
# Credential Stuffing Attacks

**Description:** This week we look at more privacy fallout from our recent coverage of Facebook and Google. We examine the uptake rate of recent Windows 10 feature releases. We finally know the source of the AV troubles with the April Patch Tuesday updates. We look at the NIST's formal fuzzing development, consider the source of a massive and ongoing database data leak involving more than half of all American households, note that Windows Insiders are already finding that their systems won't update to the May 2019 feature update, and address the concerns of United Airlines passengers who have noticed and been understandably upset by seatback cameras pointing at them. Finally, we have the "Cranky Old Guy Tip of the Week," touch on a bit of miscellany, then take a look at what many in the security industry are watching with concern: the large and emerging threat of website credential stuffing attacks.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-712.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-712-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots to talk about, including the congressional questions for Google about Sensorvault, the changing velocity of Windows updates, and NIST goes a-fuzzing. It's all coming up next on Security Now!.

**Leo Laporte:** This is Security Now! with Steve Gibson, Episode 712, recorded Tuesday, April 30th, 2019: Credential Stuffing Attacks.

It's time for Security Now!, time to protect yourself online with the Securer in Chief, Mr. Steven Gibson of the Gibson Research Corporation. Hi, Steve.

**Steve Gibson:** And Leo, the microphone is back on the side where it's supposed to be. Actually it's our inclement weather today has got - the house painters have taken a day off because it looks like it's going to, well, actually it has been raining overnight the last two nights down here. So I don't have to have the microphone over on the other side to keep us from hearing that. The weekly garbage dump truck pickup is over. So, boy, we're just raring to go here with Episode #712 for the final day - my god, where is this year going? It's the end of April already.

**Leo:** Wow, yeah.

**Steve:** I remember it was, like, no time ago it seemed like we were talking about April Fools and how nice it was that the industry had avoided any April Fools stuff.

I want to talk about something that I've referred to only a couple times before. But a report, an extensive report on the economic power of credential stuffing attacks is what named today's podcast "Credential Stuffing Attacks." It turns out that another thing we've never really talked about in detail is this notion of this dark web. And I know Brian Krebs spends a lot of time skulking around in there, I'm sure anonymously. But sort of like the Silk Road operation years ago that got a lot of press attention, how there actually is a dark web where there is commerce going on, the exchange of value, illicit products. And, for example, famously, large numbers of presumably still valid credit cards and the required information in order to fraudulently use a credit card has long been one of the commodities being traded there.

Well, it turns out that, as a consequence of the growth of an automated marketplace, there are now like the equivalent of eBay and Amazon, where there's no human interaction necessary. There's, like, automated stores set up in this dark web. And that has made practical, along with all of these database losses of LinkedIn and Twitter and endless companies who have lost control of their databases, and now the proliferation of botnets whose agents can serve as proxies so that username and password guessing can be distributed. If one person is sitting on a website trying to guess a username and password, they'll get locked out, if the security on the site is any good, after X number of attempts. But if those guesses are coming in from hundreds of thousands of individual bots, then there's no practical way to lock them out. It looks like a large number of users. And you can't block valid users from maybe making a mistake. So it's a mess.

Anyway, that's how we're going to wrap up, but there's lots of other stuff to talk about. We want to look at more privacy fallout from our recent coverage of the problems at Facebook and Google, specifically the vault that we talked about last week. Congress stood up and took notice, or sat up and took notice. We're going to examine the uptake rate of recent Windows 10 feature releases, which has a really cool graph that we'll get to later. We now finally know the source of all those AV troubles that the Patch Tuesday April 9th of this month had, like exactly which one of those updates caused everybody trouble.

We're going to look at NIST, the U.S. organization; formal fuzzing development; consider the source of a massive and ongoing database data leak involving more than half of all American households. There's an open database which has not yet been attributed to any owner which has got a ton of information on it. And actually, the people who found it have put out a call to help them identify the owner. So we'll share that with our listeners. And maybe one of our listeners has an idea who might be behind this.

We also note that Windows Insiders are already finding that their systems won't update to the next major feature update. That's this forthcoming May 2019 feature update. You're not going to believe why. And we address the concerns of United Airlines passengers, who noticed and have been understandably upset by seatback cameras pointing at them. Finally, we have, from me, the Cranky Old Guy Tip of the Week. We touch on a bit of miscellany. And then, as I said, we're going to take a look at what many in the security industry are watching with some concern, which is this large and emerging threat of website credential stuffing attacks. So I think another great podcast for our listeners.

**Leo:** Excellent, excellent. Steve?

**Steve:** So that would be V-E-E-P-E-E-N-N-N-N-N.

**Leo:** No, VPN is spelled V-P-N. Yeah, you're right. Somebody's going to spell it no matter what.

**Steve:** Oh, yeah. So our Picture of the Week relates to one of the articles we're going to get to about midway through. The chart that we'll see later shows the time varying effect of Windows 10 feature updates over time. This one is a pie chart showing, as of this month, right now, what the distribution is. And it's probably, I mean, because the - well, I should first say that basically two thirds of the Windows 10 systems are still running the October 2018, that is, the 1803 release, have not moved...

**Leo:** It's actually even worse. It's Spring 2018.

**Steve:** Oh, that's right.

**Leo:** It's 1803, yeah.

**Steve:** As opposed to 1809.

**Leo:** Yeah.

**Steve:** Yes. Right, right, right, Spring of 2018, 1803 - have not yet moved to the problematical, which is the most recent, October 1809. And of course we've been talking extensively about what a fumble that was. And the reason I'm not sure this is really at this point very representative is it has only relatively recently been made more broadly available. And what's interesting, and we'll be discussing this a little bit later, is that it's not exactly clear why, since it is now available, and Windows 10 systems will be updating people, why it's stuck at about a third.

And what's interesting is that we're, what, a few weeks away from its replacement. So there's been some conjecture in the industry that many people, and maybe enterprises, may just be thinking, you know, let's just skip 1809 altogether and wait to see how the next one fares. So it'll be interesting to see what - and what is that, 1904, 19 oh something. I thought it was 1903, but...

**Leo:** Well, it's been 19H1 for the longest time, but now I think it's 1903 or maybe 1904, yeah.

**Steve:** Yeah, yeah. Okay. So we'll be getting to that in a minute. As we know, we talked about it last week, Facebook was hit with the heavy fine from Russia of $46.53. Wow.

**Leo:** That's like one nanosecond interest on their cash.

**Steve:** Yeah. However, the news came as a result of last Wednesday's quarterly report to the shareholders, that they're expecting that they may be hit with a little bit more substantial fine from the U.S. Federal Trade Commission, maybe as much as $5 billion, related not even to the most recent problems, but several years back to the whole

Cambridge Analytica data disclosure. So what appears to be happening is that, as has to be done when you're a publicly traded company, you generally want to paint a very sober picture of your publicly traded company's expectations so that you're not hit with a bunch of shareholder lawsuits. I mean, that happens anyway, but you certainly want to minimize it and establish the best position you could.

