The Dark Escrow

Description: This week we examine Firefox's just-released and welcome re-architecture under codename "Fission." We look at a new and recently active ransomware player named "Conti" and at a recently paid, high-profile mega ransom. We then ask the question, "When they say IoT, do they mean us?" We examine the implications of a new industry term, "mean time to inventory." We'll then lighten things up a bit with a new form of CAPTCHA and, of all things, a screensaver I discovered that I cannot take my eyes off of. (Leo, it's not quite as bad as whatever that game is that you cannot stop playing, but still.) We'll then share an ample helping of closing-the-loop feedback from our terrific listeners, after which I want to conclude by predicting what I would bet we're probably going to next see emerge from the evolving ransomware business model sad though it is to utter the phrase "ransomware business model."

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-820.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-820-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We've got a great show for you. Brand new feature in Firefox. It's called Firefox Fission. Steve explains it and shows you how to turn it on. I did immediately. We'll also talk about a new form of CAPTCHA that involves playing Doom. I like this. Even has the sound. And then Steve's vision for the future of ransomware. Sort of. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 820, recorded Tuesday, May 25th, 2021: The Dark Escrow.

It's time for Security Now!, the show where we cover your security, your privacy, your safety online with Mr. Steven Gibson of GRC.com. Hello, Steve.

Steve Gibson: And yes, ladies and gentlemen, those of you listening and watching live, check your clocks. It's actually not quite 1:30 a.m. Wait, p.m.

Leo: We're early. Shall we just stop and wait for three minutes?

Steve: No, no, no, no. Let's see how much we can cram in before the clock hits 1:30 and the official start time. This is Episode 820 for May 25th. And I titled this "The Dark Escrow" just because dark is cool; you know? I mean, you can kind of have like dark anything, and it's neat. Like the dark melon. Well, maybe not that. But anyway, as a consequence of some of the aftermath of what happened with the DarkSide ransomware

when they tiptoed into the Colonial Pipeline mess, and some things that have happened since, which we're going to talk about, I have sort of some thoughts about where we're probably going to go next. So I wanted to share that.

But we've got lots of interesting stuff to talk about. We've got Firefox's just-released and welcome re-architecture under codename "Fission." We've been waiting for this for a while. They're a little late to the game. It turns out - we'll talk about this - it's difficult to do what they've done. Chrome did it, and the Firefox, we've been waiting for them to follow. And it's now something that all Firefox users can turn on. I immediately did. And it's a good thing.

We're going to look at a new and recently active ransomware player named Conti since they're looking like they're going to be something we'll be talking about to some degree in the future. So I thought I would introduce our listeners to Conti. And also at a recently paid, high-profile, what you would have to call literally a "mega ransom." We then ask the question, when they say IoT, do they mean us? And we examine the implications of a new industry term, "mean time to inventory." And it's not good. We'll then lighten things up a bit with a new form of CAPTCHA. And anyone who wants to cheat can - I'll tell you that it's this week's GRC shortcut. So those who've been paying attention know what the URL would be for this week's shortcut, and it's kind of fun.

Then, of all things, I want to share a screensaver that I discovered, that I cannot take my eyes off of. And Leo, it's not quite as bad as whatever that game is that you cannot stop playing. But still it's absorbing more of my time than I want to admit, just sitting there, like, whoa, look at that.

Leo: I don't know if staring at a screensaver's worse than playing a game. But, all right.

Steve: Yeah, I'm not getting any points for it, that's for sure. Okay. Then we're going to share an ample helping of closing-the-loop feedback from our terrific listeners because there was just a bunch of stuff that came in that I wanted to share. And then we're going to conclude by predicting what I would bet we're probably going to next see emerge from the evolving ransomware business model, sad though it is to even utter the phrase "ransomware business model." Oh, and we do have a great Scott Adams cartoon as our Picture of the Week. So I think another great podcast for our listeners. And we can begin officially, now that it is now 1:30.

I've had this one in the queue for a while, and I thought, yeah, now is the time. This is on the theme of being careful what you ask for. So we have a three-frame cartoon by the brilliant Scott Adams. We have the boss is addressing three of his employees. And so in the first frame we have the boss saying, "Our goal is to write bug-free software. I'll pay a $10 bonus for every bug you find and fix." And then in the middle frame we have the three employees. The first one, "Yahoo!" And then the second one says, "We're rich." And the third one says, "Yes! Yes! Yes!" And then, sort of wondering if he's come up with the right policy, we see the boss saying, "I hope this drives the right behavior." And one of the employees says, "I'm going to write me a new minivan this afternoon." So, yeah, if your employees are able to find their own bugs, well, that may not be the best way to produce bug-free software, is the point.

So Firefox has finally achieved sustained Fission. Fission is Mozilla's name for full-site isolation. That is, fission of course being breaking things apart. The original design for all web browsers was as a single app running on an operating system. Or a bit more formally, a process being hosted by the OS. And the trouble was browsers have kept

growing more and more complex, to the point that the line you know has become blurred between an OS app and what a browser can do with a web page.

Case in point, the other day I followed a link to a drag-and-drop mechanical electrical circuit design page where, for example, a capacitor was emulated by a balloon being filled up. An inductor was created by a flywheel. A resistor was a narrowing pipe and so on. And I initially took the page for granted, until it occurred to me that this amazing, fully animated creation was not an app in a classic sense, but this was being entirely done on a web page in my browser. And we're all kind of getting used to that.

So yeah. Today's browsers have become incredibly complex and capable. But as we know, complexity is the enemy of security. It's not impossible to have both. But it turns out that it's impractically expensive to actually have both. So we get features that mostly work. And as for security, we hope for the best while we fix the flaws that are later uncovered on the fly, after the fact.

But browsers are sort of a special case for problems because the presence of the near certainty of flaws as we just keep fussing with our browsers, making them more and more capable, but more and more complex, raises the specter of containment. If flaws are inevitable, then what we want is to at least contain them. The last thing we want is for a flaw to be leveraged to reach across web browser pages, or really across domains, to infect an unrelated page to leverage its access to some other site, for example, where we're logged into our banking site. We don't want some sketchy page we're visiting to be able to peruse the other tabs we have open and say, oh, look what we found here. Let's go there.

And a perfect analogy is, for example, of like this problem, the need for separation is my insistence upon applying network segmentation for IoT devices which are inherently high risk. If we cannot guarantee that our light switches and our blow dryers will not attack us, at least we can contain any potential damage by placing all of those less trustworthy devices onto their own low trust network segment so they can't even see the rest of our internal network. So it's a form of sandboxing. And this exactly mimics the designs that browsers have been working to adopt for several years.

