Even though I'm sending signals and operating normally, some data readouts don't exactly match what's happening out here. While they investigate, I'll keep doing my thing." Okay, Voyager.

So this made me curious about what was going on, so I dug around and found some commentary which I've lightly edited for the podcast. It reads: "NASA Voyager 1 space probe was launched 45 years ago and continues its journey as the first-ever human-made object to leave the vicinity of our solar system." Now, this is, I think, critical to my theory. It says: "It's heading out there to study the outer heliosphere and the interstellar medium. The iconic probe has sent hugely important data back to NASA since its launch on September 5th, 1977. But now the new strange data being sent by the Voyager from the edge of the solar system has left scientists shocked since, until now, there have been no significant errors reported by the probe." And I have a somewhat chilling theory for what might be going on, but I'll finish this before sharing that.

So this continues: "Since the Voyager 1 data is of critical importance, the engineering team is trying to solve the puzzle of mysterious data now being sent by the space probe. NASA said in a statement: 'The interstellar explorer is operating normally, receiving and executing commands from Earth, along with gathering and returning science data. But readouts from the probe's Attitude Articulation and Control System (AACS) do not reflect what's actually happening onboard.'

"But what does this data received back on Earth actually mean? NASA says an antenna attached to Voyager, which is pointed at Earth to send data back, appears to be working, but is sending back invalid data. The AACS controls the 45-year-old spacecraft's orientation. Among other tasks, it keeps Voyager 1's high-gain antenna pointed precisely at Earth, enabling it to send data home and to receive new instructions. All signs suggest the AACS is still working, but the telemetry data it's returning is invalid. For instance, the data may appear to be randomly generated, or does not reflect any possible state the AACS could be in. Thankfully, the issue with the NASA Voyager 1 has not triggered the probe's onboard fault protection system, the system which is designed to keep the spacecraft in its 'safe mode' which maintains a state where only essential operations are carried out, while giving engineers time to analyze and diagnose the issue. And Voyager

1's signals have not weakened, which suggests the high-gain antenna indeed remains in its prescribed orientation to Earth."

Okay. So Leo, nothing manmade, and certainly nothing sending back telemetry from such a distance, has ever been this far from Earth and our sun. What if Voyager is approaching the maximum radius supported by the simulation?

Leo: Are you saying it's going to fall off the edge?

Steve: And that's why...

Leo: It's falling off the edge of the simulation.

Steve: It's leaving reality, Leo.

Leo: Well, well, well.

Steve: That's why its reality is beginning to break down, and why it appears to be returning random data that makes no sense. Everything else we see out there appears to just be inbound radiation. It's not interactive. Voyager is, or at least it was.

Leo: I like that theory. I like it.

Steve: So it'll be interesting to see what happens, and whether perhaps the same thing happens to Voyager 2 once it gets around the same distance away.

Leo: How long will that be?

Steve: It's about 2.4 billion miles closer to Earth than Voyager 1. So we have some time to wait. I do have a link in the show notes. You can click on that, and it will bring up the Voyager Mission Status that is live from JPL, showing their respective distances, the time that they've each...

Leo: That's cool.

Steve: That each mission's been running and how far away they are from the sun and so forth.

Leo: They need an entry for how close to the edge of the simulation are they.

Steve: I know, yeah.

Leo: I'm sure there's a simpler explanation than that, Steve.

Steve: That is the simple explanation. What could be simpler than you've gone too far away? And the simulation can't handle distances that great. It's doing everything up to that, I mean, think about that, 14.5 billion miles.

Leo: It's a lot, yeah.

Steve: What's the cubic space of that?

Leo: Right.

Steve: I mean, that is a ton of space to simulate, with all the chemical reactions and molecules bouncing around in there. That would tax any big simulator. So it makes sense that there would be some maximum radius; right?

Leo: Very 2001. I like it.

Steve: So I love this Voyager probe. It's reminiscent of the Mars Rovers, right, that we were talking about for quite a while, Spirit and Opportunity. Okay. Remember how "Gilligan's Island" was an adventure that started off as a three-hour tour. Right? Similarly, the two Voyager spacecraft were originally intended to only study Jupiter and Saturn, their moons, and Saturn's rings. And for that two-planet mission, they were built to last five years.

