Password Cracking Update: The Death of "Clever"

Description: After catching up with a collection of miscellaneous and interesting security-related news, Steve and Leo take a close look at the long-term consequences of the many massive password leakages which have occurred. The upshot? Hackers are getting MUCH better at cracking passwords, and "clever" techniques can no longer be regarded as safe.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-366.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-366-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about, well, he calls it "The Death of Clever," what password crackers know about your password and how easy it is to crack. It's coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 366, recorded August 22nd, 2012: The Death of Clever.

It's time for Security Now!, the show that covers security and all that great stuff with this cat right here, the Explainer in Chief, Mr. Steven Tiberius Gibson. I gave you a middle name.

Steve Gibson: I watched a fun movie the other day, which had a sort of a surprise for me, sci-fi, of course. Well, not of course, but in this case it was: "The Puppet Masters."

Leo: I don't know that one.

Steve: Donald Sutherland and some other people who we've not really seen much of. And I watched it with Jenny. Jenny had never seen it before, and she really enjoys sci-fi as much as I do. But she was sort of thinking, well, okay, all of these plot vehicles are sort of familiar. Like, well, you know, this is from "Aliens" and this is from this and this is from that. And I said, yeah, but this was a lot - this, like, predated them all. Well, I didn't know by how much until I found my copy of the paperback, which was copyright 1951.

Leo: Wow.

Steve: And I was like, whoa, Robert Heinlein really was ahead of his time. That was just amazing.

Leo: And when was the movie made?

Steve: Oh, the movie was in the, like, late '80s, I think, or maybe early '90s. Wasn't a long time ago, although I didn't remember how hokey the cinematography was. We were sort of laughing about it because it was kind of cheesy.

Leo: You know, it's funny, it doesn't take long. I was watching "Total Recall" with Arnold Schwarzenegger, the original, and that was late '80s, something like that. And it really looks dated. It looks like a bad '70s cop show, practically. I mean, it's funny how far we've come.

Steve: Yes. This had exactly the same feel. And I was asking Jen, I said, what is it? It's like little subtle things, I guess, that have evolved in our understanding of how to make something...

Leo: And fast.

Steve: Yes.

Leo: IMDB says 1994 for that movie. So that was only 18 years ago. That's not a huge difference. But things have, you know, it's computers. You've heard about them; right? They've made a big difference in everything, including filmmaking.

Steve: There was sort of some exposure, like it was kind of overexposed, or a little too much brightness or contrast. And, oh, the other thing I noticed was that they seemed to be much more kind of in our face, as if we were looking at a smaller screen, much further away...

Leo: Interesting.

Steve: ...than we are now. Now we've got much bigger screens. They generally are closer. And of course the resolution is way up, so we don't need things really big in our face any longer.

Leo: Yup. We've gotten subtler, and we have more lens flares, that's right, chipo burrito. That's thanks to Joss Whedon, I think. Today the show is called, not "The

Puppet Masters," "The Death of Clever."

Steve: Well, yes. This is a - I guess I would call this a "crowd-sourced topic" because there has never been an article that received so much tweeting to my attention as a recent article in Ars Technica. Their cybersecurity guy, Dan Goodin, did a very nice, very comprehensive, four-page piece about sort of a snapshot on where we stand with password hacking. And obviously this is a recurring topic for us. But I decided that we really needed - well, first of all, many people were tweeting, Steve, what do you think, what do you think, is he right, blah blah blah.

Now, he didn't really draw any conclusions. We will, because this is what we do in the podcast, have some takeaways from this. But what I felt as I was reading it was that I needed to make sure our listeners understood something that I don't think I've ever made sufficiently clear. So I titled this "Password Cracking Update: The Death of Clever" because what this really drives home is that - and this is really interesting, too, because this is a consequence of the breaches that we've had. We've covered these massive breaches for the last few years.

Well, it turns out having access to more than 100 million actual in-use passwords, which is what are now available, freely downloadable over the Internet, having those actual passwords has changed the complexion of password-cracking. All of those things that we sort of "wink, wink" about doing, like changing alphabetic characters into the numbers that they resemble, or those sorts of things - and we'll talk about what those are because they've all been analyzed now. What's happened is, and this is another thing you would expect, over time there's evolution of the technology. The cracking is really getting better.

And so today's podcast, driven by so many questions that were tweeted to me about it, will update everyone. And if anybody still thinks that they're being cute with the way they're designing passwords, I hope to be able to increase their security further by putting them off of those habits because they're just not working any longer. And we've got a bunch of interesting news, as well.

Leo: You must be hitting something people want to know because we have five people in studio here watching. We've got the Gale [ph] family. They're visiting from Vancouver. And we've got Joe visiting from Richmond. And six more people have arrived. I guess they heard you're going to be talking about this. And we have an overflow audience here at the Brick House. We're going to have to put them in the living room.

Steve: You need stadium seating.

Leo: Boy, you're turning out to be quite the popular fellow, Steve. Good. This will be very - I'm looking forward to it. I actually saw that article. I read that article, and it scared me a little bit because, well, he mentions a well-known person who has a password that is, I would have thought, secure, and how easy it was to crack it. It reminded me a lot of passwords that I have used.