It's taken years because switching from a mono process architecture, which is where they all began, to a process per domain is far more easily said than done. Google initially released an experimental site isolation feature that they had already been hard at work at for quite some time in Chrome 64, back at the end of 2017, and it became generally available essentially about half a year later in May of 2018. Mozilla realized that was where the future would be. So they began work on the same thing the month before Chrome's general availability, in the Mozilla case in April of 2018.

They formally announced their plans to do it nearly a year later, I mean, it's really that hard, in February of 2019. And then that's when they disclosed the code name Fission. And a little more than two years later, here we are, last week, it's finally here. But even so, it's not yet turned on by default. Anyone listening to this podcast who's a Firefox user - I know, Leo, you are - who has a current release of Firefox, which of course you get just by making sure you go to like About Firefox, and then it'll update it, if it hasn't already, can turn this on.

Go to about:config, which we're often visiting, about:config in the Firefox URL bar. Then search for "fission," F-I-S-S-I-O-N, and locate fission.autostart. I think it's like the fifth or sixth one down. And it should now - you should set it to true. And after making those changes, Firefox won't auto prompt for a restart, as Chrome does. But if you go to about:profiles, that page in the upper right contains a manual restart button that will do the trick, that'll restart Firefox. You'll come up, and from then on every tab you open -

actually, it's more than that. It's even more granular. Every domain you visit will get its own process.

So by creating a separate standalone OS process for each domain, which the browser pulls content from in order to prevent anything nasty from escaping, essentially what we're doing is we're leveraging the mature process separation that our OS has long been enforcing. Basically the browser sort of is turning over responsibility for really keeping the containment of domains by creating a domain per process. And the reason it's tricky is that if you then, as many pages do, have frames that are pulling in off-domain content, then that needs to be in its process. Yet even so, it needs to get rendered in the page in this process.

So again, what we've seen is years of work going into this for our browsers. Safari has something, kind of a weak version of this, but hasn't really stepped up to do it because, I mean, it really is, it's like starting over from scratch in order to do it. Firefox has it. Now that Chrome has it, of course that means that Chromium has it, which means that all the other, well, Edge for example from Microsoft, and Brave and Vivaldi and all the Chromium-based browsers, also get it just as a consequence of being Chromium-based. So clearly it is more resource intensive.

One of the problems early on had been that a lot more RAM was being used because basically what looks like one thing that you're running on your desktop, especially considering the number of domains that your typical web page pulls from now, well, if we really want containment, we've got to fire off an OS process for every one of those. So you just can't throw this together. In order to make it work and make it fast and continue to make it small, that requires some tricky work. Anyway, we've got it now in Firefox. It's been available in Chrome for a while. It's the standard moving forward.

The latest ransomware on the block is called Conti. And based on the fact that the FBI has recently identified 16 different successful penetration attacks by Conti, starting up from like nothing a short time ago, I have a feeling we're going to be talking about those attacks in the future. So I thought we ought to put it on our map. Conti is believed also to be a Russia-based, if not backed, cybercrime group known as Wizard Spider. So Conti technically is the ransomware, also sort of the initiative. And the terms, they get kind of mixed up.

But Wizard Spider is the group that is behind Conti. They use phishing attacks to install the TrickBot and the BazarLoader trojans that subsequently provide remote access to the infected machines. Then using that remote access, they move laterally through networks, stealing credentials and harvesting unencrypted data stored on websites and servers, as we know. That's now the first move is to exfiltrate everything they can off to the cloud somewhere.

Once the attackers have stolen everything of value and gained access to Windows domain credentials, what's interesting is they now explicitly wait for a quiet time on the network based on whatever entity it is that they have infiltrated. You know, probably like early, very early on a Sunday morning local time. Or maybe very early on a Saturday morning, if they note that Saturdays stay quiet. The point is they're wanting as much time as possible before they're discovered, after triggering and deploying the encrypting Conti ransomware throughout the network.

So the Wizard Spider Conti gang then use, of course, the stolen data as leverage to force their victims into paying a ransom by threatening to release it on their ransom dark leak site if they are not paid. Recent high-profile Conti ransomware attacks include the FreePBX developer Sangoma, the IoT chip maker Advantech, Broward County Public Schools, and the Scottish Environment Protection Agency. The big attack that recently captured everyone's attention was their infiltration, data exfiltration, and subsequent

encryption of Ireland's Health Services Network. That was actually on the radar when I did the podcast last week, but it didn't really stand out except for the fact of what it is.

Ireland's Health Services Network is known as HSE, for Health Service Executive. It's Ireland's publicly funded healthcare system, which was forced to shut down in its entirety Friday before last. The Irish National Health Service said: "We have taken the precaution of shutting down all our IT systems in order to protect them from this attack and to allow us to fully assess the situation with our own security partners." That IT outage led to widespread disruption across the country for all of its healthcare, resulting in limited access to diagnostics and medical records, transcription errors which occurred due to the use of handwritten notes, and slow response times for healthcare visits.

Naturally, in a system which had become dependent upon IT, if you lost all your IT, you were in trouble. We found a snapshot of the dialogue which then ensued between HSE and the bad guys. They said: "As you already know, we infiltrated your network and stayed in for more than two weeks, enough to study all of your documentation, encrypted your file servers, SQL servers, downloaded all important information with a total weight of more than 700GB - personal data of patients (home addresses, phone numbers of the contact), employees (home addresses, employment contracts, scans of personal documents, phone numbers), contracts, customer bases, consolidated financial statements, payroll, settlements with partners, bank statements. The good news is that we are businessmen. We want to receive ransom for everything that needs to be kept secret, and don't want to ruin your business. And amount at which we are ready to meet you and keep everything as collateral is $19,999,000."

Leo: We are businessmen. We just want to tell you...

Steve: We are greedy, greedy businessmen.

Leo: Maybe their definition of businessmen is not quite the same as mine.

Steve: Wow, $19,999,000. So for some reason...

Leo: Oh, oh, a dollar off.

Steve: Yeah, well, a thousand dollars off.

Leo: Thousand dollars off.

Steve: One thousand less than 20 million. So apparently, sensing that they had hooked a big fish, the attackers demanded, as this note says, $1,000 shy of $20 million ransom. Now, get this. The Prime Minister of Ireland said they would not be paying any ransom. Then, as part of their ongoing negotiations - and perhaps in part due to the long shadow just recently cast by the DarkSide attack which shut down, of course, as we know, the Colonial Pipeline - last Thursday the gang behind Conti, these Wizard Spiders, posted a link to a "free," as they put it, "free" decryptor. Get your free decryptor, which will work for all of the HSE encrypted data. Researchers have confirmed that the provided tool will decrypt the files, though it is still being inspected more closely to make sure it doesn't contain any other malicious content.