That was 45 years ago. Their 45th anniversary of launch will be coming up this August and September. And after their initial successes, NASA's engineers doubled the missions' objectives to include two more giants, Uranus and Neptune. So now, between the two craft, they've explored four planets, 48 moons, a host of planetary magnetic fields and rings.

They are each powered by a trio of radioisotope thermoelectric generators, known as RTGs, which convert the heat generated by the radioactive decay of plutonium-238 into electricity. And while they are super reliable and have no moving parts, I mean, just super robust electricity generators, that decay itself decays over time. Right now, the Voyagers' power is gradually dropping at a rate of four watts per year.

This has been limiting the total number of systems that can be powered up at once, and the Voyager mission teams have been turning off more and more non-essential equipment to reserve power and keep these things going. But so far none of Voyagers' science instruments have needed to be powered down. And the latest of the continually extending goals has been to keep both of these Voyagers running for at least another three years, beyond 2025.

Now, I did not do the velocity calculations to tell us whether Voyager 2 would be at the same end-of-simulation radius by 2025, when presumably we would still be watching it. But, you know, if it spontaneously disappears, well, that will be because it crossed the other side of the simulation boundary.

Leo: So the simulation is just the solar system. The rest of it is just pictures. It's like a scrim, a backdrop.

Steve: Yeah. It's to give us something to scratch our heads and think about, and something to aspire to. But, yeah, it pretty much covers everything it needs to, to keep us busy running around in our little cage.

Leo: It is kind of amazing. To think that there is a manmade object out there that far away is amazing.

Steve: Oh, Leo. And it's alive. It's still...

Leo: And it's still reporting back, yeah.

Steve: Wow, 14.5 billion miles away. So that light takes 21-plus hours to get back.

Leo: Wow, that is far away.

Steve: And think of the accuracy of that antenna. That's the dish; right? Think of how, I mean, it's just astonishing to me that it is able to maintain alignment on the Earth at that distance. You couldn't calculate the fraction of a degree off it would be, in which case its beam would just completely miss us.

Leo: Really an amazing story. All right. I like your theory. It's as good as any. I'll go for it.

Steve: Yeah.

Leo: We don't know.

Steve: And the good news is none of us have to worry about getting that far away.

Leo: Right. That's right.

Steve: So we're all good.

Leo: Yeah. We're still well within the simulation.

Steve: You do not want to go further than 14.5 billion miles away. Could be bad.

Leo: You'll start speaking random gibberish.

Steve: So Tom Feller said: "Hi, Steve. I came across an interesting article I ran across on quantum computing." And the article was titled "Q-Day Is Coming Sooner Than We Think."

Okay. As we've discussed before - I should mention this was published in Forbes. And Forbes has had a somewhat spotty record for their presentation of technical subject matter over time. But assuming that the editors don't butcher it, the content's veracity will be a function of its author. In this case the article's author is a guy named Arthur Herman whose bio tells us that he's a Senior Fellow at the Hudson Institute, the Director of the Quantum Alliance Initiative, and the co-author of "Risking Apocalypse: Quantum Computers and the U.S. Power Grid." So it sounds like he may know what he's talking about.

Arthur is very clearly of the opinion that we should be much more worried than we currently are. Having read what he wrote, and even though I'm far from able to render any opinion on the subject, I know that our listeners will find this as fascinating as I did. So once again, I've edited what Forbes published so that it better fits our higher end audience.

Arthur said: "'Q-Day' is the term some experts use to describe when large-scale quantum computers are able to factor" - and he said "factorize," and I thought, I don't know, maybe he does not quite know what he's talking about. But anyway, I fixed it. Or maybe "factorize" is just British or something. I don't know. "...[T]he large prime numbers that underlie our public encryption systems, such as the ones that are supposed to protect our bank accounts, financial markets, and most vital infrastructure." And this is where he's tying it into the power grid, saying that the power grid may not be safe if you can break crypto. I'm sure he's right about that.

He says: "That's a feat that's all but impossible for even the fastest supercomputers, but which the unique features of quantum computers" - which I've looked at several things he's written, and he certainly understands this - "using the physics of superpositioning and entanglement, will be able to deliver. There's a growing consensus that this quantum threat is real; there's no agreement on how long it will take before a quantum computer has the 4000 or so stable qubits it will need to meet the requirements of Shor's algorithm for cracking those encryption systems."

