SECURITY NOW!

Transcript of Episode #818

News from the DarkSide

Description: This week we look at a new (and old) thread to our global DNS infrastructure. We ask what the heck Google is planning with two-step verification, and we examine a huge new problem with the Internet's majority of email servers. We look at the reality of Tor exit node insecurity, touch on a new sci-fi novel by a well-known author, share a bit of closing-the-loop feedback, then take a look at this latest very high-profile ransomware attack from a previously low-key attacker.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-818.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-818-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Some serious security issues with the Exim email server. We're going to talk about a big infrastructure problem, the Colonial Pipeline hit by ransomware. What's it mean for infrastructure in general? And then Steve's got a Picture of the Week that's actually - I think it's an IQ test. It's all coming up next - you'll pass - on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 818, recorded Tuesday, May 11th, 2021: News from the DarkSide.

It's time for Security Now! with this fellow right here, we call him James Tiberius Gibson, the captain of the good ship Security Now!. Steve Gibson is here. Hi, Steve.

Steve Gibson: Yo, Leo.

Leo: What's up?

Steve: Once again, well, you know, I did not want to talk about DarkSide. But there was no way not to...

Leo: I think you have to. I think you have to.

Steve: There was no way not to talk about DarkSide. And what was interesting - because how much time have I spent promising our listeners that we wouldn't keep talking about ransomware? But when this thing moves from an incidental concern from IT people to something where our parents or grandparents or those who predate the Internet are like, what? Ransomware? What's that? I mean, and when it steps out to dramatically affect our infrastructure - oh, and this group has a weird twist also, like they have an ethics page posted on their site on the dark web about their intentions. Anyway, we'll get to that. There was enough interest about this, like enough insider information that our listeners would not have picked up from the mainstream media that I thought, okay, we've got to talk about that.

But this is Episode 818 for Patch Tuesday of May, which we'll be talking about next week because we have to wait to see what happens. We're going to look at a new and old threat to our global DNS infrastructure. We also ask what the heck Google is planning with their so-called two-step verification. We examine a huge new problem with the Internet's majority of email servers. Microsoft Exchange, that was March. And they're by no means the biggest player. It turns out that the biggest player, Exim, has some really bad problems. So buckle up. We're also going to look at the reality of Tor exit node insecurity, Leo, and really substantiate the statements you've been making when you're talking about our VPN sponsors, that just using Tor doesn't do the problem.

Leo: Yeah, yeah, yeah.

Steve: We're also going to touch on a new sci-fi novel from a very well-known author, share a bit of closing-the-loop feedback from our listeners, and then we're going to settle down and take a look at this arguably highest profile ransomware attack ever from what was previously a low-key attacker. We've never talked about DarkSide before. We're talking about Ryuk and all these other guys. And this player's sort of interesting. And for those listeners who haven't, well, actually, you and I, all of our conversation about our Picture of the Week was before you hit the record button. We have a picture that we're not going to explain. And we will explain why.

Leo: It's an IQ test. Actually it's not. It's a test of your educational levels, maybe. I don't know. I don't think it's an intelligence test, but it is a test. So we'll have that in a moment. All right, Steve. Are you ready for the IQ test?

Steve: So the bad news is nothing we could say, like I can't even describe this because if I were to describe it, I fear that I would say something that would provide a clue. So I'm not going to do that. Everyone knows where the show notes are. You can get them at GRC.com/sn or /securitynow.

Leo: Or if you're watching the video you're seeing it. But, yeah, for our audio listeners, we don't want to describe it, yeah, because we don't want to give it away.

Steve: Yes, yes. If you're watching it, it's onscreen right now. It's just fun. It is a two-frame cartoon, very clever, and you'll enjoy the fact that you get it.

Leo: Let me just check the poll.

Steve: And probably be annoyed if you don't.

Leo: In our Discord right now we're asking do you get the Picture of the Week. In other words, do you get the joke, because it's a joke. 16 do; 9 do not. And it's got to be frustrating for those who don't because it's just not obvious. Unless it is. It's one of those things. If you know, you know.

Steve: And for what it's worth, it's well done. I mean, it's just...

Leo: Oh, yeah. They got it all right, yeah. I know what you're talking about, yeah.

Steve: Exactly. Exactly. It was done correctly. So we'll just leave that as a puzzle for the listeners.

Leo: We'll tell you next week, how about that.

Steve: Oh, I like that. Very good. You've got one week before the spoiler hits. So see if you can take a look at the picture and test yourself.