In their written release - which corresponds to this set-aside of $5 billion, which they said was "in connection with the FTC's investigation of its user data practices." In their written release they said: "In the first quarter of 2019" - I love the wording of this - "we reasonably estimated a probable loss and recorded an accrual of $3 billion in connection with the inquiry of the FTC into our platform and user data practices, which accrual is included in accrued expenses and other current liabilities on our condensed consolidated balance sheet. We estimate that the range of loss in this matter is $3 billion to $5 billion. The matter remains unresolved, and there can be no assurance as to the timing or the terms of any final outcome."

So anyway, as we know, a year ago, last March in 2018, the FTC announced that it was launching an investigation into Facebook's data privacy practices, which were triggered by the Cambridge Analytica revelations. The FTC is specifically investigating whether Facebook has violated a consent decree - I mean, so this is the problem is in 2011 Facebook agreed that they essentially would not do this. And so now the question is, having made that agreement eight years ago, did they violate that? And the consent decree that Facebook agreed to said that they would agree to receive explicit permission from users before sharing their data with third parties.

And then Mark Zuckerberg, in a conference call last Wednesday with the shareholders, said that he plans to forge ahead with what he calls, and we talked about this also last week when I said they ought to just, you know, scrap it, I mean, it's just a disaster. They're discovering logs? Like, oh, look. I mean, if we're to believe this, there's no way to paint this so that it seems good, they're discovering logs of millions of additional usernames and passwords that they didn't know they had. Anyway, he said he wants to create a "privacy-focused social network." He did not refer to the FTC fine or the company's even more recent data handling problems. However, the company's CFO, I guess you pronounce his name Wehner, W-E-H-N-E-R? He commented that: "We anticipate ad targeting-related headwinds," Leo.

**Leo:** [Makes wind sounds]

**Steve:** Isn't that a shame? They're having to tack against the headwinds. "We anticipate ad targeting-related headwinds will be more pronounced in the second half of 2019." Wow. Ad targeting-related headwinds. Anyway, he noted that he expects new regulations to affect the company's business going forward. Uh-huh, meaning they may not be able to play as fast and loose as they have before. So anyway, I just wanted to note that the world is paying attention to these things, and the U.S. government is taking this a little more seriously than Russia has so far.

And also last week we talked about Sensorvault, which is the geofencing concept where an article, I think it was The New York Times' revelation that Google was honoring subpoenas for geofenced data turning over anonymized exact location records for some period of time to law enforcement, who would then study them based on the facts of the case that they were trying to solve; return to Google with a much narrowed list of devices by anonymized ID and then get expanded coverage; look at those, whittle it down further, and then finally get the actual identities of the people.

Well, this generated all kinds of concern and resulted in a letter written by representatives in Congress, the Committee on Energy and Commerce that oversees this kind of stuff. I'm going to just share this with our listeners because it was addressed to Mr. Sundar Pichai, CEO of Google: "Dear Mr. Pichai: We are writing in response to concerning reports about a massive database of precise location information on hundreds of millions of consumers known inside Google as 'Sensorvault.' According to recent reports, Google tracks and stores precise location information on a huge volume of consumers, including practically every consumer with an Android mobile device, in some cases storing information dating back to 2009. The potential ramifications for consumer privacy are far reaching and concerning when examining the purposes for the Sensorvault database and how precise location information could be shared with third parties.

"First, according to the reports, Google collects precise location information in numerous ways, including from the location history function on Android Phones, Google searches, and Google apps that have location enabled. Second, precise location information is reportedly collected even when people are not making calls or using apps, which enabled Google to track the 'whole pattern of life' of an individual. Finally, Google reportedly never destroys any of the precise location information it captures in the Sensorvault database and has therefore compiled an extraordinarily detailed picture of the movements and whereabouts of a vast number of consumers stretching back more than a decade.

"As part of the Committee's ongoing commitment to protect the privacy of the American people and to understand the benefits and risks of various data collection and use practices, we would like to know the purposes for which Google maintains the Sensorvault database and the extent to which Google shares precise location information from this database with third parties. To that end, please provide answers to the following 10 questions by May 7, 2019." So what's that, like next Monday or something.

"1. What information does Google store in the Sensorvault database, and for what purposes does Google use this information? Please describe each use in detail. If the types of information in the database or the purposes for which such information is used have changed over time, explain any such changes in chronological order.

"2. Please describe which affiliates and subsidiaries of Alphabet have access to or use the data or analytics derived from the data in the Sensorvault database.

"3. Does Google maintain other databases of precise location information? If so, how does the Sensorvault database differ from other such databases, and how is the data from such database used?

"4. Who is able to access the information in the Sensorvault database? Include in your response both the number of Google employees with access to the Sensorvault database and the roles and responsibilities of any persons with access.

"5. What are the sources from which Google collects the information maintained in the Sensorvault database and any other database identified in response to question 3? Specifically, describe any Google services, mobile applications, devices, or any other means through which Google obtains the information." Then we have two parts of 5, "a" and "b." "If consumers are required to opt-in to the collection of precise location information, describe with specificity how consumers opt-in to the service, and the notice provided to such consumers about the purposes for which Google collects the information. If any such opt-in has changed over time, describe those changes." And, b, "If consumers may opt out of the collection of precise location information, describe with specificity how consumers opt out and the notice provided to such consumers about the

purposes for which Google collects the information. If any opt-out has changed over time, describe those changes.

"6. To the extent that a consumer has requested that precise location data not be shared with Google, through opt-outs or other mechanisms, do Android phones continue to collect precise location data on the device or store location information on the device? If so, to the extent that the consumer subsequently allows location data to then be shared with Google, is that formerly stored precise location information transmitted to Google? Under what other circumstances would the device continue to transmit precise location information to Google when a consumer has requested that precise location information not be shared with Google?

"7. How accurate is the precise location information stored in the Sensorvault database? Include in your response both the accuracy of the precise location information and the accuracy of attributing such information to a single individual.

"8. What controls, if any, does Google provide to consumers to limit or revoke Google's access to the information stored in the Sensorvault database? Does Google provide consumers a means to delete data stored in Sensorvault? If so, describe in detail how consumers may do so.

"9. What is Google's retention policy with respect to precise location information stored in the Sensorvault database? To what purposes does Google maintain precise location information going back to 2009?"

And, finally: "10. Does Google share, sell, license, or otherwise disclose precise location information, including de-identified data, from the Sensorvault database with any third parties other than law enforcement? If so, identify the types of businesses receiving such information and the purpose for disclosing any such information. If the data is de-identified, provide a description of how it is de-identified."

And they said, finally, wrapping up this letter: "In addition to providing the written responses requested, please make arrangements to provide Committee staff with a briefing on these topics to occur no later than May 10, 2019. Thank you for your attention. If you have any questions, please contact" and then blah blah blah, a couple phone numbers, and signed by four people on the Committee.

So it looks like Google is going to be explaining to Congress exactly what this is, why they have it, how they use it, and how that's changed over time, and who they make it available to, and under what terms and conditions. So I think we're going to be finding out much more about this in the near term, and that's probably all for the good because this is apparently - I didn't realize it went back 10 years. But it's every movement that Android users and now - and they didn't mention it, but we know that Google apps on iOS are doing the same thing. So it's not just Android devices. And as you said, Leo, it's a little distressing that Google has this and is making it available.