But it turns out there's a different way to do this than using Shor, which we'll get to here in a second. He says: "For example, it would take a classical computer 300 trillion years to crack an RSA 2048-bit encryption key. A quantum computer can do the same job in just 10 seconds with 4099 stable qubits." Again, stability. So the count and stability are two important criteria. He says: "But getting to that number is the main problem quantum computer engineers face, since the stability or coherence of qubits lasts for only microseconds." And again, we need 10 seconds of stability from 4099 stable qubits. He says: "Today's most entangled computer, Google's Bristlecone, has just 72 stable qubits.

"Nonetheless," he says, "I have been arguing for the past four years, including in this column, that Q-Day is likely to come sooner than even quantum scientists predict, and that the time to get ready to protect our vulnerable data and networks is now. Others prefer to procrastinate, citing experts who say that a threat is at least a decade or more away. The fact that the National Institute of Standards and Technology (NIST) won't have its quantum-resistant algorithm standards ready until 2024, and expects the rollout to take another five to 15 years, has helped to encourage complacency disguised as confidence.

"But new developments in quantum science suggest that this complacency may be misplaced. So-called quantum annealers like the one Canada-based D-Wave Systems,

Inc. uses, are able to calculate the lowest energy level between the qubits' different states of entanglement, which equates to the optimal solution. These machines have proven their worth in solving optimization problems that usually stump classical computers, as I explained," he wrote, "in a previous column."

He says: "Not surprisingly, scientists have been quietly finding ways to turn factorization into an optimization problem, thus bypassing Shor's algorithm, the paradigm for discussing quantum decryption since the 1990s." And of course which we talked about last week. "In 2019, scientific papers emerged that showed how to do this, including factoring integers using 'noisy' qubits, i.e., swarms of quantum bits that aren't perfectly entangled the way a large-scale computer requires. In other words, no longer requiring stable qubits.

"One was authored by Chinese scientists who found a way to factor a large number using only 89 noisy qubits. They then showed it's possible to factor an RSA 768-bit private key, which is the current factorization record accomplished using classical computers, with 147,454 noisy qubits a fraction of the millions of noisy qubits a large quantum computer would need to reach the 4000 stable qubit threshold, and within reach of the architecture of an annealer like D-Wave Systems.

"Also in 2019, a pair of Google researchers and the Royal Institute of Technology at Stockholm published a paper showing how to crack 2028-bit RSA integers in eight hours using 20 million noisy qubits. Given the fact that in 2012 scientists speculated it would take one billion qubits to perform this feat, it will likely not be long before researchers show they can get there with a lot fewer than Google's 20 million noisy qubits.

"And sure enough, in 2020 three Chinese researchers found a way to use the D-Wave quantum computer to factor large integers, completely bypassing Shor's algorithm. 'Thus,' they concluded, 'post-quantum cryptography should consider further the potential of the D-Wave quantum computer for deciphering the RSA cryptosystem in the future.'

"In effect," Arthur writes, "these researchers found a way to turn decryption using quantum technology into a straightforward process on a timeline much shorter than ten years. Perhaps four to five years is more likely. This was what Chinese scientists were openly publishing," he writes. "We don't know what's happening behind the scenes, but we can bet if there's a short cut to achieve what a large-scale quantum computer can do using annealing technology, their military and intelligence services will want to find out.

"All of this should change the timetable for Q-Day and for our strategic calculations. Not only is quantum-based decryption coming our way sooner, but thanks to annealing, this code-cracking feature will be more accessible to other machines than the hugely expensive large-scale computers Google, Microsoft, and others are working on, which in turn puts the capability within reach of small-state or even non-state actors."

He finishes: "Why gamble with the quantum future? Annealing technology makes becoming 'quantum ready' more important, and getting started now more imperative than ever." So, okay. I'm going to take his word for it. I can't argue the pros and cons. I'm programming in assembler on a machine with 32 or 64 bit bits, not cubits, binary bits. But, wow, interesting stuff. And the good news is I'm sure that the academics and those who are deciding these things are aware that these other breakthroughs are occurring.