Okay. However, the bad guys still insist upon being paid just shy of $20 million ransom, threatening to go public with their trove of extremely confidential Irish citizen health records data. They said: "We are providing the decryption tool for your network for free. But you should understand that we will sell or publish a lot of private data if you will not connect us and try to resolve the situation."

Okay. Now, in an odd-seeming bit of whimsy, the High Court of Ireland issued an injunction against the Conti ransomware gang, demanding that the 700GB of stolen HSE data be returned and not sold or published. What? The High Court has issued an injunction against a hidden, probably Russian, cybercrime organization? To what end?

The injunction was received by the HSE against the Conti ransomware gang from the High Court of Ireland. But without any formal method to service the Court's order, government representatives uploaded it to the Tor dark website associated with the gang, I guess as their means of serving them. The order prohibits the attackers from publishing, selling, or sharing any of the stolen data with the public. And then, get this. It also demands that they return the stolen data and identify themselves by revealing their full and true names...

Leo: Identify yourselves. You vagabond.

Steve: ...email addresses, and physical addresses.

Leo: Sure, why not?

Steve: Which, you know, yeah, exactly. You can ask. Which led me to wonder what sorts of mind-altering substances might be in use by the High Court. This feels like something more than a long afternoon in an Irish pub.

Okay. So what's up with this apparent loss of sanity? No one believes for a moment that the Conti gang will acquiesce to the demands of the High Court of Ireland, no matter how sincere they might be. But there is some outside hope that the country providing cover and domicile to the attackers might be willing to track them down - uh-huh - and at least prevent them from leaking the stolen data. You have to imagine that Russian intelligence knows exactly where these guys are holed up.

But the next day, last Friday, the Irish Times Security & Crime Editor, Conor Lally, explained in a pair of tweets that the injunction is never intended to prevent the Conti gang from leaking the data. Rather, it was issued to prevent the press, or anyone else, from publishing the contents of any stolen data if it were to subsequently be leaked by the ransomware gang.

In Conor's first tweet, he said: "The injunction is not a super injunction in the traditional sense. We can report about it and report if data is published. The injunction is designed to stop/limit the distribution of the data/docs in Ireland after they are published by the attackers." And then in his follow-up tweet he said: "The injunction doesn't stop the media reporting if the attackers leak the data, which one assumes is likely. It just means media, and anyone else, cannot publish/share the actual documents when they are leaked."

So I checked just, I think it was this morning again, to see if there had been any update since late last week. There was some buzz about some disclosure several weeks ago, but it seems unlikely because there isn't any, you know, nothing has happened since, and

this all happened afterward. So anyway, we'll see what happens with this gang. But Conti is now in the big-time, and they've gotten the attention, presumably deliberately. I don't know if these guys are a Ransomware as a Service. I don't think they are. So there's no middleman that they're dealing with. And we will be talking about the whole RaaS stuff here shortly.

I did, however, want to take note of literally a mega ransom. A huge firm, CNA Financial, based out of Chicago, has paid up big-time. And I'll also note that insurance companies in general have begun backing away. The U.S. insurance giant CNA Financial reportedly paid $40 million to a ransomware gang to recover access to its systems following an attack back in March. And this registers as one of the biggest ransoms paid to date. This was reported first by Bloomberg, citing "people with knowledge of the attack."

The adversary that staged the intrusion is said to have allegedly demanded 60, six zero, million a week after CNA, which as I said is based in Chicago, began negotiations with the hackers. So they talked them down 20 million, which, you know, is good. And that culminated in the payment two weeks following the theft of company data. In a May 12th statement, CNA Financial said it had "no evidence to indicate that external customers were potentially at risk of infection due to the incident." Likely there was exfiltration, and they are assured that the data will go no further.

The attack has been attributed to a new player on the scene, the "Phoenix Cryptolocker," according to a BleepingComputer report back in March at the time. The strain is believed to be an offshoot of WastedLocker and Hades, both which are known to have been used by the Russian cybercrime network, Evil Corp. A year and a half ago, back in December 2019, U.S. authorities sanctioned Evil Corp, we talked about this at the time, and filed charges against its alleged leaders Maksim Yakubets and Igor Turashev for developing and distributing the Dridex banking trojan to plunder more than $100 million over a period of 10 years. Law enforcement agencies also announced a reward of up to $5 million for information leading to their arrest. Today, both remain at large.

And remember that last October 2020, the U.S. Department of Treasury issued a "guidance," as they called it, warning of penalties against companies making ransom payments to any sanctioned person or group. This prompted ransomware negotiation firms to avoid dealing with blocked groups such as Evil Corp. The Treasury Department said: "Companies that facilitate ransomware payments to cyber actors on behalf of victims, including financial institutions, cyber insurance firms, and companies involved in digital forensics and incident response, not only encourage future ransomware payment demands, but also may risk violating the Office of Foreign Assets Control regulations."

So the surge in ransomware attacks has always had an impact on the cyber insurance industry. And this is something we have been expecting. One firm, AXA, announced earlier this month that it will stop reimbursing clients in France should they opt to make any extortion payments to ransomware cartels. This underscores the expense of ransomware indemnification, where insurance firms grapple with successfully underwriting ransomware policies while confronting the rising payout costs that threaten their bottom lines.

And to that end, a report just released by the U.S. Government Accountability Office (GAO) last Thursday revealed that the soaring demand for cyber insurance has driven insurers to raise premiums while limiting coverage. The amount of total direct premiums written jumped up 50% from 2016 to 2019, from $2.1 billion to #3.1 billion. And those numbers of course are two years old now. So one imagines the rate of payment is even higher. The GAO noted that: "The continually increasing frequency and severity of cyberattacks, especially ransomware attacks, have led insurers to reduce cyber coverage limits for certain riskier industry sectors, such as healthcare and education

[unfortunately] and for public entities, and to add specific limits on ransomware coverage."

So we've been watching this unfold on this podcast. It was inevitable, and now it's happened. So when they say "IoT," do they mean us?

Leo: UoT.

Steve: A headline at Threatpost from last Wednesday caught my eye. It read: "Keksec Cybergang Debuts Simps Botnet for Gaming DDoS," with the subhead "The newly discovered malware infects IoT devices in tandem with the prolific Gafgyt [spelled G-A-F-G-Y-T] botnet, using known security vulnerabilities." So I thought, IoT devices. What IoT devices exactly? Are they some obscure widget that no one we know uses? Turns out no. They're referring to the incredible number of Linux-based NAT routers that virtually everyone is using.