Okay. So the best name - the best name. The best thing about this flaw is I think its name. The flaw is TsuNAME, obviously meant to be tsunami, so with a little bit of fudging of the spelling, because it's about name servers. So tsunami, or TsuNAME. And this is one of those clickbait-y stories, but it's still interesting and I think educational. When I first encountered the industry's coverage of this, with its portents of doom, I thought that some new nightmare must have been found with DNS, just when we needed Dan the most.

But when I dug into the story, I learned that it boils down to an interesting way for a domain's DNS records to be misconfigured such that when a naive, and I'll explain what I mean by that, a naive recursive DNS resolver is asked to resolve one of these misconfigured domains, that recursive server, serving as a DNS resolver, will get itself into a name resolution loop, which causes it to pound away on that domain's authoritative DNS servers without end. It turns out there's a way to put DNS into an infinite name resolving loop.

Now, if this had never occurred to anyone since man walked the Earth, it might be somewhat more alarming. But not surprisingly, this had previously occurred to the guys who built DNS. RFC 1536 - yes, four digits, it's an oldie, 1536 - published way back in October of 1993, was titled "Common DNS Implementation Errors and Suggested Fixes." So, yeah, things can go wrong, and how to fix them.

Section 2 of that RFC 1536 bears the title "Recursion Bugs." And after a bit of shortening for the podcast, it reads: "When a server receives a client request, it first looks up its zone data locally and in its cache to check if the query can be answered. If the answer is unavailable from either location, the server seeks names of servers that are more likely to have the information in their caches or zone data. The server chains this request to these known servers closest to the queried name. This process repeats until the client is satisfied.

"Servers might also go through this chaining process if the server returns a CNAME record" - we've talked about that, the canonical name record, which is an alias - "for the queried name. Some servers reprocess this name to try to get the desired record type. However, in certain cases, this chain of events may not be good," is what they wrote in 1993. May not be good. "For example, a broken or malicious name server might list itself as one of the name servers to query again. The unsuspecting client resends the same query to the same server. In another situation," they wrote, "more difficult to detect, a set of servers might form a loop wherein A refers to B and B refers back to A. This loop might involve more than two servers."

Okay. So with that bit of background, here's what the guys who reminded us what was written 28 years ago said in their published paper's opening abstract. They said: "The Internet's Domain Name System is one of the core services on the Internet. Every website visit requires a series of DNS queries, and large DNS failures may have cascading consequences, leading to unreachability of major websites and services." Okay. That we all know. They said: "In this paper we present TsuNAME, a vulnerability in some DNS resolvers that can be exploited to carry out denial-of-service attacks against authoritative servers. TsuNAME occurs when domain names are misconfigured with cyclic dependent DNS records. And when vulnerable resolvers access these misconfigurations, they begin looping and send DNS queries rapidly to authoritative servers and other resolvers." And they said: "We observe up to 5,600 queries per second."

They said: "Using production data from .nz, the country-code top-level domain of New Zealand, we show how only two misconfigured domains led to a 50% increase in overall traffic volume for the .nz's authoritative servers. To understand this event, we reproduce TsuNAME using our own configuration, demonstrating that it could be used to overwhelm any DNS Zone. A solution to TsuNAME requires changes to some recursive resolver software to include loop detection and caching cyclic dependency records. To reduce the impact of TsuNAME in the wild, we have developed and released CycleHunter, an open source tool that allows for authoritative DNS server operators to detect cyclic dependencies and prevent becoming victims of TsuNAME attacks themselves."

And they conclude with the abstract: "We used CycleHunter to evaluate roughly 184 million domain names in seven large, top-level (TLD) domains and discovered 44 cyclic dependent name server records, likely from configuration errors, used by 1,400 domain names. A well-motivated adversary could easily weaponize this vulnerability. We have notified resolver developers and many TLD operators of this vulnerability. Working together with Google" - and actually also with Cisco, I'll get to that in a second - they said, "we helped them to mitigate their vulnerability to TsuNAME."

So later in the paper they discuss their use of this CycleHunter tool and show that they found a total of 3,696 DNS resolvers which were not protecting their queries from this cyclic DNS misconfiguration. They manually tested the DNS resolvers Unbound, BIND, KnotDNS, spelled K-N-O-T, which is DNS, Quad9 and Quad1. All of those passed. But Cisco's OpenDNS and Google's DNS both got themselves caught in cyclic lookup loops. They informed both companies, and both fixed their problems quickly. That's, you know, it's an internal thing. I don't know if Cisco's OpenDNS, I mean, presumably they make that available to people, whereas Google's DNS is a service of Google, so they would have fixed that in-house. Anyway, and interestingly, DNS developers, it turns out, do need to always be, and generally are, on the lookout for DNS looping errors. They note that the changelog for the Unbound DNS resolver contains 28 entries related to looping.