Unfortunately, I think Dis-CONTI-nued is the best way to say that. I've been looking at that.

Leo: Dis-CONTI-nued.

Steve: Dis-CONTI-nued, yeah.

Leo: Discontinued.

Steve: So last Thursday, Advanced Intel is the name of this organization; advintel.io is their domain. Advanced Intel's Yelisey Boguslavskiy tweeted: "Today the official website of Conti Ransomware was shut down" - this is last Thursday - "marking the end of this notorious crime group." He says: "It is truly a historic day in the intelligence community."

And the day after that, last Friday, they published their report about exactly what happened. There's so much more to it than just someone turned the site off that I felt certain our listeners would find the details fascinating. And their report is titled - don't blame me, although I did perpetuate it - "DisCONTInued: The End of Conti's Brand Marks New Chapter for Cybercrime Landscape." And the top of their report teases, reading: "From the negotiations site, chat rooms, messengers to servers and proxy hosts, the Conti brand, not the organization itself, is shutting down. However, this does not mean that the threat actors themselves are retiring."

Okay. What does it mean? Advanced Intel apparently rushed out their report. It contains some typos, misspelling, and grammatical awkwardness. And they may not be native English speakers. So in order to share it with the podcast, I cleaned it up a bit. But otherwise it remains what they wrote. And I think everyone's going to find it interesting.

They said: "On May 19th, the admin panel of the Conti ransomware gang's official website, Conti News, was shut down. The negotiations service site was also down, while the rest of the infrastructure, from chat rooms to messengers and from servers to proxy hosts, was going through a massive reset.

"Conti News, a shame blog, is the last beacon of the group's public operation where victim data was being published. It also served as a media tool that Conti used for their endless public statements, one of which led to the gang's downfall." We'll get to that in a minute. I have a snapshot of it later in the show notes. They said: "This publicity function of the blog is still technically active; and this activity, as shown below, is highly strategized. At the time of this publication, May 20th, 2022, Conti was even uploading anti-Americanist hate speech claiming the USA to be 'a cancer on the body of the earth.'

"This, however, only manifests that the website became an empty shell. At the same time, the crucial operational function of Conti News, which was to upload new data in order to intimidate victims to pay, is defunct, as all the infrastructure related to negotiations, data uploads, and hosting of stolen data was shut down."

Okay. "And this shutdown," they wrote, "highlights a simple truth that has been evident for the Conti leadership since early spring of this year. The group can no longer sufficiently support and obtain extortion. The blog's key and only valid purpose is to leak new datasets, and this operation is now gone. This was not a spontaneous decision," they write. "Instead, it was a calculated move, signs of which were evident since late April. Two weeks ago, on May 6th, Advanced Intel explained that the Conti brand, and not the organization itself, was in the process of the final shutdown. As of May 19th, 2022, our exclusive source intelligence confirms that today is Conti's official date of death.

"In this retrospective analysis, we will not only take an in-depth look into the reasons behind the Conti shutdown but, perhaps most importantly, assess and project the future of a new threat landscape that is already on the horizon. But first we need to review how Conti prepared for its own demise and how this group, notable for its sophistry,

continued to utilize information warfare techniques to orchestrate the shutdown until its final days in order to ensure the legacy of its surviving members."

They explained: "Shutting down ransomware's iconic criminal brand is a long and complicated venture. A notorious and prolific threat group cannot simply turn off its servers, only to pop back up the following week with a new name and logo design. Even a whisper of novel threat group activity following the announcement of Conti's demise would likely spark immediate accusations of poorly executed identity theft. At best, immediate comparisons between the two would permanently leave the new group in Conti's ghostly shadow, the collective that fell and the one which emerged." And I'll just note that we've seen and commented on exactly this with previous ransomware operations.

So these guys said: "REvil, DarkSide, and countless other collectives attempted the disappearing act. The simple approach failed miserably. As what was one of the predominant ransomware groups active at the time, Conti realized that an element of 'performativity,'" they wrote, "would need to be involved. Where other groups had been attempting a grand stunt with smoke and mirrors, Conti would try a sleight of hand.