Steve: Well, yes. And that's really what I want to - I think one of the takeaways is - he

actually refers to a doctoral thesis by someone who analyzed how badly human-chosen passwords are. So the idea is you just - you don't want it to be up to you. You want to turn over responsibility to something else.

Leo: I started, after reading article, more religiously using the generator built into LastPass, and I set it for 12 characters and special characters mixed and everything. Although I'm a little disappointed. I opened an account at a new bank the other day, and I was actually quite disappointed. First of all, I could only use, I can't remember what it was, 12 or 13. After that it stopped. I couldn't use more characters, which I know from listening to the show means that they're not hashing passwords, or it wouldn't matter. And second, that they wouldn't allow me to use special characters.

Steve: Yeah, I get reports of these sorts of bad security practices constantly, so it's amazing how widespread this is.

Leo: And yet they use two-factor authentication. They do all sorts of jump-through-hoops things. So I'll just turn that on, and now I feel a little bit better about it.

Steve: That's good. Well, so one bit of news that I thought was important because it had been underreported and misreported, and it potentially bears on our listeners. And that is that our friend Moxie Marlinspike, along with an associate of his, Dave Hulton, gave a presentation at the DEF CON 20 conference some time ago, I think it was in June. But their particular talk was just, I don't know if it was that it was not understood or what the problem was. But, for example, the thread post site said "New Tool From Moxie Marlinspike Cracks Some Crypto Passwords." And I saw some other sort of lukewarm explanations of what Moxie and Dave did. And I thought, okay, I need to give this a little bit of time because people need to understand what happened.

What they looked at was a longstanding mutual authentication protocol that Microsoft created called CHAP, or in Microsoft's version, MS-CHAP, which did have security problems in v1, which they updated to v2. And it's made people uncomfortable, but not sufficiently uncomfortable. The reason this is important is that it is still the authentication protocol used in almost all VPNs. Most VPNs are still using the so-called Point-to-Point Tunneling Protocol which we talked about years ago, PPTP, you remember, when we were having fun with that acronym, Leo.

Leo: Yes, I do.

Steve: And many enterprise environments are using this for their radius server authentication in their WPA2 environments. So just to read from the beginning of Moxie's posting, he said, "The first obvious question is why we looked at MS-CHAPv2, given a lingering sense that the Internet should already know better than to rely on it." He said, "Unfortunately, however, even as an aging protocol with some prevalent criticism, it's still used quite pervasively. It shows up most notably in PPTP VPNs and is also used quite heavily in WPA2 Enterprise environments often in cases where its mutual authentication properties are being relied upon.

"For the talk, we put together a list of the hundreds of VPN providers which are now dependent on PPTP." That is, all of these commercial VPNs that are available online are

using PPTP, Point-to-Point Tunneling Protocol. "This included some high-profile examples such as iPredator, The Pirate Bay's VPN service, which is presumably designed to protect communication from state-level observation," and so on. And in the case of The Pirate Bay, they say on their site, "Right now we only offer PPTP," where this MS-CHAP is the authentication protocol for it.

Well, without going into infinite detail - because there's really no point. Moxie does describe this on his CloudCracker site in sufficient detail. What they essentially did is they carefully analyzed the handshake which the endpoints that are negotiating for mutual authentication go through, and looked carefully at the assumptions and what was known. And they managed to reduce it through a very clever series of analysis down to a total complexity of only 2^56. So that's 56 bits, which is the strength of a single...

Leo: That's terrible.

Steve: ...DES encryption. And we all know that DES has long since been determined to be extremely far from secure enough.

Leo: 3DES is what we use now, three times; right?

Steve: Yes, yes. And I had - I don't see it. I had it, and I'm afraid I closed it. I had the page open to - shoot - to his blog site.

Leo: I can link to it. I have your notes.

Steve: I wanted to explain what he's done so that I get everyone's attention here because this is - I found it here. This is significant, essentially, in what these guys have done because the way Moxie explained it, on his page he breaks down the protocol, shows how they divide this thing in the way the handshaking is going, all the way down to a key space, essentially, of 2^56, which we know is not strong.

So then he says, "At this point a question of feasibility remains. In 1998 the EFF used ASICs" - Application-Specific ICs - "to build [their] Deep Crack [machine], which cost [them]" - in 1998 - "250,000 and took an average of 4.5 days to crack a [56-bit DES] key. David Hulton's company, Pico Computing, specializes in building FPGA (Field Programmable Gate Array) hardware" - which is like sort of a modern-day version of an ASIC, but you're able to just load the circuitry into it through software. He "specializes in building FPGA hardware for cryptography applications. They were able to build an FPGA box that implemented DES as a real pipeline, with one DES operation for each clock cycle. With 40 cores at 450MHz, that's 18 billion keys per second" that this box, based on basically standard hardware, is able to crack.