Okay. So the chart that I have, and it's worth displaying it on the next page, Leo, it's the bottom of page 4 of the show notes. This is an interesting chart which it took me a while to visually parse what I was looking at there. What's interesting is that each of those waves is the onset at the bottom, and then the slope is the rate of adoption of successive feature updates to Windows 10. And what's significant is that, for example, what was the Fall 2018, which was the one that now has about two-thirds adoption, it was almost a straight line upwards. I mean, everybody just jumped on it and said yahoo, and up they climbed up that curve. Compared to the curve on the far right of this chart, which demonstrates a much different adoption rate. There's a tiny little upwards tick. Then it stops.

Well, we know why, because Microsoft realized, oh, crap, and they just shut it down completely. And so it goes horizontal with no additional adoption. You remember they took it back, essentially; no additional adoption for a couple months. And then, as they began to make it available, we see it beginning to sort of creep up again, not vertical by any means, at a relatively slow pace. And it actually looks like it's leveling off there toward the end of the chart. It's like it slowed down again.

So it turns out that there's a Lithuanian-based company called AdDuplex which calls itself the largest cross-promotion network for Windows Store apps and games which empowers developers and publishers to promote their apps for free by helping each other. So sort of a third-party add-on tool. They have an SDK that is in about 5,000 Windows Store apps which gives them coverage of sort of a semi-random sampling of more than 100,000 Windows 10 machines. I looked for any other, after seeing this, I looked for any other similar representation, as they have with this chart, and couldn't find any. What I did find was that they are also updating this every month in order to create this tracking database.

So it's going to be interesting, I mean, this is sort of what we would expect to see on this most recent October 2018 feature update which has been so painful. We don't know, I mean, although - well, for example, all of my non-LTSC machines, that is, that are not on the Long Term Service Channel that I've mentioned recently which is so nice, except that it is a problem because, I mean, I've run across a problem which is that the Windows Store app is not in the LTSC, and I would like to have Edge, which requires the Windows Store app in order to install it. Probably I could use the pre-release version of Edge, but I haven't gone that far. And certainly I could use Chrome or Firefox. So maybe somebody who wants to use the LTSC would just use Chrome or Firefox and not have Edge. I need to have it as a developer. So, I mean, in that sense the LTSC is almost a little too lean. But it is really beautifully feature stripped.

But my point was that I have a number of non-LTSC standard Windows 10 builds, and they're all at 1809 across the board. So whatever it is about my systems, they've immediately jumped to this latest one. Which sort of begs the question, why are we not seeing, now that Microsoft has made it available, why are we not seeing its adoption take off? And in doing some research out in the industry, I saw some speculation that users are gun shy now of 1809, with it having had such a rough start. Except, okay, admittedly our listeners, and certainly listeners of Windows Weekly and other users of Windows 10 who are very much in the loop, they would be aware of this. But most Windows 10 users are not going to be. They're just wanting whatever Microsoft says they need, and they just click "okay," and off they go.

So I think it is probably to what degree enterprise has had Windows 10 forced down their throat. Enterprise is probably saying, whoa. And there has been speculation that I saw suggesting that maybe 1809 is just going to get skipped over completely, that there's just not going to be any adoption of it because we're looking at the next one being available so soon. We'll just wait and see how the next one takes off. So anyway, just sort of some interesting dynamics of the shape of these curves. Previously they were all very rapid. But this one has given everybody second thoughts.

And speaking of trouble being caused by Windows Update, BleepingComputer's Lawrence Abrams did some creative sleuthing, and he appears to be the first person to figure out and to clearly disclose exactly which of the April 2019, that is, this last month's updates, was responsible for so many of the enterprise AV systems having trouble. As we know, in eight separate support articles for the April Windows Update, Microsoft explained that, if users have enterprise antivirus software from Avast, Avira, McAfee, or Sophos installed - and by the way, McAfee was recently added to that list - Windows may become, as in Microsoft's words, "unresponsive upon restart after installing."

Lawrence had observed that both McAfee and Avast had made passing references to CSRSS being related to the problems. CSRSS, people who are Windows savvy have probably seen it sitting in their list of processes running. It's the client-server runtime subsystem - that's what CSRSS stands for, Client-Server Runtime Sub-System - within Windows. So Lawrence looked through the list of updates released on April 9th and noticed that there was a security update released. It was CVE-2019-0735, which was a Windows CSRSS elevation-of-privilege vulnerability. The security update stated that an elevation-of-privilege vulnerability exists when the Windows client-server runtime subsystem fails to properly handle objects in memory. An attacker who successfully exploited this vulnerability could run arbitrary code. An attacker could then install programs; view, change, or delete data; or create new accounts with full user rights.

Microsoft wrote: "To exploit this vulnerability, an attacker would first have to log onto the system. An attacker could then run a specially crafted application to take control of an affected system. The update" - this update, Microsoft's - "addresses the vulnerability by correcting how the Windows CSRSS handles objects in memory." So that appears to have been the culprit that caused those third-party AV tools so much trouble after the April 9th updates. And the good news is that the vulnerability was privately disclosed to Microsoft and was not known to be actively exploited in the wild.

So I guess, had users known at the time that that one was the one causing the problems, that one could have been backed out until all of the AV guys caught up. They all have now. There were emergency updates that were pushed. I remember I saw in the notice from Avast they said restart your system and don't even log in. Just wait 15 minutes because Avast would come up, and then it would perform an emergency update of itself so that, if you gave it 15 minutes to get all that done, then you logged in, you would be okay. And of course that was only the case after they had figured out what this update that Microsoft made had broken and then were able to engineer themselves around it.

So sort of a mess, but we're beginning to see, I mean, as we've talked about this often, in order to perform the deep sorts of checking that antivirus now has to perform, because Windows has not provided, because Microsoft has not provided the kind of official hooks into the kernel that the AV companies want, the AV companies are having to reverse engineer these pieces of Windows and install their own deep hooks into the software, which inherently makes them prone to this kind of problem. Microsoft certainly has the right to change their code when they need to, especially when they're fixing an identified problem that could allow hackers to run code in the kernel. So it's a mess, and there's really no good solution for it at this point.

The NIST, the U.S. National Institute of Standards and Technology, has gone fuzzing. After more than, it turns out, 20 years of steady improvement, the U.S. NIST believes that it has reached an important milestone in the development of a technology it named "Combinatorial Coverage Measurement." So Combinatorial Coverage Measurement is the official name for what hackers commonly call "fuzzing," which as we've discussed often is just throwing a bunch of stuff at something - a web app, ports that are open on a server, whatever - and seeing if you can crash it because, as we know, a crash is always the preamble to looking then in detail at what it was that you threw at what you threw it at, see exactly what happened, and then examine whether that can be leveraged into a much, by really understanding the nature of the crash, that's where these hacks begin. You then turn that into an exploit, and off you go.

So this Combinatorial Coverage Measurement (CCM) is part of a larger research toolkit which the NIST calls "Automated Combinatorial Testing for Software" (ACTS). So CCM is an algorithmic approach used to test software for interactions between input variables that might case unexpected failures. In the NIST's coverage of this, they said: "This is potentially useful for testing and hardening complex and important mission-critical systems such as aircraft, cars, and power plants where problems can be life-threatening.