"Conti would not be itself without its project frontman, an individual operating under the alias 'reshaev,' a.k.a. 'gangster.' Besides being a talented coder" - this reshaev was behind the original Ryuk payload - "this person was an outstanding organizer. It was reshaev who set the foundation for Conti's dominance in the cybercrime business by creating an organizational system based on skill, teamwork, clear business processes, hierarchy, and clear foresight.

"It's not surprising that reshaev was the first who saw Conti's structural challenges. Due to the group's public allegiance to Russia in the first days of the Russian invasion into Ukraine, Conti was unable to be paid. Since February, almost no payments were given to the group, while Conti's locker" - the slang for malware - "became highly detectable and was rarely being deployed. The only possible decision was to rebrand.

"For over two months, Conti collective has been silently creating subdivisions that began operations before the start of the shutdown process. These subgroups either utilized existing Conti alter egos and locker malware, or took the opportunity to create new ones. This decision was convenient for Conti, as they already had a couple of subsidiaries operating under different names: KaraKurt, BlackByte, and BlackBasta. The rebranded version of Conti, the monster splitting into pieces but still very much alive, ensured that whatever form Conti's ex-affiliates chose to take, they would emerge into the public eye before news of Conti's obsolescence could spread, thus controlling the narrative around the dissolution as well as significantly complicating any future threat attributions."

And then they wrote: "This is where the plans for what was left of Conti became increasingly complex. In order to hide the fact that Conti was now dispersed and operating via smaller, more novel brands, the former affiliates of the gang had to now convincingly simulate the actions of a dead brand. Conti's remaining infrastructure operated like an army preparing for an ambush. Lingering actors were left to keep their fires lit, visible from behind enemy lines. Meanwhile, hidden from view, Conti's most skilled agents were instead laid low in a nearby encampment, biding their time while watching their great and empty camp send out smoke signals, meticulously emulating the movements of an active group.

"Conti continued to publish documents stolen from victims, most likely targets hit earlier with attacks and lined up in a sort of queue waiting for public release, and campaigned hard for themselves on criminal forums. Their public persona boasted a strong and enduring foundation, even one that was willing to further expand the group's operations. From the perspective of Conti's posting history, the group appeared to be as strong as

ever." Okay. Then they shared a snapshot of a long and quite rambling chest-thumping post from March 30th, where a Conti representative talks up the group's successes, even seeking to recruit new affiliates, all apparently just smokescreen.

Then they continue: "However, in order to pull off their ultimate tactical maneuver, the agents left behind to operate from within Conti's massive empty shell now had to ensure that their antics would successfully lure attention away from their escaping comrades. To do this, they had to be certain that they left bait big enough to satisfy all of the opposing forces," stretching this analogy. "Conti would have to perform a grand finale, one big enough to live up to the group's name.

"And finally, on May 8th, Costa Rican President Rodrigo Chaves declared a national emergency as the result of a major cyber attack executed by the Conti ransomware gang. The massive attack, which took place against multiple Costa Rican government agencies, seems almost like a last-ditch effort by the group to squeeze a few more drops of riches from foreign government funds. However, Advanced Intel's unique adversarial visibility and intelligence findings led to what was in fact the opposite conclusion: The only goal Conti had for this final attack on Costa Rica was to use the platform as a tool to publicly perform their own death and subsequent rebirth.

"Advanced Intel has been tracking the preparations for this attack since April 14th, days before even the initial compromise. Our prevention alert was sent on April 15th, three days before the first incident compromising Costa Rica's Ministry of Finance occurred." Okay. So they said that. Then their report links to a tweet thread in Spanish, but it appears to be dated from the 18th. But they then provide a screenshot which indeed appears to substantiate a three-day early warning of an impending attack.

So they explain: "In our pre- and post-attack investigation, we have found three things. First, the agenda to conduct the attack on Costa Rica for the purpose of publicity instead of ransom was declared internally by the Conti leadership. Second, internal communications between group members suggested that the requested ransom payment was far below $1 million USD, despite unverified claims of the ransom being $10 million USD, followed by Conti's own claims that the sum was $20 million. A low demand such as this, made to a state entity no less, was only made with the knowledge that the group would never see payment for the ransom either way," because their payment pipeline had been completely foreclosed on by the sanctions against Russia and by their pronounced affiliation with Russia. "And third, Conti was very vocal about the attack, constantly adding new political statements." And that's this kind of junk that we talked about last week.