"With 48 FPGA [chips], the Pico Computing DES-cracking box gives us a worst case of about 23 hours for cracking a DES key, and an average case of about a day" - I'm sorry, an average case of half a day, worst case 23. So that's what we would expect. You might get lucky and hit it on average of half of the worst-case time. So the worst case, about 23 hours, is the entire 56-bit key space, and you're probably going to get it in half that.

Leo: Wow. That's the average, right, is half that.

Steve: Right. But these guys, what I love about Moxie is he drives it, because he has moxie, all the way home. He says, "With Pico Computing's DES-cracking machine in hand, we can now crack any MS-CHAPv2 handshake in less than a day. It wouldn't be a ton of fun if only David or I could crack MS-CHAPv2 handshakes, however. So we've integrated the DES-cracking box with CloudCracker, in order to make David and his team's genius and skills and resources available to everyone. We've published a tool called 'chapcrack,' which will parse a network capture for any MS-CHAPv2 handshakes. For each handshake, it outputs the username, the known plaintext, two known ciphertexts, and will crack the third DES key. It will also out put a CloudCracker 'token,' which is an encoded format of the three parameters we need for our divide-and-conquer attack," which is what they called this attack.

"When this token is [then] submitted to CloudCracker, the job is transmitted to Pico Computing's DES-cracking box, and you receive your results in under a day. What do you win? At this point, you can plug the cracked MD4 hash CloudCracker gives you back into chapcrack, and it will decrypt the entire network capture and all future captures for that user." So he says, "Alternatively, you can also use it to log into the user's VPN service or WPA2 Enterprise radius server."

So what they have done is completely reduced this protocol to rubble, which is the No. 1 most used protocol for VPN authentication today, used by thousands of commercial VPN services. So this means you capture the traffic. You've got to jump through some hoops, so this is not yet single pushbutton ease, but we know that that's the way these things begin is not being that way, and they evolve into that. And this will decrypt the entire VPN encrypted communications. So I thought that was just worth mentioning, Leo.

Leo: But does anybody use, what is it, CHAP, does anybody use it anymore?

Steve: Every VPN.

Leo: Oh, crap.

Steve: Yeah.

Leo: CHAPv2. That's how VP - oh. I thought it was just - not just Microsoft.

Steve: No. The Point-to-Point Tunneling protocol is...

Leo: All PPTP does it, not just Microsoft's implementation?

Steve: Yes, all PPTP.

Leo: Oh, crap. But it is Microsoft's CHAPv2.

Steve: It is, but it's an RFC standard that everybody else has adopted.

Leo: Wow.

Steve: Yeah.

Leo: And we thought VPN was safe. So how practical is this?

Steve: Again, this is, well, for example, you want to use an SSL VPN. That's the buzzword you look for. You don't want to use PPTP. Unfortunately, most VPNs today, and for example, The Pirate Bay's VPN, that's all they offer for their service, is they're all using PPTP. Just because of inertia, basically. It's stable. It's reliable. It's free. It's online and open source, and that's what people set up when they're creating VPNs. Well, if your traffic is captured, Moxie and his team now make it trivial to basically decrypt the entire conversation.

Leo: Wow.

Steve: Yeah.

Leo: Wow. That is kind of surprising. Wow.

Steve: Now, I saw something really nice. Reuters on Monday put up a story about how a group is working on strengthening the security of cars, which I'm really glad for because I'm really worried about automotive security as everybody rushes to compete with each other headlong into more technology. What caught my eye in the article, and the reason I wanted to mention it, was a reference specifically to Ford, and it's good news. So I'm just going to share that much of the top of the article. This was written by Jim Finkle for Reuters. He said, in Boston:

"A team of top hackers working for Intel Corp.'s security division toil away in a West Coast garage, searching for electronic bugs that could make automobiles vulnerable to lethal computer viruses. Intel's McAfee unit, which is best known for software that fights PC viruses, is one of a handful of firms that are looking to protect the dozens of tiny computers and electronic communications systems that are [today] built into every modern car. It's scary business. Security experts say that automakers have so far failed to adequately protect these systems, leaving them vulnerable to hacks by attackers looking to steal cars, eavesdrop on conversations, or even harm passengers by causing vehicles to crash." And of course we've talked about these problems, those specific problems, through the last few years as this has begun to make the news.

"'You can definitely kill people,' said John Bumgarner...."

Leo: Oy.

Steve: Sorry?

Leo: Oy.

Steve: Yeah. Oh, yeah. You can slam on the brakes.

Leo: You could, like, lock up the brakes, yeah.

Steve: Yes, "...said John Bumgarner, chief technology officer of the U.S. Cyber Consequences Unit, a nonprofit organization that helps companies analyze the potential for targeted computer attacks on their networks and products. To date there have been no reports of violent attacks on automobiles using a computer virus, according to SAE International, an association of more than 128,000 technical professionals working in the aerospace and the auto industries. Yet, Ford spokesman Alan Hall said his company had tasked its security engineers with making its Sync in-vehicle communications and entertainment system as resistant as possible to attack. 'Ford is taking the threat very seriously and investing in security solutions that are built into the product from the outset,' he said." Which was, frankly, that was good news.