Many complex systems operating in real-time must continually receive inputs from large arrays of sensors which might generate unexpected conflicts the software cannot resolve. The systems' designers work to identify and counteract these problems by modeling as many interactions as they can before the software is used in the real world. But what if something that could happen is missed?" They say: "This is where ACTS and CCM come in."

And of course what happens is, as the system grows, we've seen how this exponentiates, in the same way that adding each binary bit to a number doubles the number of possible states in a system. In other words, the length of a binary number grows linearly, while the possible combinations grows exponentially. So the problem becomes one of modeling enough interactions from enough variables to somehow spot all the possible combinations that might lead to an issue. And the problem is, when the system gets sufficiently complex, while it would be nice to try them all, you just can't. There are just too many possible combinations.

So over the course, they say, of two decades of work, their solution to this problem that they've been working on has been improving. They said as the result of a collaboration with the University of Texas, Austria's SBA Research, which describes itself as the research center for information security in Austria, and Adobe, who is one of several large companies using the toolkit, NIST now thinks that the 2019 version of this CCM, the combinatorial solution, has made a significant leap forward. With the help of a new algorithm developed by this Austrian SBA Research, NIST's tool has jumped from being able to model a few hundred variables to up to 2,000 from five-way combinations of inputs. This means that it can now be practically applied to much more complex systems than previously.

So NIST is making this available to partners that it's working with in situations that call for, essentially, sort of more formalized academic rooted combinatorial testing, which we know as "fuzzing." And in fact these systems really are far more complex. Fuzzing isn't dealing with 2,000 inputs. Fuzzing is dealing with a few inputs. And that's enough to typically bring our systems down to their knees. NIST is looking at, for example, all the sensors on an aircraft and verifying that, in every situation that they can find, they're able to root out a problem that could cause the system trouble before it ends up being life-threatening. So it's nice to see that we're having this kind of research at the academic level, basically an outgrowth for what the hackers have been doing sort of on an ad hoc basis until now.

So an organization, vpnMentor, has reported an unknown data breach which is exposing 80, eight zero, million U.S. households. It's hosted by a Microsoft cloud server, exposing 24GB of data.

**Leo:** Holy cow.

**Steve:** I know.

**Leo:** That is a ton of data.

**Steve:** 24GB is exposed, which includes the number of people living in each household with their full names, their marital status, their income bracket, their age, and more. It's got the full street address, cities, counties, states, and zip codes; the exact longitude and latitude; their full names (first, last, middle initial); age; date of birth; title; gender; marital status; income; homeowner status; and dwelling type.

**Leo:** See, all that information is very interesting because you log into a dating site, they're not going to get that information.

**Steve:** Exactly.

**Leo:** So what could it be?

**Steve:** Well, and this is the question. And the reason this kind of hit me is, in scanning through interesting news of the past week for this podcast, I'm continually seeing articles about data breaches. I mean, they're just kind of boring, normally. We know that they're happening all the time. It just gets to be another form of Internet background radiation. It's like, yeah, yeah, yeah, another company has lost its data. For example, I just thought, okay, like what?

And so there are three recent headlines that I skipped over: medical information of almost 150,000 rehab patients exposed; 540 million Facebook records leaked by public Amazon S3 buckets, it's like, yeah, okay, fine; 257,000 legal documents leaked by unprotected Elasticsearch server. And it's like, these are things I've encountered in the last few weeks. And it's like, yeah. I mean, what is there to say about it? But here we've got more than half of U.S. households with, as you said, Leo, a really interesting set of specific data. And normally it's possible to identify the owner of the data and then notify them that they've left their fly open. But in this case there's no attribution to this data. We know what it is. We don't know whose it is.

So the research team at vpnMentor has been unable to figure out who all this data might belong to. And they've essentially created an open call for help. They said: "The research team is currently undertaking a huge web mapping project." And it's going to be sort of interesting to see what else they find. "They use port scanning to examine known IP blocks. This reveals," they wrote, "open holes in web systems which they then examine for weaknesses and data leaks. Usually the researchers suspect where the leak is coming from. They can then examine the database to confirm its identity." Then they wrote: "We then reach out to the database's owners to report the leak and, where possible, alert the people affected. This helps build a safer and more protected Internet," yada yada.

"Although we investigated the database online, we didn't download it." They wrote: "Our researchers felt that downloading it would be an ethical breach as they would then illegally own personally identifiable datasets without people's consent." On the other hand, it's flapping in the breeze. But anyway, they said: "This time it's different. The database that the team discovered includes identifying information for more than 80 million households across the United States. As most households include more than one resident, the database could directly impact hundreds of millions of individuals.

"vpnMentor is calling on the public to help identify the database and close the leak. Unlike previous leaks we've discovered, this time we have no idea who this database belongs to. It's hosted on a cloud server, which means the IP address associated with it is not necessarily connected to its owner. The data includes uniform entries for more than 80 million households, making it almost impossible to narrow down. The only clue we found lay in people's ages. Despite searching thousands of entries," they write, "we could not find anyone listed under the age of 40."

**Leo:** Oh, interesting.

**Steve:** Yeah. They said: "Interestingly, a value for people's income is given. However, we don't know if it's a code for an internal ranking system, a tax bracket, or an actual amount. This made us suspect that the database is owned by an insurance, healthcare, or mortgage company." And mortgage is what I'm thinking, like something related to housing, because they talk about, what was it, they had a dwelling type and homeowner status. And they said: "However, information one may expect to find in a database owned by brokers or banks is missing. For example, there are no policy or account numbers, Social Security numbers, or payment types."

So they said: "Help us identify this database. We want to contact the database's owners." And I wonder why they just can't talk to Microsoft because Microsoft must know, be able to equate an IP in their cloud to its owner. They said: "We want to contact this database's owners and let them know that their data logs are exposing millions of households. Help us solve this riddle. What service is used by 80 million homes across the U.S., but only the U.S., and only by people over 40? What service would collect your homeowner status and dwelling type but not your Social Security number? And what service records that you're married, but not how many children you have? If you can help us identify this database or know who owns it, please contact us at info@vpnmentor.com. The 80 million families listed here deserve privacy, and we need your help to protect it."

And notice they have not indicated what the IP is or where it is. They don't want to expose this because unscrupulous types will immediately suck out that 24GB and put it to malicious use. So they've found it. And if anybody has an idea, info@vpnmentor.com. So an interesting call for - and this is the world we're in today, as I find myself often saying on this podcast. It's like, wow.

And speaking of the forthcoming May 2019 Windows 10 feature update, Leo, you're not going to believe this one. Microsoft is explaining to those on the Windows Insider program, who are trying to apply that May 2019 feature update ahead of its official release, as they are able to, that the reason for the "This PC can't be upgraded to Windows 10" is because it has a USB device or SD card attached. You heard that correctly. I'm not kidding. Microsoft's exact text reads: "If you are part of the Windows Insider Program, and you are trying to upgrade to the May 2019 Update for Windows 10 (Windows 10 version 1903), you may experience an upgrade block and receive the following message: 'This PC can't be upgraded to Windows 10.' An external USB device or SD memory card that is attached to the computer could cause inappropriate drive reassignment on Windows 10-based computers during the installation of the May 2019 update. For this reason, these computers are currently blocked from receiving the May 2019 Update.