They say: "The attack on Costa Rica indeed brought Conti into the spotlight and helped them to maintain the illusion of life for just a bit longer, while the real restructuring had already taken place. While Conti had been busy with its diversion tactics, other brands such as KaraKurt, BlackByte, and numerous other groups which existed as extensions of Conti, but without taking the group's name, were extremely operationally active, although working in silence.

"Working concurrently with them, talented infiltration specialists who were ultimately the backbone of Conti's gang were also more active than ever, forming alliances with BlackCat, AvosLocker, Hive, HelloKitty/FiveHands, and a whole other cadre of ransomware groups. These pen testers maintain personal loyalty to the people who created Conti, but ultimately continued their work with other gangs in order to fully shed Conti's name and image. The situation presents the first and foremost reason for Conti's timely end: toxic branding.

"Indeed, the first two months of 2022 left a major mark on the Conti name. While there is no tangible evidence to suggest that the well-known Conti leaks had any impact on the

group's operations, the event which provoked the leak, Conti's claim to support the Russian government, seems to have been the fatal blow for the group, despite being revoked almost immediately."

And we noted this at the time. Remember that Conti posted: "The Conti Team is officially announcing a full support of Russian government. If anybody will decide to organize a cyberattack or any war activities against Russia, we are going to use our all possible resources to strike back at the critical infrastructures of an enemy."

"That statement had several key consequences," Advanced Intel wrote, "all of which deeply reshaped the environment Conti was operating within. First, by engaging in political discourse, Conti broke the first unspoken rule of the Russian-speaking cybercrime community, which is not to intervene in state matters. In Advanced Intel's public blog regarding REvil's ultimate takedown by the Russian government, AdvIntel provided an in-depth analysis of this unspoken agreement, making case studies of the two most notable groups to break it, Avaddon and REvil. With the ongoing Russian invasion of Ukraine, it may be very plausible that Russia's state security apparatus is attempting to exert governmental control over its cyberspace, even taking down groups that appear to have been allies, but who exhibited undue independence with their actions.

"Advanced Intel has seen internal communication of the Conti leadership suggesting that the Russian FSB had been pressuring the group. And even though non-factual evidence was involved, the REvil scenario may have simply repeated itself with Conti, the group's brand becoming a target for Russian authorities despite their pledged loyalties.

"Second, Conti's allegiance to the Russian invasion of Ukraine provoked internal conflict and brought shame on the Conti name from members who were either ethnically Ukrainian, or were Russian but supported Ukraine, or simply wanted to maintain an anti-war ethic. Considering that one of these members decided to betray the gang and leak private Conti chat logs" - we talked about that, too - "not long after the conflict began, this illustrated the final nail in Conti's self-made coffin. The third and most important factor: By pledging their allegiance to the Russian government, Conti as a brand became associated with the Russian state, a state that is currently undergoing extreme sanctions.

"In the eyes of the state, each ransom payment going to Conti may have potentially gone to an individual under sanction, turning simple data extortion into a violation of OFAC regulation and sanction policies against Russia. This liability came to a head on May 6th, when the U.S. State Department offered rewards up to $10 million USD for information that led to the takedown of the Conti group. As a result of these limitations, Conti had essentially cut itself off from the main source of income."

They wrote: "Our sensitive source intelligence shows that many victims were prohibited from paying ransom to Conti. Other victims and companies who would have negotiated ransomware payments were more ready to risk the financial damage of not paying the ransom than they were to make payments to a pro-Russian state-sanctioned entity.

"As Advanced Intel previously stated, the end of the Conti brand does not equal the end of Conti as an organization. As seen with the Costa Rica case, Conti has been carefully planning its rebranding for several months, preparing a comprehensive strategy to execute it. The strategy is based on two pillars. First, Conti is adopting a network organizational structure, more horizontal and decentralized than the previously rigid Conti hierarchy. This structure will be a coalition of several equal subdivisions, some of which will be independent, and some existing within another ransomware collective. However, they will all be united by internal loyalty to both each other and the Conti leadership, especially reshaev. At this point, this network includes the following groups,

the first type being autonomous, no malware locker involved, pure data stealing. That's KaraKurt, BlackBasta, and BlackByte.