Leo: That's a relief. No, I know Alan, and I know this because I've talked to their technologists, as you know, for years. And they have a very clear separation between the in-car computer, the real computer, and the - it's not the same computer, even.

Steve: There was one instance where an infected CD was played in the CD player that allowed it to get out of the player and into the car's automotive electronics.

Leo: Not in a Ford.

Steve: Not in a Ford.

Leo: I just want to point this out.

Steve: Not in a Ford.

Leo: In another company's setup. No, and we've seen all sorts of weird - you could drive by, and they can reprogram the car. But that's not in a Ford because they understand this, and they very intentionally separate the two systems.

Steve: Yup. And just to wrap up, it says, "And a group of U.S. computer scientists shook the industry in 2010" - and we covered it then - "with a landmark study that showed viruses could damage cars when they were moving at high speeds."

Leo: Yup. Yup.

Steve: "Their tests were done at a decommissioned airport."

Leo: In case the car took off. Yeah, no, I remember we did that story. In fact, it was after that that I asked the Ford folks. And they said, oh, no, no, no, you couldn't - no, no, couldn't happen with us.

Steve: Yup. And I was just glad to see that Ford and Alan were quoted as saying, yes, we're on top of it, because this is something you need to get ahead. You need to be proactive rather than reactive.

And finally, Amazon has a new service that I just - every time I look at the name, it makes me grin. And I'm glad for it because I think it's really interesting. It's called Glacier. What it is...

Leo: Oh, I like this.

Steve: Yes. It is long-term...

Leo: I like the logo, too.

Steve: ...archival, yes, long-term archival storage for an amazingly low price. The idea is it is glacial in its response.

Leo: Brilliant.

Steve: Yes. So again, quoting from them, they said: "Amazon Glacier is an extreme" - it's aws.amazon.com/glacier. "Amazon Glacier is an extremely low-cost storage service that provides secure and durable storage for data archiving and backup. In order to keep costs low, Amazon Glacier is optimized for data that is infrequently accessed and for which retrieval times of several hours are suitable." So it's offline storage. "With Amazon Glacier, customers can reliably store large or small amounts of data for as little as" - and here it is - "$0.01 per gigabyte per month."

Leo: That's awesome. That means it's 10 bucks a terabyte. That's great.

Steve: Yes. Let's see, so...

Leo: Is that right?

Steve: Let's see. So 100GB for $1 per month.

Leo: So 1000GB for $10, yeah. Terabyte is $10. That's amazing. That's great.

Steve: Yes. So the way they have it organized, data are stored as archives. Retrieval of an archive takes, they're quoting, three to five hours. Archives are organized in vaults. And there is zero transit cost to upload.

Leo: Unlike Amazon S3, where you pay for...

Steve: Actually S3 has gone to zero.

Leo: Oh, it has, too. Because you used to pay for transmission plus storage. Now it's just storage.

Steve: Yes. Well, no. Or retrieval. You do pay for transit down to you, but not up to them, which I love because my main application is sending images up to them after I encrypt it locally.

Leo: So that's what I'm going to do. This is where I put my photos. And then my children, and my children's children, will be able to ignore those photos for generations to come.

Steve: So I don't know what the model is. I don't think it's tape. I think it must be powered-down, offline hard drives, that they must actually take them and pull them out of servers and store them somewhere so that, if you submit a request, maybe it's a bot or a robot or I don't know if it's a person or what the model is, but...

Leo: Somebody goes down - I know what it is. It's an Indiana Jones-style giant warehouse.

Steve: Yes. I think it probably actually is, Leo.

Leo: And as a machine goes down, a robot goes down and gets the hard drive, pulls it out, and runs it down to the data center where - I bet that's what it is.

Steve: I bet it's not a bot, though, because that wouldn't really explain that time delay, unless they just want to give themselves some buffer.

Leo: They want to give you the time delay so that you don't use this instead of S3. That could be completely arbitrary.

Steve: That's a very good point.

Leo: It's just a distinction. I love this. This is great. Who doesn't - who, I mean...

Steve: I know. I'm never needing to access these snapshot images fast. But you don't want them to break the bank, and three to five hours, that's fine for a big image. So anyway, I really think this is a nice service.

Leo: Very, very, very cool.

Steve: And then I have...

Leo: Is it going? Is it started?

Steve: Yeah, it's available now.

Leo: That's great.

Steve: And I did get a nice note, a short one, from a Wayne Scott in Australia, who said, "I've known about SpinRite for many years, using it on my own PCs after deciding to purchase a copy when I had to do a recovery on a customer's hard drive. Another company had attempted data recovery before that and returned the drive as faulty and unrecoverable. So I wanted to have a go. Where the data recovery company had failed, other" - he says "other SpinRite," it's a little typo here. "Where the data recovery company had failed, SpinRite made the drive readable once again, and I got the data off. The customer was very happy because it was for the tax department." So once again...

Leo: Ooh, wow.

Steve: ...SpinRite to the rescue.

Leo: You don't want to lose your tax information.

Steve: No, you don't.

Leo: Unh-unh. All right. We're going to talk about "The Death of Clever." I like this

subject. Are you ready? Let's talk about the end of clever.