"This generates the error message that is mentioned in the 'Symptoms' section if the upgrade is tried again on an affected computer." And they said: "Example: An upgrade to the May 2019 Update is tried on a computer that has a thumb drive inserted into a USB port." Yeah. Mine does. "Before the upgrade, the device would have been mounted in the system as drive G, based on the existing drive configuration. However, after the upgrade, the device is reassigned a different drive letter. For example, the drive is reassigned as drive H."

And then they said: "Note: The drive reassignment is not only limited to removable drives. Internal hard drives can also be affected." Whoops. Anyway, this will presumably be fixed before the big rollout of next month's release. Otherwise, the adoption of this release may not proceed any faster than the last one. So this appears - Windows 10 just in general appears to be an increasing challenge for Microsoft. And Leo, I've heard Paul and Mary Jo just sort of, I mean, at this point Paul has his head in his hands, saying, "What is going on up there?"

**Leo:** Well, I mean, to be fair, I remember this is a problem on mobile devices, too, is removable USB storage. It often confuses them. In fact, Windows Phone for a long time wouldn't work with SD cards, micro SD cards. Apple doesn't even give you the option. And on Android they go back and forth about how to treat SD cards. So it's not unreasonable to say, because we can't be sure that that device will be attached, you don't want to have it attached during installation. You can plug it in afterwards and use it. You just can't have it attached during installation because it's removable. I don't know if that's - I'm sure they could fix it because we've never seen it before. But I can understand why that might be an issue. And as for drive letters, that's just a convention anyway, what the drive letter is.

**Steve:** Well, except that it has never been an issue before. We've had all of the previous feature updates.

**Leo:** Right, right. It has been an issue in other operating systems, though. It's not unusual to hear about problems with removable storage devices because, if you're going to use it, then you have to know that it's going to stay there. So, I mean, admittedly, it's not a good thing, and they'll fix it.

**Steve:** Right.

**Leo:** But I can understand how it could happen, I guess is what I'm saying.

**Steve:** Yeah. So United Airlines made a little mistake. They purchased seatback displays that had embedded cameras. Passengers became quite upset at the idea of a camera staring at them throughout their flight. And this is exactly where and when and, I mean, I hope you want to have a physical mechanical sliding door, a shutter to shut the camera, and have it on the seatback, and have it labeled with clear instructions, and maybe even train your maintenance staff that goes through between flights and puts all the seatbelts in the correct place to just get into the habit of closing all of those shutters.

Anyway, so what happened was it turns out that they purchased a bunch of these, didn't think twice about it, installed them in seatbacks, and there was an immediate hue and cry among their passengers, wondering why there was a camera staring at them, which did not provide any means for it to be closed. They were quick to respond because of course United Airlines has had some PR problems lately with that passenger who was forcibly removed, and then unfortunately there was I guess a puppy that was forced to be stored in the overhead, and it did not survive the flight, unfortunately, or tragically. So United Airlines has apologized for the confusion, said that these cameras were intended to be deployed in some future mode for teleconferencing.

But our takeaway is, and we've talked about this often, there is no good way to have a camera aimed at someone where you don't have a mechanical shutter. It's nice to have a little light next to it that shows when it's on, but that's just not as good as a mechanical shutter. And years ago we talked about cutting off the sticky part of a Post-it note and just putting it over the camera. And in fact you often now see pictures of people's laptops that have a little chunk of a Post-it note over their camera because it's just simple and easy. And there are, you know you can buy aftermarket low-profile shutters, if you are in the routine of actually using your camera, and you want to be able to open and close it because there have also been lots of examples of bad guys getting into people's computers and turning on their cameras. We've talked about them years ago.

So this week I offer what I call my "Cranky Old Guy Tip of the Week." This just happened this morning, when doing some web browsing to a site where I was pursuing something to flesh out the coverage of one of these things. I just got tired of Firefox popping up that question about whether I wanted to allow this website to send me notices. I cannot imagine a single instance where I want a website to have the privilege to put notices down in the notice area of my OS. Just not a single instance where I want that.

So I thought, you know, I'll bet Firefox, I'm hoping Firefox has a way just to turn this off, this whole web notifications thing. So I opened another tab, put in about:config, and then in the search bar I put in "notifications." And sure enough, in fact, if you put in, and I would recommend for any of our listeners who feel similarly, if you type "webnotifications" as all one word, you get four hits. The first one is dom, D-O-M, as in document object model, dom.webnotifications.enabled. It defaults to true. Double-click it, and flips to false.

And what's interesting is that I also went to - I did some googling because I was curious about the API. And I chuckled because Mozilla's page about the web notifications API itself pops up a web notifications prompt. So you can verify, I have the link in the show notes for anyone who's interested, you can verify that you have turned them off by refreshing, for example, you could go to this notifications API page first. And sure enough, Firefox will pop up a little "Do you want to allow this website to send you notifications?" And you've got Allow, Not Now, or then in a dropdown you can say Never. I don't even want to be asked. And so this turns off the ask. And once that's been set to webnotifications.enabled to false, then sure enough, you're no longer asked, which is a great relief.

And Chrome makes it possible, too. In the case of Chrome, under settings, go to Advanced, which extends the settings page to reveal a whole bunch more stuff. And the first section there is Privacy and Security. Under Privacy and Security, click on Site Settings, and then five entries down you'll find Notifications. It's a little less clear how to do this there. But the first option at the top, there's a switch that says "Ask before sending," and it says, "(recommended)." Anyway, when you toggle that little slider switch to off, the text changes to blocked, and it has the same effect.

So if all of these, I mean, to me this feels like what is going to be one of the most abused things that we can imagine. I mean, as bad as unsolicited pop-ups. And we're seeing those more and more now, too. But the good news is, for both of these two browsers, it's pretty simple to turn them off and just not even be asked anymore because I just can't imagine an instance where I would want to do that.

And a couple bits of miscellany for our listeners. Leo, for the last year and a half I have been spending my time with a fabulous woman in a location that has essentially zero bars of Verizon.

**Leo:** You must really like her.

**Steve:** Oh, my god. We're both Verizon users. And this place, there's an area known as Turtle Rock. And when I mentioned it to the Verizon employees, they all go, oh, my god. Apparently the associations up there, the housing associations refuse to permit cell towers. You know, it's NIMBY.

**Leo:** Not in my backyard, yup.

**Steve:** Exactly. And as a consequence it just has this reputation, like, forever, that it's just like the cellular graveyard. It's the cellular dead zone. And what occurred to me is that, okay, 10 years ago maybe that was feasible because people were still heavily reliant on landlines. But that's going away now, and the world is going to cellular. And it's becoming much less practical. Well, I knew this was a problem when we decided this is where we wanted to settle down. And so I thought, okay, fine. We've got Verizon WiFi; right? We can do WiFi calling.

Well, I don't know what the problem is. It works on my iPhone 6. My iPhone 10 refuses to have it work, as does Lorrie's iPhone 8. It just - I managed to get it on. Then it went away. And it's like it's always on on my iPhone 6, but not - and of course the idea is that, if you have that, then you're able to make your calls over your WiFi and your cable modem bandwidth or whatever. It may be that the reception is so poor on these phones that it doesn't even know you're in a Verizon zone enough to hand off to WiFi. Anyway, this is my long-winded introduction to the Samsung 4G...