"The second type being semi-autonomous, acting as Conti-loyal collective affiliates within other collectives in order to use their malware locker. That's AlphV/BlackCat, Hive, HelloKitty/FiveHands, and AvosLocker. The third type being independent affiliates, working individually, but keeping their loyalty to the organization. And finally the fourth type being mergers and acquisitions where Conti leadership infiltrates a preexisting minor brand and consumes it entirely, keeping the small brand name in place. The small group's leader loses their independence, but receives a massive influx of manpower; while Conti obtains a new subsidiary group.

"This is different from Ransomware-as-a-Service since this network, at least at the time of writing, does not seem to be accepting new members as part of its structure. Moreover, unlike Ransomware-as-a-Service, this model seems to value operations being executed in an organized, team-led manner. Finally, unlike Ransomware-as-a-Service, all the members know each other very well personally and are able to leverage these personal connections and the loyalty they bring." And implied in that of course would be some protection against U.S.-based bounties against their members, if they maintain a loyal, cohesive group. You know, one turns one in, and they're subjecting themselves to similar reprisal.

And finally they finish: "This model is more flexible and adaptive than the previous Conti hierarchy, while also being more secure and resilient than Ransomware-as-a-Service." And finally: "The other major development for this new ransomware model is the transition" - and this is really interesting - "from data encryption to data exfiltration, covered extensively by Advanced Intel in our analysis of KaraKurt and BlackByte. In a nutshell, relying on pure data exfiltration maintains most major benefits of a data encryption operation, while avoiding the issues of a locker altogether. Most likely, this will become the most important outcome of Conti's rebrand. The actors that formed and worked under the Conti name have not, and will not, cease their forward movement within the threat landscape. Their impact will simply leave a different shape."

So to our listeners, if anyone in your cyber sphere announces that Conti has shut down and disbanded, well, now we know better. It appears that earlier this year, as a consequence of, you know, we've talked previously about the entire reason that ransomware has come into existence, whether it be encrypting malware or exfiltrating and holding that data for ransom. It's the ability to get paid thanks to cryptocurrency, which has made that practical from an underworld standpoint. But the sanctions against Russia, Conti's original proclamation that they were standing with Russia essentially cut them off from extra-Russian payment of cryptocurrency into them. And that set them on a multi-month course to basically kill Conti off while continuing to function as a viable ransomware organization, learning from the mistakes they'd made before, changing their structure, and probably, apparently, changing the nature of what they do maliciously.

Leo: Well, they're not fooling anyone; okay? That's the truth.

Steve: We know better.

Leo: We know better thanks to you, Steve Gibson, and this fabulous Security Now! program every Tuesday about 1:30 Pacific, 4:30 Eastern, 20:30 UTC. You could tune in at live.twit.tv and watch it happen, audio and video available at that site. You could chat while you're watching at irc.twit.tv or in our Club TWiT Discord. After the fact, on-demand versions are available from a variety of places. Steve has two

unique versions of the show. He has a 16Kb audio, which he's been doing for years, for the bandwidth-impaired. That's the smallest audio form of the show. There's an even smaller version, the transcripts, which are plaintext, and you can read along. They're well done because they're written by an actual human being, a court reporter who can keep up with Steve. Thank you, Elaine. Those are available at GRC.com. That's Steve's website.

While you're there, pick up a copy of - oh, he also has a 64Kb audio version, by the way. While you're there, pick up a copy of SpinRite, the world's best mass storage recovery and maintenance utility. Version 6 is current, but soon 6.1 is going to come out. If you buy today, you'll get a copy of 6.1 the minute it's available. You also can participate in the development of 6.1. Leave messages for Steve at GRC.com/feedback or on his Twitter page. He's @SGgrc, and he takes DMs, so you can ask him a question in his direct messages, as well.

We have audio and video of the show at our website, TWiT.tv/sn. There's also a YouTube channel devoted to Security Now!, the video. And of course you can also subscribe, in fact that's probably the best thing to do, in your favorite podcast application. Pick the audio or video version, and you'll get it automatically every Tuesday evening after the show's over so you can listen at your leisure.