Steve: So Dan's article, or his security blog posting, was "Passwords Under Assault." Anyone who wants to read the entire four-page piece can just Google "Passwords Under Assault," and it's the first link that comes up. And he titled it, "Why passwords have never been weaker and crackers have never been stronger." Which sort of reminds us of the famous Bruce Schneier quote, where he noted years ago that attacks never get weaker, they only get better. And Dan said, "Thanks to real-world data, the keys to your digital kingdom are under assault."

So essentially what's happened is there have been consequences, there's evolutionary effects that we would expect, that is, passwords are very tasty fruit for hackers to try to grab. And, unfortunately, websites have proven themselves surprisingly inept at managing user logon credentials. We're routinely, actually, covering the major breaches in passwords. It was just a couple months ago, in June, that LinkedIn famously lost control of 6.5 million passwords. What's happened is, as a consequence of those and other breaches - there was another major gaming site that lost, I think it was 32 million of their user passwords all at once. And so what's happened is it's moved the hackers' understanding of what passwords people are using from theoretical, like the planets of the Klingon universe, to the actual. And we've learned weird things, like "monkey" is used unusually often, Leo.

Leo: It's so embarrassing because it in fact was my default password many years ago.

Steve: But how obscure is that? But the point is that, for some bizarre reason, lots of people chose the word "monkey." Well, nobody would guess that. So it's only by looking, doing statistical analysis of actual password databases, that these sorts of things come out. Another thing that is often occurring is that people capitalize words, instead of them being all uppercase or all lowercase. They tend to - first character is capital, then the rest of them are lowercase. Many times people create passwords which are word followed by four numbers, like their date of birth, for example, or 1492, something that is memorable to them, but they think, oh, this is clever.

Leo: But that's what you're talking about with Password Haystacks. That's padding. That's not a bad thing, as long as it's not guessable. Right?

Steve: Well, okay. So the problem with patterns, like the idea of eight characters where the first one is uppercase and the other ones are lowercase and then, for example, a four-digit number, if you made it five digits, that is, if you broke the pattern, then you get security. If you don't, what analysis of databases have shown hackers is that, in the same way that for some bizarre reason the password "monkey" gets chosen way more often than randomly, people are using eight-character alphabetic words followed by four-character numbers, I mean, exactly that pattern. And so what happens is, if that's known, or even just believed, that is, if it's tried for, then it completely changes the math.

For example, say that you didn't know what a 12-character password was, and that it could use the full alphabet and special characters and numbers. Well, any one character,

as we've talked about many times, could have approximately 96 different possibilities. So 12 of those would be 96^12, since it's 96 for the first character, 96 for the second character, 96 for the third. But we also know that that really only applies if the 12 characters are really random. They could be anything. And 96 raised to the power of 12 is 612.7 times 10^21. Huge number. That's 612,700 billion billion possibilities for 12 characters.

But people don't choose their 12 characters randomly. And what statistical analysis of these captured online databases have shown hackers is that, as I was saying, for example, there's a huge preponderance of first letter is capitalized, the next seven are lowercase alpha, and then they're followed by four digits that is, like, a year. It's something generally in the 20th Century. So what that does is that dramatically changes the math. Now that means you only have 26^8 power since you have only - you know you're going to have capital A through capital Z, then lowercase A through Z for the next seven characters. Then say that you didn't even constrain it to a modern-era year, but you just did 0000 to 9999, so now you're at 26^8 times 10,000. Well, that's only 2.08 million possibilities, compared to 612,700 billion billion possibilities.

So the point is that, what hackers have done is, by analyzing the actual databases of captured passwords, they have found all of these tendencies. It is absolutely no longer the case that we can do anything clever. We cannot use, like, "Prince$$," where we change the S's into dollar signs. They got that. You can't use...

Leo: Cracked.

Steve: Sorry, Leo. You can't turn your E's into 3's. They got that, too. I mean, all of the kinds of things that people typically do, thinking that they're being clever, trying to sort of - essentially we're trying to compromise. We're trying to come up with something that's sort of ours and that we think nobody else is going to do. Well, surprisingly, because we're all human, and we have similar experience, we're generally doing the same things, it turns out. When you statistically look at 100 million passwords, there aren't that many possible things that people can do that meet these criteria. And of course there's certainly some communication among people. Not everyone is coming up with these things on their own. They're talking to their friends about, oh, what do you do, how do you make passwords? And so they share some of their ideas.

Oh, the site was RockYou.com which, in 2009, through a SQL injection attack, lost their 32 million plaintext passwords, which all went into this huge 100 million-plus hopper for statistical analysis.

So the other thing that has happened is, and this is the evolutionary part, not only are hackers really focusing on this, but as we know, there's been huge movement in technology over time. We've talked about how GPUs, the graphic processing units that are now powering our graphics cards in order to give us the 3D realism and high frame rate performance that we want for gaming, those can be repurposed to create essentially cryptographic pipelines which are able to run cryptographic algorithms at very high speed.