**Leo:** I was going to say femtocells, yup.

**Steve:** Oh, my lord, is it a win.

**Leo:** Yeah. We use a femtocell here. We don't have - for some reason T-Mobile doesn't work here. Yeah, it's great. It uses your Internet. It's like your own cell tower.

**Steve:** It is life-changing. So for any of our listeners, I mean, it's like, why didn't I do this a year and a half ago? Lorrie spends a lot of time there on the phone, and it's just been a constant headache for her. And I'm feeling, like, negligent that I allowed this to go for 18 months without fixing it. Because finally I said "enough" last Friday. And so this is where I went to the Verizon store, purchased one of these things. There's no service fee.

**Leo:** Actually, it's always a good idea to say first, "I can't get your service in my house, so I'm quitting," because they will often give you one for free. I mean, they're not hugely expensive, but still. Because they don't want to lose you.

**Steve:** Interesting.

**Leo:** We got our T-Mobile femtocell for free.

**Steve:** Interesting. Didn't even - I'm too much of a...

**Leo:** 250 bucks, big deal.

**Steve:** It was, it was 250 bucks, one time. And it's funny, too, because we are planning eventually to move somewhere more permanent, and Lorrie's been worried. She says, "Well, what about..."

**Leo:** Wait a minute, Steve. What? You and Lorrie together?

**Steve:** Well, we're getting along really well.

**Leo:** Oh, my god. You've not left your apartment in 23 years. Wow.

**Steve:** So anyway, so she's worried, like this has been such a problem for her that it's like, what if this is a problem when we move? So I've explained that we will absolutely never again have this problem because basically we get to take this with us wherever we go.

**Leo:** Well, now, one little caveat. Before you move anywhere, check to make sure they've got good high-speed Internet.

**Steve:** Yeah, yeah.

**Leo:** Because you do have to have that to use this. I wouldn't move anywhere that didn't have good high-speed Internet. My god.

**Steve:** No, no. We couldn't possible live without that.

**Leo:** You couldn't do the show, Steve, more importantly.

**Steve:** No. Could not, well, I'll probably still keep this place. This is like my man cave, and then - because it's nice...

**Leo:** I have a relief, yeah.

**Steve:** ...to go somewhere to make a mess.

**Leo:** You need your man cave.

**Steve:** That's right. Anyway, I just wanted to share this fabulous experience. I was a little annoyed, for what it's worth, that you cannot do a whitelist. I wanted to whitelist our phones because I didn't want to be giving absent Verizon 4G LTE coverage to my neighborhood. It's like, they can get their own. And I didn't want them using my bandwidth. It turns out this is on purpose. Basically I'm subsidizing...

**Leo:** You're doing them a favor, yeah.

**Steve:** Yes, I'm subsidizing Verizon's lack of 4G LTE coverage throughout my neighborhood. Well, actually it turns out...

**Leo:** Turtle Rock, we now have a cell tower. It's in Steve's house.

**Steve:** Turns out it's not a problem because, if I go out into the street, it's already, you know, it falls off.

**Leo:** It's only 50 feet. It doesn't go very far, yeah.

**Steve:** Yes. So it is, it's 50-foot radius, 7,500 square feet. So it beautifully, it's like all bars anywhere in the house now. And Lorrie is just dancing. So anyway, I wanted to share the fact that...

**Leo:** Every cell company generally offers something like that. Technically it's a femtocell because femto is tiny.

**Steve:** Yes. Femto is part of the default password. I was very impressed with the security, too. You log in with a fixed-string prefix, and then the last four digits or the last four characters of the MAC address printed on the bottom. So there's no default password. You've got to have physical access to the device in order to know what the MAC address is, in order to log in. And then of course, once you do that, then you change it to whatever you want.

**Leo:** Oh, okay, yeah. You can get somebody's MAC address if it's on the air; right?

**Steve:** Yup, yup.

**Leo:** Yeah, femto is 10 to the negative 15th. So it's a really small cell tower. But they don't call them femtocells. AT&T has another name. Every company has another name. But they all offer them because they don't want to lose you as a customer. If you can't use it in your house, what good is it?

**Steve:** Well, it works. I just wanted to tell everybody, this little thing has, like, changed our life.

**Leo:** Nice.

**Steve:** So, yay. Let's see. From my blog on the 28th, which was last Sunday, Derek Hamilton posted, he said: "Hi, Steve. You've mentioned many times on Security Now! about the big release event you are going to do on SN. Are we counting weeks, months or longer here?" He says: "I can't wait to watch. Thanks." And so I just thought I would take that opportunity to note that we are at Release Candidate 3 of the SQRL app for Windows. And it's way ahead of the other clients just because it had a much longer head start. But I'm like, as far as I know this is done. And I want to play a little bit with the text. But we're, like, there. And then you and I, Leo, need to figure out when we can come up - and you'll get to meet Lorrie - and put together a little bit of a show of what is SQRL and how it works and so forth.

**Leo:** A ketogenic dinner is on us. Actually, no. You eat pasta now.

**Steve:** I do.

**Leo:** I'll eat the keto. You eat the pasta.

**Steve:** I've been rescued from ketosis. And speaking of SQRL, I wanted to also mention that last week's mess with the just-minted EV code signing certificate appears to be history. Windows Defender was first to recognize that SQRL was not malicious and to start believing my newly minted EV code signing cert. And then the Windows SmartScreen folks responded quickly to my notifying them that this was from and by a legitimate developer. And so all those problems went by the wayside as quickly as I could ask. So I was very impressed by that.

So, website credential stuffing attacks.

**Leo:** Mmm, stuffing.

**Steve:** Leo. I think you need to have your lunch, Leo.

**Leo:** I haven't had lunch, you got it.

**Steve:** So credential stuffing attacks - it's sort of an odd term, but it's what the industry has adopted - are Internet-based mass username and password logon guessing at scale. So this requires not only the technology to implement the attacks, which is now widely available, with so many hundreds of thousands of Internet-connected computers, routers, IoT devices, whatever, available to compromise. But to be profitable these attacks also require the presence of a sufficiently mature marketplace into which it turns out a frighteningly large number of discovered working logon credentials can be resold.

So I know we've never really talked much about the so-called "dark web," but it's a real place. This is, as I mentioned at the top of the show, this is where Silk Road was conducting its commerce. And the point is that this is not a hacker trying to get into some person's account by guessing a known person's unknown password, or knowing something about them or their lives and so forth. This is a shotgun scattershot numbers-based attack where the millions now we're at of paired usernames and passwords that have leaked out onto the Internet and have been sucked up into databases, those are being broadly used, blindly, against websites with the hope that they get in. And if they get in, they log out because that particular individual, that is, this is an automated attack. They have no interest in what's there. What they want to do is they want to put that up for sale.

So the site is Recorded Future. We've mentioned them from to time in the past. They're a well-known security research group. They titled this, I think very appropriately, "The Economy of Credential Stuffing Attacks" because it's the economics of this which it turns out has made this very profitable. And unfortunately, as we have found, when it's possible to make money at something, you get more of it rather than - like once upon a time viruses were just sort of, as we've often talked back in the beginning of the podcast

14 years ago, it's like, okay, well, the good news is they're not really doing anything malicious. They're just annoying.