So anyone doing recursive DNS needs to clearly make sure that they don't get themselves chasing their tail endlessly. Given the numbers, it seems unlikely that this would have happened. But somewhere in their report, I read the whole thing, they noted that while they were observing some range of domains, one new problem appeared. Like somebody brought up some new zone records and apparently just made a mistake. And sure enough, a new recursion problem occurred.

So this is happening from time to time. As long as resolvers don't chase their tail endlessly, but realize, wait a minute, I'm just caching my lookups, and I've just been asked to look up the same thing I was asked a moment ago, I'm in a loop. And so this is the Kobayashi Maru. I'm not going to proceed any further. In which case these occasional lookup problems, they'll probably get found because lookups will be broken if they recurse and never complete. But certainly DNS, you know, anyone operating a DNS server wants to make sure they don't have one, which is just going to sit around, I mean, you're using up your own local network bandwidth when you are making 5,600 queries per second to other servers out on the Internet. So that's not something that you want to have happen.

So what this all boils down to is that two of the industry's many DNS server families were failing to detect DNS lookup loops. And, sure enough, there were a few definitions out there that would cause those servers to become stuck. The benefit of this research is that it identified those servers and got them patched, and they did develop this CycleHunter tool of theirs to allow administrators of DNS to check up on their own DNS zone definitions for any cyclic lookup trouble. It's TsuNAME, T-S-U-N-A-M-E dot io. You can go there, and they have the full tech report for anyone who wants it, and also a pointer to their freely available tool CycleHunter to allow people to make sure that they're not stuck, and they don't have like a misconfigured DNS that could be loading down their servers without them knowing.

Okay. So I labeled this one, "Huh, Google?" Last Thursday Google's Mark Risher, R-I-S-H-E-R, their Director of Product Management for Identity and User Security, posted to the Google blog under the "Safety & Security" section an entry titled: "A simpler and safer future without passwords." Okay, now, unfortunately, that's not what his blog post addressed. And no one seems to be exactly sure what his blog was trying to say and what it did address, since it led to many confusing and misleading tech press headlines. I saw a headline, "Google wants to enable multi-factor authentication by default"; and another headline, "Google is turning on two-factor authentication by default"; and another one, "Google will start automatically enrolling users in two-step verification soon."

And on top of that, I saw many users who read this to mean that Google would be requiring the use of two-factor authentication. And I can certainly see how one might get that, you know, come away with that feeling from the confused headlines. It's also not helpful that Google has apparently decided to create a new term and abbreviation. Everyone already knows what two-factor authentication is. In fact, the headlines, generally two of the three used that, even though Google didn't, because that's what we call it. We call it two-factor authentication, typically abbreviated 2FA. But now we have Google's 2SV, which is what they're using, which is two-step verification. Okay. But if you first put in your email address, then you put in your password, then you're asked to do something else, aren't we already up to three steps of verification?

Leo: That's a good point. They just don't want to act like it's two factors; right? They just want to say it's another step; right?

Steve: Yeah, okay. But if you need to go get your phone, arrange to unlock it with your identity, then respond to a prompt or a text message or a one-time password, we're up to about six or seven steps by that point. I've lost count. Anyway, so I read through Mark Risher's blog posting, and here's the problematic paragraph that no one is quite sure how to interpret. He wrote: "Today we ask people who have enrolled in two-step verification (2SV) to confirm it's really them with a simple tap via a Google prompt on their phone whenever they sign in." Here it comes. "Soon we'll start automatically enrolling users in 2SV if their accounts are appropriately configured." Uh, what? So I have no idea what he means when he says "We'll start automatically enrolling users in two-step verification if their accounts are appropriately configured." What does "appropriately configured" mean?

Leo: [Mumbling]

Steve: Yeah, huh?

Leo: I wish they were clearer on this. I mean, they in their minds know what that means, but they haven't told us.

Steve: Well, yes. And you're reading my mind, Leo, because in the show notes I wrote, "And that's the problem. It apparently means something to Mark. But it's gobbledy-gook to the millions of people who read Google's blogs, and also apparently to the tech press, which tried to write news stories around it."