One of the takeaways from all of this is that hashing was never the right thing to do. Hashing was better than leaving things in plaintext, certainly. But hashes were designed, as we have said before, for speed. They were designed to be efficient. But efficiency is exactly what you don't want in password security because it allows brute-forcing to run at tens of millions of guesses per second. So while it's certainly better that sites have

been hashing their passwords than not, it turns out that we no longer should consider that very useful. Certainly not if they are unsalted hashes.

The LinkedIn breach that we talked about in June, where 6.5 million passwords were lost, to give you some sense for this, for what this really means in the real world, independent security researcher Jeremi Gosney took the leaked LinkedIn unsalted but hashed, it was hashed with SHA-1, database. He applied it against his 500 million strong word list of common words, using a block of GPUs, which are able to make 15.5 billion guesses per second. This is not the NSA. This is some guy in his bedroom who can do 15.5 billion guesses per second. Against LinkedIn's 6.5 million passwords, he cracked the first 20 percent in 30 seconds. He had one out of every five of that 6.5 million passwords cracked in 30 seconds. The next 33 percent took two hours, so in two hours and 30 seconds he had 53 percent of them cracked. It began to slow down exponentially so that, after a day, 24 hours, he was at 64 percent of the 6.5 million passwords cracked. And after five days, he had an additional 24 percent.

So we're not talking long-term protection here, if a database gets loose, even if it is salted. That's no longer the case. The other interesting thing to Google is a new, open source, free, GPU-based cracking facility called Hash Cat. You should bring it up onscreen, Leo, H-a-s-h and then space, Cat, Hash Cat, calls itself "advanced password recovery." It's the first result in Google, and it's just HashCat.net, also. And it says, "Download the latest version." The requirements are, for NVIDIA users, you need to have their ForceWare 290.40 or later; for AMD users, you need to have Catalyst 12.4 or later. And it looks like a very nice, professional piece of work. Under features they claim the world's fastest md5crypt, phpass, mscash2 and WPA/WPA2 cracker; the world's first and only GPGPU rule-based engine. Its multi-GPU support can run 16 graphics processing units in parallel; has native binaries for both Linux and Windows. Low resource utilization - you can still watch movies or play games while cracking in the background. Isn't that convenient.

Leo: There's plenty of CPUs to spare.

Steve: Absolutely. Focuses on highly iterated modern hashes; uses dictionary-based attacks. Oh, you can even pause it and resume it while cracking. So it has all the features of a modern password-cracking system. It can read words from a file, so you can have dictionaries; can read from standard input. It has an integrated thermal watchdog, just in case you overheat your system by running too many hashes too quickly.

Leo: This is nicely done. Nicely.

Steve: Isn't it nice. More than 20 algorithms: MD5, Joomla, osCommerce, SHA-1, Base64, Oracle 11g. Here we have OS X v10.4, 10.5, 10.6, and a little bit lower is 10.7. Not to be outdone, we've got Double MD5, SHA-256. Oh, there's NT LAN Manager, Microsoft's NTLM is there. And on and on and on. So, yeah. Oh, and runs in both 32- and 64-bit OSes, tested and fully supported. And free. Did I mention that? Free.

Leo: Yes, free, free.

Steve: So you no longer need to be a GPU programming guru. And of course this is the same pattern that we see over and over and over. Remember when Firesheep was released, which allowed anyone to download this add-on for Firefox, wander over to Starbucks, and people's pictures and logon credentials started popping onto the screen. What we're seeing is the standard evolution in password-cracking technology that once truly was rocket science. Now it's turnkey.

It's not quite where Moxie is with capture packets through the ether, dump it into his CloudCrack, and he'll handle all that for you. But somebody who is interested in playing with this no longer needs to write a lot of code or understand it. There are videos on that site, how-tos, forums, and an offer to download the latest version into your GPU, and you, too, can start cracking like crazy. And of course those forums will have links to the 100-million-plus password databases and 500-million word lists and so forth. So, I mean, this really has gone exponential in terms of the fun that people are having and how easy it is...

Leo: That's the way to put it, "fun."

Steve: ...to get into the password-cracking business. So looking at these lists, essentially no one is any longer believing that people's passwords are truly maximum entropy random. What this says is, I mean, as you said earlier, Leo, using a LastPass-generated, long, absolutely unmemorizable password is the best thing you can do. Now, my Haystacks notion was a compromise, admittedly. It was the recognition that, in the face of brute-forcing, length trumps complexity because, if you're off by one character, you get no result whatsoever. It's got to be an exact match, so close doesn't count. So the Password Haystacks idea was to get you something long, if you couldn't use LastPass, or for whatever reason you didn't want to, you needed something memorable that would not be quickly crackable.

Leo: That's the problem, memorable. And that's where we get this complication; right?

Steve: Yes. Now, it is still the case, and I think on maybe the last page of Dan's four-page piece he shows a very interesting chart which you should put on the screen if you can find it there, Leo, where it goes exponential. There is still the so-called "password cracking wall," which means, if none of these dictionary attacks work, if your password isn't something, a normal word where the E's are changed into 3's or three exclamation points are added to the end, or if it's not something where you have been clever, but in fact the password you're using doesn't fail in any of those ways, and you have to assume now clever is broken, clever is no longer good enough, if it doesn't match that, then you're back to brute-forcing.