The Threatpost piece goes on to talk about a recently developed botnet named Simps, S-I-M-P-S, which has emerged from the cyber underworld for the purpose of carrying out DDoS attacks aimed at gaming targets and others. It's hosted on other people's consumer routers and forms part of the toolset being used by this Keksec cybercrime group. The Simps botnet was first spotted in April being dropped on IoT devices, and by that I mean NAT routers, by the Gafgyt botnet. And I suppose one good botnet deserves another. So this was one botnet planting a second one.

Gafgyt is a Linux-based botnet that was first seen seven years ago, back in 2014. It targets vulnerable IoT devices such as routers made by Huawei, Realtek, Asus, and Dasan's GPON home gateway devices, of which a quarter million are on the 'Net. In other words, not obscure, unsupported light bulbs somewhere, but our NAT routers. And in the present infection campaign, this Gafgyt botnet is compromising, looking for and compromising Realtek and Linksys endpoints. It then fetches and installs the Simps bot using Wget. Simps itself then uses the Mirai and Gafgyt modules for its own DDoS functionality.

So our takeaway here is these routers, I mean, our NAT routers are only being discovered because they are in some way responding to incoming packets. Way back in 1999, on October 8th, my god, nearly 22 years ago, some guy named Paul Thurrott wrote the very first article for a site called WUGNET, the Windows User Group Network. Paul's article was titled "Protect your Windows PC with Steve Gibson's ShieldsUP!."

Leo: Nice.

Steve: Paul was the first person to write anything.

Leo: Really? Oh, that's awesome.

Steve: Yup. So the case I made back then, when I was, as far as I know, the first to use the term "stealth" to refer to an Internet-connected device that did not respond in any way to incoming probes, was that it was worth deliberately violating a de facto rule of the Internet, which was that all devices having TCP/IP stacks must respond to a ping, and that closed ports should respond with either a TCP RST or an ICMP Port Unreachable.

That nearly 22-year-old advice has aged well. It's withstood the test of time. So when they refer to "IoT devices," they do mean us.

At my other location I have an Asus router. It's not on the front line. It's safely positioned behind a FreeBSD pfSense router since I need features such as pfSense's powerful static port translation and its IP-based incoming packet filtering. But after a tweet I received this morning, I'm now excited to get home to update my Asus by hand because it may be the last time I ever need to do so. I'll have more to say about that when we get to this week's listener feedback. In the meantime, please, please, please make absolutely certain that the routers you're responsible for and the routers of those you care about do not have any connection-accepting ports statically exposed to the Internet. Any appearance of convenience that something that might be listening might be offering is just not worth the risk.

So we have a new term introduced: Mean Time To Inventory. Everyone's familiar with the abbreviation MTBF, right, Mean Time Before Failure. Now the industry is coining a new abbreviation, MTTI, Mean Time To Inventory. This refers to the startling speed with which bad guys have begun to scan the Internet for newly released vulnerabilities after that vulnerability's first public announcement. In this case the term "inventory" refers to them adding penetrated devices to their "inventories."

Now, in both cases, MTBF and MTTI, we'd like them to be as long as possible. But at least in the case of this new MTTI, it turns out that a study recently released by Palo Alto Networks' Cortex Xpanse research team reveals for the first time just how startlingly short today's MTTI actually is. They frame their research by explaining: "Malicious actors are opportunistic predators, constantly searching for vulnerable targets. Unfortunately, adversaries are much faster at finding vulnerable assets to attack than defenders are at finding those same assets to secure. It's not just an arms race," they write, "in terms of conducting cyberattacks and protecting against them. There's also a sprint to detect systems vulnerable to cyber threats."

They said: "To help enterprises gain ground, the Palo Alto Networks Cortex Xpanse research team studied the public-facing Internet attack surface of some of the world's largest businesses. From January to March of 2021, we monitored scans of 50 million IP addresses associated with 50 global enterprises, including a subset of the Fortune 500, to understand how quickly adversaries can identify vulnerable systems for exploitation. In this report, we share our key findings, information on the top threats in attack surface management, and insights on how to ensure your organization remains secure."

Okay. So we've talked about how this race to patch versus race to penetrate is happening for quite a while now. But we've been lacking in metrics. What the Cortex Xpanse team found was that potentially juicy zero-day vulnerabilities can prompt attackers to begin scanning within as few as 15 minutes following first public disclosure. Now, think about that. This is not just a scan for a port. This is a scan for a specific vulnerability that they didn't know about 15 minutes earlier, located at a specific port. That means that attackers are writing and deploying custom scanning code within 15 minutes in order to be scanning for not-yet-patched vulnerabilities that fast. This really does change the landscape for serious remotely exploitable vulnerabilities.

And get this. The researchers noted that attackers worked even faster when it came to Microsoft Exchange, with first vulnerability scans detected no more than five minutes after the release by Microsoft of the patches. Now, recall that I talked about how it seemed clear that the bad guys must already have a mature database. All these different groups must already have databases indexed by port and probably also by what's known to be answering queries at that port for the entire Internet.

So the moment a new vulnerability appears, for instance in Exchange Server, they're not scanning all 4.3 billion IPv4 addresses. They're able to immediately extract the list of all known Exchange servers currently accepting connections over SMTP, POP, and IMAP, whatever the vulnerability is listening on, in order to then immediately begin exploitation. And today this literally happens in the blink of an eye. When this happened to Microsoft Exchange, the researchers over at F-Secure commented that vulnerable servers were being hacked faster than they could count.

It's also been noted that the general availability of inexpensive cloud services has helped not only well-established APT groups known to be using them, but also smaller cyber criminal groups and individuals, to take advantage of new vulnerabilities as they surface. This Cortex Xpanse group report notes: "Computing has become so inexpensive that a would-be attacker need only spend about $10 to rent cloud computing power to do an imprecise scan of the entire Internet for vulnerable systems. We know from the surge in successful attacks that adversaries are regularly winning races to infect systems before they can be patched against new vulnerabilities."

And, interestingly, the Palo Alto Networks group's research also highlighted that, not surprisingly to any of our listeners, remote desktop protocol, RDP, was the most common vector for security intrusions among enterprise networks. It alone accounts for one third of all security problems overall. The report says: "This is troubling because" - well, it's just troubling, period. But they said: "This is troubling because RDP can provide direct admin access to servers, making it one of the most common gateways for ransomware attacks."

And I don't know if I mentioned, actually, publicly on the podcast that there are instances where - there are databases of compromised RDP login credentials which are for sale on the dark web. So affiliates of Ransomware as a Service who aren't themselves super skilled at engineering intrusions, they'll buy a credential and use it to log in, and then begin their penetration. And you know, if it weren't for the blessed pervasive ubiquity of NAT routers behind which all of us with Windows systems are able to hide, the world as we know it today would have already ceased to exist. Can you imagine if the world's entire inventory of remotely accessible devices weren't just enterprises who have to have a public presence, but were also every single last powered-on Windows machine and IoT device.