Well, of course that changed. As soon as you could have a way of transferring funds, like through bitcoin, then you started getting crypto malware because you could hold somebody ransom for the data on their computer. And of course we know what has happened there. This promises to be something similar.

So these guys begin by explaining. They said: "This report covers the current threat landscape of credential stuffing attacks. It reviews the most popular tools used by cybercriminals to initiate credential stuffing, and describes some of the most popular marketplaces that sell compromised credentials. This report contains information gathered using the Recorded Future Platform, as well as additional open source, dark web, and underground forum research; and will be of most interest to analysts protecting e-commerce, telecommunications, and financial organizations from credential stuffing attacks, as well as those looking for investigative leads on threat actors performing such attacks."

And then to set up their presentation they explain: "The rapid proliferation of automated marketplaces" - and that's the key - "on the dark web, fueled by the widespread availability of support infrastructure such as account-checking software" - meaning brute forcing usernames and passwords - "email and password combo lists, and proxy service providers, has created the perfect attack landscape for the abuse of thousands of popular web services such as ecommerce, financial services, travel websites, and telecommunications companies." They say: "It's safe to assume that almost every large organization with an online retail presence has had their users exposed to credential stuffing attacks in the past few years, with some companies having upwards of millions of exposed login credentials available for purchase on the dark web at any given moment."

Then they have five main takeaways. They said: "The first widespread credential stuffing attacks were observed in late 2014, coinciding with the proliferation of automated underground marketplaces. When selling accounts, attackers offered the quick and easy monetization of compromised account credentials. Some actors who engaged in credential stuffing attacks remain active today." So this has been - they've been present for five years.

Second: "With an investment of as little as $550" - and we break this down in a minute - "criminals could expect to earn at least 20 times the profit on the sale of compromised login credentials." And actually the return actually shows it turns out to be much greater return for investment. They said: "Third, the overall supply of compromised login credentials across several large marketplaces exceeds tens of millions of accounts." I'll say that again. The overall supply, that is, for purchase of working compromised login credentials, across several large marketplaces, exceeds tens of millions of accounts.

They said: "Fourth, we have identified at least six popular variants of account-checking software," that is, basically that's the software that the bad guys run where they feed in a database that they have purchased into the software which they have purchased, which operates against a network of proxies which they rent, in order to generate hits, successful login hits which then this software collects, which they then compile and turn around and put up for sale on the dark web. And there are purchasers, not surprisingly, for this freshly farmed information. So they said: "We have identified at least six popular variants of this software used by cybercriminals." They said: "However, dozens of lesser known variants can also be found on the dark web."

And, finally, fifth: "While some companies may choose to implement multifactor authentication, which blocks the credential stuffing attack vector, organizations may not be prepared to choose security over convenience." And probably the number one

takeaway from this is that anything you do to block the feasibility of blind username and password login - oh, and CAPTCHAs don't solve the problem. All of these software frontends have capture-defeating features. It has to be something like multifactor authentication, something that removes the vulnerability from brute forcing.

So in their report they write: "Around late 2014 and the beginning of 2015, we observed the widespread adoption of new dark web business models specifically tailored to facilitate a high volume of trades in a fully automated manner. Designed to emulate legitimate retail platforms such as eBay and Amazon, these so-called 'automated shops,' as they're known, allow even low-level criminals to become vendors of stolen data, such as in this case compromised login credentials, without having to worry about maintaining their own infrastructure or marketing campaigns."

I mean, and so think about that. That's like the eBay model. If somebody's got some garage junk that they want to sell, if it weren't for eBay, how would you do that? You'd have to put up a website somewhere and somehow tell people to go to your PayPal account. And if that didn't exist, that would be a problem. So eBay has solved that problem for anyone who wants to sell some random thing. As they say, there's a buyer for everything. So the equivalent now exists recently, as of about four years ago, on the so-called "dark web."

They said by and large the adoption of account marketplaces, that is to say, marketplaces for selling login account information, was made possible primarily by the proliferation of account-checking software, which is known by the slang "checkers," used as the main tool in credential stuffing attacks. And then I'll note that we previously heard about the related marketplaces for stolen credit card information. So what we have on the dark web is an underground, effectively retail, automated marketplace for buying and selling of stolen goods - in this case large quantities of recently tested and verified usernames and passwords for specific websites, the pairing of all of those together.

Anyway, they continue to say: "Compromised account credentials were always a valuable commodity in the dark web. The number of transactions was relatively small, and they were primarily conducted either on a peer-to-peer basis or via semi-automated markets such as AlphaBay, Silk Road, and Hansa Market. In older models, buyers received their wares only after the seller manually approved the deal and delivered the purchased data. Moreover, sellers had to maintain the listings and communicate with the buyers personally. However, with the advent of automated shops, the need for manual engagement was eliminated, and the business of compromised accounts fully transitioned from peer-to-peer dealings to a much more democratized, open to everyone, marketplace.

"For a nominal 10 to 15% commission deducted from the amount of each sale, members can upload any number of validated compromised accounts which, in addition to email and password, often include data such as the accountholder's city or state of residency, transaction history, and/or account balance, all of this valuable data to fraudsters seeking to buy accounts tailored to meet their specific needs. The vendors' main focus is replenishing the stock; while all customer support, remittances, and dispute resolutions are handled by the shop's support team.

"At first, only a handful of select vendors became the primary suppliers of stolen data; but as the tradecraft was shared among members of the criminal underground, the business of stolen credentials has grown exponentially. Since regular Internet users tend to reuse the same passwords across multiple websites, threat actors quickly learned that instead of attempting to obtain access to an individual account, which may take a very long time, they should instead focus on hacking multiple random accounts, that is to say across websites, thus reducing their efforts." In other words, if they find one particular username and password that works at one site, quickly check that same pair at all the

other sites that they are checking against on the off chance that the user has reused their password, as certainly used to be the case and, among uninformed, un-security-informed users, is still the ongoing practice.

Then they said: "A combination of several elements made the hacking of various online service accounts not just effortless, but also incredibly lucrative." And this reiterates what they said: "To launch account brute forcing, also known as credential stuffing attacks, an attacker only needed" - and here are the requirements - "brute-forcing software, a database of email and password combinations, and access to a pool of proxies."

They said: "Early versions of checkers [as they're called] were made to target a single company and were sold for between $50 to $250, depending upon the tool's capabilities. These tools would attempt to log into a website using an email and password combination obtained from a random database often obtained on the dark web. If a combination worked, it would be marked as valid. If not, the software would simply pick another combination from the list and attempt to log in again.

"For valid logins, more expensive and complex checkers would also collect additional information from the compromised account, such as linking banking and payment card information, account balances, the owner's address, and even transaction history. Until this day, the ingenuity of the method truly lies in the economy of scale, allowing criminals to process hundreds of thousands of combinations in a very short period of time. Eventually, several dominant players such as Storm, Black Bullet, and Sentry MBA entered the market with more robust tools, supporting an unlimited number of custom plug-ins."