Leo: And boy, does it go up. After seven characters of true randomness, it gets impossible.

Steve: Yes. And it doesn't matter if you use a GPU or if you use CloudCracking or anything. So to put a number on it, there's a picture shown of a homebrew, $12,000 machine containing eight AMD Radeon HD7970 GPU cards, running Hash Cat. It requires 12 hours to brute-force the entire eight-character password keyspace.

Leo: Of random numbers, random letters.

Steve: Yes. Now, remember that what I said at the very top of the Password Haystacks page was, if every single one is tested, sooner or later they will get yours. So this thing, for $12,000, eight AMD Radeon GPU cards running Hash Cat, takes 12 hours to test every possible eight-character password.

Leo: That's upper and lower, digits and symbols. The whole thing.

Steve: Yes. But now remember, you add one character to that and it's 96 times longer. One more character, 96 times again. One more character, 96 times again. So that's why this thing still exponentiates. It goes straight up because, if you really have very high entropy, if you have not - if your password hasn't crumbled because you did something that you thought was clever - oh, another one that is mentioned here I thought was interesting, that apparently, again, lots of people think, oh, I'm being tricky, no one's going to think of this. It is to spell a word forwards and then concatenate that word backwards. Whoops. They know about that, too.

Leo: That's not that tricky. Really it shouldn't be the death of clever, just the death of kind of clever, or maybe clever, or you think you're clever.

Steve: Yeah, I just don't think you can be clever enough. That's the problem.

Leo: Maybe that's it, yeah, random is better than clever.

Steve: What has been learned is that we're just not very good at coming up with something really clever. The classic was transposing the keys on the keyboard. Yeah, they know about that, too. There are dictionaries of all the words shifted one up and over to the left, or one down and over to the right and so forth. All of that. So the idea is that what we as users need to appreciate is that this day and age, this is the low-hanging fruit. The hackers are just having a ball, literally spending their time thinking, okay. They'll look at a password that was captured and think, that looks random. Where did that come from? And they'll realize, oh, look, that's shifted down and to the right from a normal word in the English language. So they add that strategy to their cracking library, and suddenly all passwords of that form fall to the addition of that strategy. And this is only going to get better in the future. So if you haven't yet switched to something that will not fall to this kind of attack, the sooner you do, the better.

Leo: Wow.

Steve: And one other piece of analysis showed that the typical web user is logging onto - I'm trying to find the number here. I just saw it. I think it's 26 different sites, but only using 6.5, on average, that is, between six and seven different passwords. So it is still the case that we're seeing cross-site - oh, it's 25, 25 separate accounts, but uses between six and seven passwords for protection. So there is still a substantial amount of

password reuse going on. And we know why that's not safe because, if a site like LinkedIn, with its 6.5 million passwords and associated email addresses, if those passwords get cracked, and 90 percent of them have been now after a couple months, and you use the same credential elsewhere, then you're very vulnerable to impersonation, which is of course all of what this is supposed to be protecting us from.

So the argument is that, yes, over time, we are moving to multifactor authentication. But unfortunately, today, in this day and age, we're still being forced to authenticate with passwords. And this is where the action is. People are having fun just with the idea that a GPU has this much computing power, and all these resources are available on the Internet. You no longer need to be a rocket scientist in order to play these games and play with this stuff. And the consequence is that more and more people are going to be doing so, and freely downloadable software is going to be getting more and more clever. So that anything that you've thought of that you think is like your trick, your tricks have gotten loose, or people like your tricks have gotten loose. They've been analyzed and added to the strategy. So that it's no longer just simple, try every possible password, aaaaaa, aaaaab, aaaaac and so forth.

So what we need to do is abandon this and just use entropy, ultra-high-entropy passwords, and something then to manage them, like LastPass of course is what I use, 1Password, and there's a collection of great utilities to help people remember. I haven't looked at any of the others, that is, the security of any of the others other than LastPass. So that's the one, as we know, that I've looked at closely. And as far as I can tell, they've done everything right.

But I would say, from this point on, and as you have the chance, you really want to migrate away from things you did that you felt were clever because, if those get loose - and unfortunately that is the attack model today, it's not somebody logging in through the web interface, guessing your account. No. It's that a database on the backend escapes, and then millions of credentials are being cracked in parallel.

Leo: [Sigh]

Steve: Yeah.

Leo: I have all these tricks that I use. But really you're convincing me that tricks are just a bad idea. You just should use totally random passwords.

Steve: Who would have guesses that everyone would choose "monkey"? We weren't telling each other.

Leo: I didn't think that was a trick. I knew that was bad. A dictionary word, yeah. But then, I mean, all these other tricks are bad, too, apparently.

Steve: Yes, yes. Well, I mean, think about it. Anything you can think of, they can, too. But more importantly, you thought of it, and you used it. And then some website where you used it got cracked. What happens is the hackers look closely at the ones they could not crack, and they go, hmm, why couldn't we crack...