The nutty IP purists, with their heads well positioned far up their you-know-whats, where the sun don't shine, have always decried the use of NAT. They say that the Internet was designed for every device to be directly addressable and accessible to every other. Well, thank god that never happened. Just because every IPv6 user will be receiving their own personal 64,000 IPv6 address space, don't ever consider directly mapping those external IPs through to your devices on the inside. It's already bad enough that Microsoft gave us UPnP so that Xboxes could autonomously solicit incoming traffic. The last thing we need is to step out from behind the protection of those billions of little hardware firewalls that everyone is using today. With a mean time to inventory numbered in the low minutes, none of us would stand a chance. So, whew.

Okay. Two bits of fun.

Leo: Yabba Dabba Doo.

Steve: Oh, and the sale of SpinRite 6. Very cool. Okay, Leo, you're going to want to go and test your skill. It took me an embarrassing number of times to prove I was human: grc.sc/820. It is our Shortcut of the Week.

Leo: I hate these CAPTCHAs. I just hate them.

Steve: Well, this is the Doom CAPTCHA. It's a joke.

Leo: My doom or your doom? Oh, my doom, huh? Okay.

Steve: Well, no. It is the game Doom.

Leo: Oh. Oh.

Steve: Yes.

Leo: Okay. So kill four enemies. I think a computer could do this very easily. One, two, I mean, really, how hard is this? Oh, game over.

Steve: Yeah, you didn't do it quick enough, Leo.

Leo: You've got to do it fast. Computer can do it slowly. I see. Whoa. Oh.

Steve: Okay, see, now...

Leo: You know what? I like that.

Steve: This tells me you really have been spending a lot of time shooting stuff.

Leo: I spend a lot of time playing Doom.

Steve: It took me, like, 10 tries.

Leo: What?

Steve: To get that green checkmark.

Leo: One, two. Oh, I see, the red progress bar is my time.

Steve: Yeah.

Leo: I get it, yeah.

Steve: Yeah.

Leo: Yeah. You don't like this, then. I get it. I take it.

Steve: No, I just thought it was just - it occurred to the guy last Saturday morning.

Leo: That's funny.

Steve: He coded it by the end of the day. It made #1 Product of the Day over on Product Hunt.

Leo: Product Hunt, yeah. It's just a little embed. Look at that.

Steve: Yeah.

Leo: So simple.

Steve: It's just a cute little thing. So anyway, I just - I wanted to share it with our listeners. I thought, I knew that a lot of us old-timers would recognize that and get a kick out of it.

Leo: Oh, wait a minute, I didn't have the sound turned on.

Steve: Oh, yeah, yeah, yeah.

Leo: Is that Doom sound in this?

Steve: Oh, yes.

Leo: Oh, yeah? Let's see. Oh, yeah, baby. Oh, yeah, baby. It's Doom. Okay. Wait a minute. But I don't mind because I got the Doom sound effect. That's hysterical. I wonder how Carmack is at this. That's great. How fun is that? It's just hysterical. Love it.

Steve: Okay. A little less overtly entertaining, but this is the one that I just find myself staring at. I know this is completely random, but a few months ago I looked at Windows 10's built-in screensavers and realized that as far as Microsoft appears to be concerned, screensavers have fallen into disfavor. Or perhaps they've migrated over to the Windows Store. I don't know. I didn't go there. Since I often leave my Win10 machine on and unattended, like during dinner, when I'm taking a walk with Lorrie after dinner or whatever, and since my workstation is in our family room so that I'm able to be working while still being nominally present, clanky mechanical keyboard notwithstanding, I

decided that I wanted something fun on the screen when I wasn't using the machine. So I went looking for a satisfying screensaver.

Several months later now, I'm enjoying what I found, so much that I wanted to share it with my listeners, who often tell me that my taste matches theirs. Maybe that'll happen again. A developer named Terry Welsh has written a collection of open source, open GL screensavers. He wrote them for Windows, but this one and most others have also been ported, because they're open source, to OS X and Linux. This one is called Helios, which is the one that for some reason I find transfixing. So they're all at ReallySlick.com, that's Terry's site, ReallySlick.com/screensavers.

But I have to say I was not a fan of Helios's default settings. Out of the box, it was way too busy for me. So it needed some tweaking to match my taste. And I've captured the settings panel I use with it so that others can see it the way I see it. So it's in the show notes for anyone who's interested. So, let's see, I have ion size at 10. Number of emitters, I left it at 3, I think that was the default. Number of attracters at 3, I think that was the default. Animation speed is 10. Camera speed is 1. I think I slowed that down to 1. Motion blur is 50. I think I turned that up to 50. And then the frame rate has no limit. It's at zero.

Anyway, see what you think. Watch it for a while. You can download a zip of his that contains all 12 of them. And then if you just drop it in your Windows System32 folder and then go to select your screensaver, it may be set to none right now, you can see all of them and poke around. So there is another one that I just started looking at, Microcosm, which Lorrie likes a lot better. It produces some really cool-looking kind of liquid 3D objects which, you know, if that's more your thing, check out Microcosm. Anyway, just to point people there. You know, see what you think.

Okay. Some closing the loop stuff from our listeners. JP wrote: "On today's Security Now! you discussed Tor and HTTPS." He says: "By definition, Tor hidden sites are not going to use HTTPS. Getting a cert would make mockery of being hidden. And a self-signed cert these days just pops up flags."

So JP's assertion was interesting to me since I had never looked into the issue of Tor .onion sites themselves using HTTPS. JP was a bit confused about our previous discussion, however, since in that discussion of the problems with Tor exit node security we were talking about non-HTTPS connections being made through Tor to external websites on the Internet, not internal .onion sites. But that left the interesting question about obtaining TLS certs for .onion Tor hidden sites. It seemed to me that any ACME-based TLS certificate issuer such as Let's Encrypt would be what one wanted; right? I mean, you don't want to identify yourself. But you do want to assert your domain name, which is what ACME-based automation of TLS certificates allows.

So it turns out that, until recently, only EV, believe it or not, Extended Validation certs, could be issued for .onion domains. And it was truly by coincidence that it appears DigiCert, as we know, my chosen certificate authority, appears to be the choice for EV .onion certs. It was initially unclear to me why EV was required. And it would seem that needing to authenticate oneself to the level required to obtain an Extended Validation certificate could hamper some of the value provided by .onion domains. But as we'll learn in a moment, there was a rationale behind that.