Naturally, as we've seen before, the software is only going to get better in time. So what they're saying is that it used to be that you'd get like the PayPal cracker, or the Amazon cracker, or the Twitter cracker or something. Now they are now multisite crackers that are able to span a wide number of ecommerce-related sites. And I have a picture on the page below where I am, showing a bunch of the sites that are attacked and the value that a typical attacker can expect to get per site.

Anyway, they said they supported an unlimited number of custom plug-ins, also called "configs," which essentially offered hackers the capability to target almost any company with an online retail presence. What had initially started as several hundred or several thousand compromised accounts quickly ballooned to hundreds of thousands, or even millions of accounts. I mean, this feels like science fiction, but these guys have the numbers to demonstrate it. And they've been poking around in the dark web, where this is all going on, and seen the size of the databases of compromised accounts that are on offer for sale.

They said: "Some of the most prominent account shops have tens of millions of compromised accounts for sale at any given moment." If this isn't enough to get people to change their passwords to random strings, I don't know what is. "Although the competition quickly brought the average price of a single compromised account down from over $10 to a mere $1 to $2, the overall profitability of credential stuffing attacks increased significantly through sheer volume. According to underground chatter observed over time, the average success rate for credential stuffing is anywhere between 1 to 3%. Hence, for every one million random combinations of emails and passwords, attackers can potentially compromise between 10,000 and 30,000 accounts. Moreover, the same database could then be reused over and over again to hack dozens of different websites, yielding even higher profits."

So this page, or this graph from their report on the next page shows that the so-called "checking software" would cost somebody wanted to set himself up as an attacker $150. They would then purchase an email password database containing 100,000 records for

another $150, and then would rent a proxy network, rent access to a proxy network through which to launch these attacks for $250 per week. What this then generates is - assuming that you are attacking Amazon, PayPal, eBay, Expedia, Airbnb, FedEx, Credit Karma, Online Video Service, and Xfinity. They sum up the average price available for each one of those different services' compromised accounts, and they vary.

For example, eBay you get $3.50. For Amazon, you can sell an Amazon-compromised account for two bucks on the dark web. Credit Karma brings two bucks. Xfinity brings 3.50. Airbnb $1.50. Expedia only $0.50. PayPal, you can sell a compromised PayPal account for a buck. Anyway, they look at between 1 and 3% versus 100,000 records, show the amount of profit per. And they end up with a gross profit of just shy of $20,000 for the accounts that you can generate from this attack using that investment. So they're talking about they look at the gross profit margin. And so for a direct cost of $550, your return on the dark web is $20,000.

They go on in their report to give the - they fully examine the capabilities and features of six specific tools which are being used to carry out these attacks. I won't go into each of them in detail since anyone who's interested can follow the links that I have, both to their site and to their PDF of this report at the top of this coverage in the show notes. My goal here has been to share and develop a real awareness of the existence of this very active underground market and, for bad guys, the compelling financial incentive that promises to keep it alive and growing for the foreseeable future.

And you can imagine now that what this means is that, when bad guys discover email and password databases, they're not going to immediately publish them. They may not even immediately offer them for sale. They may use these tools to exploit them themselves before - because the more that same database gets resold and reused, there will be a drop in the value as people realize that somebody got into their account. They quote, "I've been hacked, my eBay account was hacked" or whatever, and then they'll change their password, thus rendering that entry in the database no longer effective.

But still, to me, it was fascinating that, as a consequence of the creation of an automated eBay-style marketplace on the dark web, the fact that we now have botnets which can serve as proxies so that these probes and attacks to log in are coming from hundreds of thousands of different IPs, not just one IP that it would be easy for a website to blacklist after 10 failed account login attempts. You can't do that with hundreds of thousands of IPs because then you risk blacklisting legitimate users. And of course it spreads out the guesses so that you're not going to be hitting from a single IP multiple login failures. This is real. And unfortunately, it's only going to grow over time because, as I said, it's not just anymore for the lark, it's for turning these things into money. And here is now a way to turn a leaked database into cash on the dark web. Wow.

**Leo:** Yeah, it sounds like it's really almost hacking as a service.

**Steve:** Yeah.

**Leo:** It's like you can really - everything you need.

**Steve:** You buy one of these, you buy one of these, you buy one of these, and you set yourself up, exactly.

**Leo:** Yeah. It's all cloud-based these days.

**Steve:** Yup.

**Leo:** It's modern. So really you wouldn't need much of a fancy system at home to do this.

**Steve:** In your mom's basement, as they say. Yup.

**Leo:** It's the will and the knowledge, and the knowledge is out there.

**Steve:** Well, and willing to put yourself at risk because you need to protect your IP. The feds are going to come looking for you, too. So, I mean, it's not like it's free money that you're getting. You're committing a crime when you're doing this. So, yikes.

**Leo:** Yeah, although the feds in Bulgaria are maybe not as scary.

**Steve:** Yeah, well, and Russia, if Russia's hitting Facebook with a $48 fine, they're not that concerned, either.

**Leo:** They don't care, right, exactly.

**Steve:** And in fact some of these tools were translated from Russia.

**Leo:** Russia, yeah, of course.

**Steve:** Because the market is there.

**Leo:** They're perfecting it for us.

**Steve:** Wow, just amazing.

**Leo:** Yeah. It's a brisk market.

**Steve:** The world we're in today, Leo.

**Leo:** It's amazing. Steve, thank goodness we've got you to tell us all about it. So we're all set. We're all ready. This is the Security Now! podcast in a nutshell. Everything that's going on all around you in the darkness. But Steve brings light to it all every Tuesday, around about 1:30 Pacific, 4:30 Eastern, 20:30 UTC. You can tune in TWiT.tv/live and watch or listen as we do it live. If you do that, go to the chatroom, irc.twit.tv. They're fascinated by you, Steve. They're looking for pictures of Lorrie already.

**Steve:** Oh, no.

**Leo:** Might want to use a pseudonym in future. No, no, these guys are benign. They're nice. They're not the bad guys, they're the good guys in the chatroom. Once we chop it all up and make it a podcast, Steve puts it up on his site at GRC.com, plus gets Elaine to write what he said in pure English. That's an amazing feat in and of itself that she performs each and every week. And you'll find that at GRC.com, as well.

When you're there, check out SpinRite, the world's best hard drive and recovery and maintenance utility. That's Steve's bread and butter. It's the only thing he charges for. He has lots of other great stuff on the site. So browse around. You can leave him questions there, too: GRC.com/feedback. But he's also on Twitter and takes direct messages from anybody at @SGgrc. But, you know, you've got a good class of people following you, so I think that's safe, as well.

**Steve:** Yeah, true.

**Leo:** We have audio and video of the show at our website, TWiT.tv/sn. But of course, as always, we encourage you to subscribe. Whatever podcast application you use, you'll be able to find Security Now!, one of the longest running podcasts on the air still today. I was just going through my subscription list. Half the podcasts were gone. Not gone like gone, but they hadn't made a new podcast in two or three years. We make one every week. Steve gets mad if we don't make one every week. So come on by and enjoy the frothy, delicious Security Now!, now with bergamot. All right, everybody. See you next time. See you, Steve.

**Steve:** Thank you, my friend.

**Leo:** Have fun in Turtle Rock.

**Steve:** Right-o.