Leo: Oh, that's neat, yeah. And they try to find patterns in it.

Steve: They zero, exactly, they zero in on the ones they couldn't crack, and that leads them to strategies they don't yet have crackers for, and so they add crackers for those strategies.

Leo: And this is why you need regular wars so that people like this can go to places like Bletchley Park and use their genius for good, not ill.

Steve: Yeah.

Leo: I rest my case. No, I don't.

Steve: The good news is we're talking about this on the podcast. We've got tens of thousands of listeners who are hearing this, who have advance notice...

Leo: Are doing as I'm doing and changing their passwords right now.

Steve: Yeah.

Leo: Wow.

Steve: The nice thing, too, is you can use something like LastPass that has learned all of your passwords, and go through and figure out which ones you really need to change. I have had some neat feedback also from our podcast about, was it Tom Honan?

Leo: Mat Honan.

Steve: Mat, yeah, Mat. Got a little bit of alliteration there, or dyslexia. Yeah, Mat. Many people were focusing their password recovery in the same way that he was. And so I've had a lot of feedback from people who said, hey, thanks for explaining that. I was doing the same thing. I've broken my accounts apart now so that they're no longer chaining in the same way that Mat was. So that's good, too. I'm glad that we're able to help people.

Leo: LastPass has a security, a password security audit feature I'm seeing. Jesse tells me in the chatroom. I didn't know about that.

Steve: I thought it did, yes.

Leo: So try that. It finds duplicate passwords, I think, is mostly what it does.

Steve: Well, but that's good, too, because we know that you don't want to have re-use.

Leo: Unique is good, yeah.

Steve: Yeah.

Leo: Well, I'm more and more using that Generate from LastPass to generate passwords. It seems to...

Steve: I do, too. I resisted it at first because it was like, oh, this just looks like total noise. But that's the point. You want something that looks like noise and trust LastPass to remember it for you.

Leo: Mr. Steve Gibson is the man at GRC.com. That's where you can find SpinRite, his bread and butter and the world's best hard drive maintenance and recovery utility, SpinRite, at GRC.com. He's got a lot of freebies there, too, though, including lots of information on passwords and security of all kinds, freebies. It's GRC.com. That's where you'll find the most compact versions of this show. There's a 16Kb audio version Steve makes available and also transcriptions, which are even more compact. We do the big fat ones, the audio and the video, at TWiT.tv/Security Now!. And of course the best way to get it is to subscribe so you don't miss an episode.

You can watch live. And I think it's fun to watch live. I talk with the chatroom while we're doing the show and fill them in. For instance, somebody said, well, is OpenVPN safe, given this CHAPv2 crack? And of course it is because...

Steve: Yes, because it's an SSL VPN.

Leo: Yeah, yeah, so that's good news. So if you want to watch live, we do it 11:00 a.m. Pacific, 2:00 p.m. Eastern time, 1800 UTC on TWiT.tv every Wednesday. We are going to move you, Steve, just a heads-up, if Apple does this announcement September 12th. That's a Wednesday.

Steve: Ah, okay.

Leo: But we don't know yet because they haven't sent out invites.

Steve: Will I swap with MacBreak Weekly?

Leo: Yeah, we'll just swap you to Tuesday at 11:00 a.m. But we probably won't know that until September 5th because Apple doesn't like to tell anybody ahead of time.

Steve: Although I did see, you're right, I saw that a number of, like, what was it, retail staffs had been put on alert.

Leo: Verizon, yeah.

Steve: Right.

Leo: I think it's true. But I'm not going to reschedule a show based on a rumor. Andy Ihnatko's flying out based on a rumor.

Steve: Jen was asking what the iPhone 5 has. Is it, like, more than just thinner?

Leo: "It goes to 11." It'll be taller, four inches tall - well, it's all rumor - and 1136 by, what was it, 640? I can't remember. 1136.

Steve: Wow. So even more screen resolution.

Leo: Yeah, going up, though.

Steve: Yeah.

Leo: 16:9.

Steve: And thinner.

Leo: And we don't know. I don't know. I haven't seen thinner, although I think that's the speculation. And a new connector, a new nine-pin connector.

Steve: Oh, that's right, no longer that big dock, the traditional docking connector.

Leo: No 30-pin, yeah, yeah. But I don't know - LTE for sure because you've got to nowadays. But I don't know if there's much else. We've seen iOS 6. There's an improved Siri, things like that.

Steve: Yeah.

Leo: No more Google Maps.

Steve: And you've seen that the Nexus 7 is doing phenomenally well.

Leo: And rightly so, yeah. Love it. And I'm looking forward, you know the Galaxy Note comes out in a week, the Galaxy Note 2. And I love that big form factor. That must be a 5.5" screen. So it's heating up. It's getting hot in here.

Steve: It is. Fun stuff.

Leo: Thank you, Steve. Steve Gibson, GRC.com. We'll let you know as soon as Apple lets everybody else know if we're going to move shows. That's not for a few weeks yet.

Steve: Thanks. No problem for me.

Leo: All right. Thanks for joining us. We'll see you next time on Security Now!.