But first, I discovered that .onion domains have a long history of HTTPS access. Seven years ago, when Tor users would attempt to visit their Facebook accounts, Facebook's geofencing security would trigger to lock that user's account because the traversal through Tor would cause the user to appear to be connecting to their account from some foreign land. To fix this problem, while also allowing Tor users to have a better experience when connecting to Facebook, way back in 2014 Facebook launched the

dedicated Tor address https://facebookcorewwwi.onion. Using this SSL/TLS authenticated and encrypted onion site, Tor users could access the site directly without fear that doing so might freak out Facebook's geo-aware security.

So yes, traditional security certificates have long been available for Tor's .onion sites. However, I also found a CAB, you know, the CA/Browser Forum, a CAB Forum ballot, where certificate issuers, 15 different certificate authorities, and four customers - Apple, Microsoft, Google, and Mozilla - voted unanimously, with a few abstentions, but no nays, in favor of opening up DV and OV certs to .onion domains. And Let's Encrypt was among those voting in favor of having this happen.

The explanation of the purpose of the ballot was also quite informative. It says: "This ballot will permit CAs to issue DV and OV certificates containing Tor onion addresses using the newer version 3 naming format. In ballot 144, later clarified by ballots 198/201, the Forum created rules for issuing EV certificates containing onion addresses. A primary reason for requiring EV level validation was that onion addresses were cryptographically weak, relying only on RSA-1024 and SHA-1. More recently a newer 'version 3' addressing scheme has removed these weaknesses. For much the same reason that EV certificates are not always a viable option for website operators, for example, sites operated by individuals" - or I would argue sites wanting to be secure, but not identify themselves - the ballot says "many onion sites would benefit from the availability of Domain Validation and Organization Validation certificates for version 3 onion addresses." Which, by the way, is now the standard in Tor.

They said: "The Tor Service Descriptor Hash extension required in the EV Guidelines to contain the full hash of the keys related to the .onion address is no longer needed as this hash is part of the version 3 address. Older version 2 onion addresses are still in use, so this ballot does not remove the existing EV Guidelines requirements for onion names."

So this balloting occurred back in February of 2020, so more than a year ago. I haven't been able to locate any evidence of anyone other than DigiCert issuing .onion certs, and those EV. But it appears that Let's Encrypt is all onboard for this, so I would imagine that it's just a matter of time before this starts happening, if it hasn't yet. But again, I mean, I didn't, like, scour the Internet. But a quick search didn't turn up anything.

I mentioned going home to perform maybe the last manual update I would ever need to of my Asus router. Mikael Falkvidd, who is a friend of mine, he was one of the hosts of the SQRL presentation during the trip I made around Europe. I found a tweet from him this morning saying: "Hi. I just wanted to let you know that our Asus routers now support auto update. This feature was not included in the release notes, but a fix for the frag attacks were included." And he gave a screenshot of the page, which I included here, showing the ability to turn auto updates on and specify what time of the day, typically the wee hours of the morning, when you give the router permission to update itself.

And there is some talk about the way to recover in the event of a failed update. If that updates, and if it's a fallback to the previous firmware, both fail, then there is a recourse for the user. But, you know, bravo. Let's hope this becomes standard operating procedure. It's, you know, way easier for Linux-based routers because they have a lot more resources to work with. What we need is for our light bulbs and our electric plugs to be able to do that, too.

Jared Stein tweeted: "Steve, since you always talk about science fiction books, I was wondering if you could suggest a starting book, one that would either lead me into continuing to read it or determine it's not for me." He said: "I do really appreciate Security Now! and all you bring to the community. Regards, Jared."

Well, okay. So there are as many science fiction authors as there are musicians, and taste for the work of this or that author is probably as personal as taste in music. But one of my favorite authors is, as we know, Peter F. Hamilton. And there goes another sale of SpinRite 6. Thank you. One of my favorite authors, as I was mentioning, is Peter F. Hamilton and his "Fallen Dragon." It's a standalone novel, a great example of his imagination. It is a little militaristic, so that's okay. But it's not based on military, it's just it's a great read. And if you like that one, then his pair of novels, "Pandora's Star" and "Judas Unchained," would be next in line. He then gets into trilogies. And, boy, Peter has never written a short book. I guess actually there is a novella. Novella? Novella.

Leo: Novella.

Steve: Yeah, or two. But generally they're really satisfying, and they are really long, but really fun. Also I've got two tweets regarding "The Martian," which was going to remind me to ask you about how the Andy Weir interview went on Friday.

Leo: Oh, it went great. He's such a cool guy. I just love him. And, yeah, highly recommend it. We put it out on the Triangulation feed, as well as the TWiT events feed. So it's about an hour and few minutes of Andy talking about, well, the first half is spoiler-free. And then we put up a big sign that says, "Spoilers."

Steve: Good, good, good, good, good, good.

Leo: Because I had questions I wanted to ask him about the book specifically. So, but yeah, we do at least half an hour spoiler-free. Listen to the first half hour, read the book, then listen to the second half hour. It's easy.

Steve: Very cool. Kerry Blue Life tweeted: "Steve, thanks for 'The Martian' recommendation years ago. Andy Weir has done it again with 'Hail Mary.'" He said: "I was surprised how emotional it could be." And Leo, did you finish it? Because you had, like, two hours left, I think.

Leo: Oh, yeah. I finished it before I talked to him, yeah, yeah, yeah.

Steve: Okay. And then Craig tweeted: "Thanks for 'Hail Mary.' Work on Monday is so much closer now <sigh>." I guess he listened to it all weekend. And he said: "BTW, great narrator, same as the Bobiverse books." He says: "I kept waiting for Bob or piggies to show up." Now, I have no idea who Bob and the piggies...

Leo: No idea either, yeah.

Steve: Bob and the piggies. At least we know that "Hail Mary" is really good.

Leo: Really good. Yeah, the guy, I asked him about R.C. Bray, who read "The Martian." And I guess he just wasn't available. But the guy they have read "Hail Mary," you don't listen to audiobooks, so you don't care.

Steve: No.

Leo: But for those who do, Ray Porter does a fantastic job, I would say as good as R.C. Bray.

Steve: Very cool.

Leo: Yeah. And the R.C. Bray version of "The Martian" is gone. I don't know why. Wil Wheaton reread it.

Steve: And given the power of Andy's pull, I would imagine he could get anybody, I mean, that Audible could get anybody they wanted to do it.

Leo: Yeah. He did offhandedly mention, I did not pursue it, that R.C. and Audible didn't always see eye to eye. But I don't want to read too much into that. He also mentioned they've already sold the movie rights to "Project Hail Mary."

Steve: Yay.

Leo: And that Ryan Gosling will play the lead in it.

Steve: Oh, my god, they've cast it already.

Leo: Oh, yeah. Lord and Miller are directing it. Yeah, I think after - "The Martian" almost won an Oscar. "The Martian" was, you know, a huge success. In fact, it was, you know, you think about Ridley Scott, the guy who directed it, and all the great movies he's done like the "Aliens" movies, "The Martian" was his highest grossing movie ever. So, you know.

Steve: And you know, it's going to - had popular appeal.

Leo: It really did.

Steve: Yeah. And I will say the book was so much better. I read the book before seeing the movie on purpose.

Leo: It's always better to read the book first, I think, yeah.

Steve: Yes, yes. And "Jurassic Park," same way. There were so many things, it's like, wait, wait, wait, you left this out. It's like, oh.

Leo: Especially for me with science fiction. I always prefer to read the book. Either the book's going to spoil the movie or the movie's going to spoil the book. But I just feel like the book is always going to be the true vision.

Steve: Oh, "Ender's Game," so much better to read.

Leo: Oh, much better in the book. And "Ender's Game" has a twist, which will be spoiled for you. So read the book. Don't watch the movie. It's much better to be spoiled by the book than it is by the movie. Let's put it that way.

Steve: Yeah.

Leo: He said he's got a percentage of the gross for "Project Hail Mary," so he'll be doing okay.

Steve: Good for him.

Leo: You always want to do that. He's a producer on the film.

Steve: Good for him.

Leo: And Lord and Miller are - they directed "The Lego Movie," the "Jump Street" movies, the "Spider-Man: Into the Spider-Verse." So they have a good track record, mostly with animated movies. But we'll see. I think it'll be a good movie. I'm excited. It's a great book. Read the book. Always read the book.

Steve: Well, there's no way they're going to put lame directors in charge of a movie with that much potential.

Leo: Exactly.

Steve: So, yeah.

Leo: Yeah, yeah. And you're absolutely right. You should read the book first, yeah.

Steve: Okay. I think this is my last little bit of feedback. But it was an interesting question. Mike Lawrence tweeted: "@SGgrc Have you given any thought to vaccine e-passport verification/privacy?" He says: "One concern is verification can induce a record for tracking activities, but I feel like there's a public key solution to that." And yes, there is indeed a public key solution to that. In fact, there's a public key solution to most things.

So I played around with this a bit yesterday. The largest possible standard QR code for 8-bit binary data can encode 3K bytes. I took a headshot of myself and used JPEG

compression to reduce its size to 3K and pasted it into today's show notes. It's entirely recognizable as me. A QR code's digital data could be signed with a government's private key and subsequently unspoofably verified with the government's matching public key. So, for example, a credit card-size vaccination ID could be created containing its holder's photo, its signed QR code, and its signature, side by side.

Then, when needing to prove vaccination status, say at an arena, or maybe for indoor dining, or perhaps even dating, people could present their card to the ticket agent, the receptionist, or their date. Upon scanning the card, the validity of the QR code, that is, the QR code's signature, thus the QR code, could be instantly and locally, that is, without any communications, verified without any need to phone home. And the QR code would also present the user's face on the verification screen for comparison to what's shown on the card and with the face of the person presenting the card.

So to be clear, I'm not proposing that this is what we or anyone should do. I'm just noting that today's crypto technologies provide us with such a flexible toolkit that they can be used to provide solutions to nearly any problem.

Now, Mike also asked about privacy. The system I've described is able to validate that a visually recognizable and digitally scannable image of an individual can be signed using a private key, and that signature can therefore be authenticated without any communications. But to be more useful moving forward, if a query were allowed to a trusted, non-tracking, privacy advocacy group, then the individual's relevant vaccination history could be retrieved, presumably without any logging, to make this single card also useful in the future for any possible vaccination boosters or next pandemic vaccines and so forth.

So anyway, just a proof of concept that, yeah, as a country it's not clear what's going to happen in the future. Looks like hopefully we're going to get out of this current COVID-19 mess without the need of any proof of vaccination. We'll see how that goes. But for what it's worth, it would be possible to create a separate health status card whose signature could be verified locally with an optional check to see if there are any updates.

Okay. So I wanted to finish up this week by making an observation about something that's going on in the dark underworld. There are now around 20, maybe it's 21 since we began the podcast, individually identifiable ransomware groups with most now, though not all, operating under the Ransomware as a Service model because it's proving to create additional value through specialization of function. And thus arises the problem of the responsibility to pay affiliates their share of a ransom after payment has been received from the victim.

In the case of DarkSide's decision to abruptly shutter their operations and/or being forced to do so, we don't have full visibility into exactly what went on there, affiliate postings have begun appearing on dark web forums complaining about nonpayment of affiliate commissions which have been earned, so to speak, and which are now due. And yeah, I agree, "Oh, boohoo." But the existence of well-known dark web forums means that this is now happening in plain sight for all RaaS operators, you know, Ransomware as a Service, RaaS operators, and current and would-be affiliates to witness. Service operators want to attract affiliates, and affiliates want to be reliably paid. In this environment, the presence of dark web forums means that we now have a marketplace with low-friction information flow and communications. This in turn means that three things are likely to happen.

First, the notion that not all RaaS operators are the same will evolve. RaaS operators will begin to acquire a reputation for reliable and timely payment. And since to a large degree ransomware is ransomware and cash is king, reliability of payment will largely dictate future affiliate choices, and affiliates can be expected to be extremely fickle. Any mistake

or payment dispute will instantly doom any RaaS service. It's obviously over for DarkSide. No one would ever trust them again, especially as the underworld is witnessing the many complaints about nonpayment of earned commissions.

Now there's competition among RaaS service providers. So I expect that the second thing we're going to see is some jockeying over commission rates. The only thing that affiliates care about is money. So as ransomware matures, we can expect to see commission rates settle and mature, too. Those services which have built the best reputations will be able to charge a higher price for the use of their ransomware by taking a larger piece of the pie, and newer upstart services which are not yet established and trusted may need to lure new affiliates to use their product by offering the use of their ransomware at a greater discount than the more well-established operators.

Hooked a big fish? Want to keep more of what you're about to earn? Consider encrypting your target's network with Newbieware. We only take 5%. We also have the fastest encryption around so your target will never know what hit them. And we have big pipes, hosting your ill-gotten goods on AWS, the industry's most reliable cloud storage.

And, finally, the third thing I expect we're going to be learning about before long will be the emergence of another player in the evolving increasingly specialized multi-component Ransomware as a Service ecosystem: ransom escrow services. The publicity surrounding the DarkSide collapse is likely to bring about the emergence of a neutral intermediary to manage the money. Escrows are a time-honored system for holding and transferring valuable assets among untrusted parties. The cryptocurrency system makes it easy to have the ransom paid to a wallet that's not under the control of either the Ransomware as a Service provider or their affiliates. This helps to ensure that the affiliate will be paid, no matter what might happen to the RaaS service.

When tens of millions of dollars are at stake, and when payment might be an all-or-nothing proposition, shaving off a quarter point for escrow commission will probably seem like a wise investment. It will give affiliates the payment assurance that they will now be clamoring for, and it will allow RaaS services to bootstrap themselves by offering to escrow ransom payments as an option. Given the dynamics of this dark ecosystem, I'll be surprised if we don't learn about ransomware escrow services appearing before long.

For completeness, I should note that there have been some informal first stabs at something like this. We noted a year or two ago when the REvil ransomware gang deposited $1 million worth of bitcoin into a different hacking forum as a means of attracting affiliates. They wanted to show that they meant business. And to show that they, too, meant business, DarkSide placed 22 bitcoin on deposit with the admin of the XSS forum. But this was more in the form of a guarantee than an escrow. And 22 bitcoins won't begin to cover the sorts of ransoms we're seeing recently.

So this wasn't an escrow. As we detailed last week, DarkSide victim funds flowed directly into DarkSide's bitcoin wallet, which had been identified by Eclipse. And moreover, claims now being made against the DarkSide bitcoin guarantee stash, which the XSS forum is administering, are meeting with trouble being paid. This also suggests why that XSS admin may have said that ransomware is no longer welcome there. This admin is likely quickly getting fed up with the unwanted responsibility of being DarkSide's unpaid de facto guarantor. And one wonders what happens to any unclaimed guaranteed funds. Who gets to keep them?

It does appear that DarkSide did the best they could after they stumbled into the Colonial Pipeline nightmare. Their affiliates who successfully encrypted victim networks have all received the corresponding decryption keys which allow them to pursue negotiations with their victim companies independently, at least those who have not yet paid into

DarkSide's bitcoin wallet, though that's not what they signed up for, either. They wanted their own intermediary.

And after writing all of the above, it occurred to me to Google the phrase "cryptocurrency escrow," which returned more than 2.5 million hits. So escrowing cryptocurrency is obviously not a new concept. However, no reputable commercial cryptocurrency escrow service wants to receive a letter from Ireland's High Court demanding that they freeze funds in escrow. So the dark web will need to establish its own dark escrow services. I expect we'll be hearing about such a service before long.

Leo: You know the other thing this points out is there are some real good avenues for federal law enforcement to disrupt this because, as you point out, reputation is so important. We don't know who brought down DarkSide's servers and who brought down their payment system.

Steve: Right. But it hurt them.

Leo: It hurt them. Whether they did it on purpose or not, it hurt them. That's really an interesting pressure point that law enforcement could use. Instead of trying to catch these guys, damage their reputation. That puts them out of business, anyway.

Steve: Yeah.

Leo: Interesting. I hope we don't see a whole - somebody in the chatroom said, you know what Steve's left out is the corporate acquisitions and mergers. I'm a businessman here. I'm doing business. What's your problem? I think you're taking that a little too far, the business thing here. But we'll see. It is interesting, though, that they themselves want to pretend that they're a legitimate business.

Steve: And they say they are. It's like, hey, we're really sorry that your, you know...

Leo: We're in it for the money. It's not personal.

Steve: You know, you can't access your CAT scan right now. We want you to.

Leo: We do.

Steve: But, you know, we're businessmen. Just pay us. All we want is your money, and then we'll let you have your stuff back.

Leo: I'm in waste disposal. That's my business. Olive oil sales. It's a good business.

Steve: Yeah.

Leo: Yeah. Wow. Well, I hope that this is more of a humorous essay than it is a roadmap for the future of ransomware.

Steve: Lay down your bets, ladies and gentlemen. I'll guarantee you we're going to see the dark escrow appear.

Leo: Yeah. Wow. Steve is at GRC.com. That's his website, the Gibson Research Corporation. Lots of things there. But you start with SpinRite, the world's finest hard drive, sorry, mass storage recovery and maintenance utility.

Steve: And really, I want to thank our listeners. I wouldn't be at all surprised, Leo, that those two yabba-dabbas we heard weren't from somebody who said, I want to hear my own purchase celebrated. And so thank you very much.

Leo: I've long thought you should turn the yabba-dabbas up because - except that it would, eventually, you'd hear so many of them, it would really disrupt the show.

Steve: Let's hope.

Leo: I'm working on a little counter that I could put on the set that shows the current Club TWiT membership number.

Steve: Oh, cool.

Leo: Because I'm hoping people will want to get that spinning around.

Steve: Yeah, yeah.

Leo: You know, just as a fun thing to do.

Steve: Yeah, it's a good idea.

Leo: Let's go to the tote board. Yeah, make him yabba-dabba. He gets a yabba-dabba whenever somebody buys the world's finest mass storage recovery and maintenance utility, SpinRite. 6.0 is the current version; 6.1 is imminent. Join in on the fun. You'll get a free upgrade, and you get to participate in the development of this fantastic, much-needed tool. Steve's got lots of free stuff there, too, like ShieldsUP!, Paul Thurrott's favorite tool.

Steve: Thank you, Paul.

Leo: Isn't that great. I had no idea that he was the first to discover it. That's wonderful. When he was working for WUGNET, of all things. Let's see, what else?

There's lots of free stuff. But you should also go check out the podcast. It's also free. He's got the only 16Kb audio version of the show for the bandwidth-impaired. He has a very nice transcript he commissions from Elaine Farris so you can read along as you listen. He's also got the 64Kb audio. That's all at GRC.com.

We have audio and video at TWiT.tv/sn. We also of course put it out as a podcast so you can subscribe in your favorite podcast player and get it automatically, the minute it's available. If you do like the show, give it five stars, please. It really helps us spread the word. And I think everybody should be listening to Security Now!. It's that important. Certainly anybody who's responsible for anybody else's computing systems.

We do the show every Tuesday, usually about 2:00 p.m. Pacific, but today we were on time, 1:30 Pacific, 4:30 Eastern, 20:30 UTC. You can watch us live at TWiT.tv/live. There's audio and video there. Live chat during the show in two places, irc.twit.tv, and Club TWiT members also get access to our Discord server, which is fun, very active all times of the day or night, not just about our shows, but all sorts of stuff going on in there, including a Linux show we do on Saturdays. There's a TWiT+ feed, and of course ad-free versions of all the shows. If you've interested in Club TWiT, it really helps us out. Seven bucks a month, just go to TWiT.tv/clubtwit. And for you Club TWiT members, thank you for your support. We really appreciate it.

Steve, have a wonderful week. Have fun tonight. And I'll see you next Tuesday on Security Now!. Bye-bye.

Steve: Right-o. Bye.