Okay, now, as we all know, you either have second-factor authentication enabled for authentication to your Google account, as I do, or you don't. There's no third setting labeled, "Well, I'm open to the idea, hit me up when you want." We don't have that. So the only thing I can figure is that, I don't know, Mark woke up last Thursday, and his calendar told him that it was World Password Day, as indeed it was. So he thought, oh, crap, that's right, I'm Director of Product Management for Identity and Security. I'd better think of something to say. So he banged out that confusing blog post to the world.

I think what we need to take away from his aberrant posting is that Google is a fan of using more than just our email address and password for authentication. We know that's true. And that in the interests of their users they plan to arrange to somehow encourage more of their users to add a second factor. Or, as they put it, "another step" to their logons. But as for what Mark wrote last Thursday to celebrate World Password Day, I have no idea what he could possibly mean by "automatic enrollment in 2SV," two-step verification, nor does anyone else at this point. Maybe they don't know. But looking at just what a mess this caused out in the press, if they thought that removing FTP support from Chrome might cause a ruckus, just watch what happens if they start surprising their users with the presumably unwanted additional complexity of two-step verification, which sounds like it's more like six or seven steps. I guess we're going to find out. And Leo, thinking about this further, they do have, they've authorized Android phones to get involved in a simple authentication cycle; right?

Leo: I think they have - that's what they're talking about is single sign-on. It works on an Apple phone, too, by the way, if you have the Google app on your Apple phone.

Steve: Right.

Leo: And that may be what they're thinking. And that's what Microsoft uses, that single sign-on, which is great. I love it.

Steve: So would they be aware of it for users, but see that a user has the app and then say, hey, happen to know you're an Android person.

Leo: Exactly.

Steve: So you can do this.

Leo: They wouldn't even say that. They just would start using it.

Steve: Really.

Leo: Yeah. But if you didn't have the app, it wouldn't mean anything. So you'll get a notification on your phone. And so it'll say "click okay on your phone." I've seen actually this happen. I mean, it happens to me all the time with Google.

Steve: But if you're on your desktop logging in?

Leo: Yeah, it says click okay on your phone.

Steve: Oh, okay.

Leo: And then, if that doesn't work, it says "Try another way." You know, they give you - it's not like a - but I think honestly that's maybe why he calls it two-step, because it isn't - in fact it's one step because, I mean, it's one factor.

Steve: That's back.

Leo: Well, no, I like single sign-on. But with Microsoft, for instance, when you sign on to Windows, instead of saying what's your password, it says what's your, you know, you put your Microsoft account email in there.

Steve: Right.

Leo: And then it says, okay, look on your phone for the number 80 and tap it. And I think it's a far preferable way. There's no password at all. And that's kind of what Google does with their single sign-on. And I've seen this happen with single sign-on. I suspect that's what he's talking about. But the problem is it isn't clear at all what he's talking about.

Steve: No. Nor did he in any way - nowhere did he talk about the end of passwords. He just said we're going to add steps.

Leo: But that's what single sign-on in effect does. You don't enter your Microsoft password now when you first set up Windows. But you have to have the authenticator app on your phone, and Microsoft knows that you do, and knows that you've used it. Similarly Google would have to know that you have that capability. And you're right, on a Pixel phone you don't need an app. On an iPhone you need the Google app. And then I love it because then I just tap okay on my phone.

Steve: Leo, I worked on something called SQRL for quite a while.

Leo: Yeah, it's kind of SQRL-ish.

Steve: I'm well aware of the benefits, yes, indeed. Let's take a break. I'm going to sip some water, and I'm going to talk about 21 nails in Exim's coffin.

So Tor's exit nodes. Since 2015, a Tor network researcher who goes by the moniker, I guess it's Nusenu, N-U-S-E-N-U. I googled that, thinking maybe that was a term or something that I've never heard of before, as sometimes is the case. No. There was no reference to Nusenu except this guy, N-U-S-E-N-U. Anyway, whoever he is, Nusenu has been tracking the deliberate abuse of the Tor network by quite determined and lately quite increasingly determined attackers. And of course as our listeners know, through the years the TWiT network has enjoyed the sponsorship of various high-quality VPN providers, as it does at the moment. And in talking about the various benefits and reasons to use a VPN, you, Leo, often cite the dangers inherent in Tor exit nodes. Once everybody hears what this researcher has been tracking, I doubt that anyone will or should feel comfortable using Tor without added protection.

Leo: I seem to remember the NSA, or was it the CIA, ran some Tor exit nodes. So, you know, just keep that I mind, I guess.

Steve: Yes. Okay. So because this was fascinating to me, and we've talked about Tor a lot, Tor is a cool concept. It used to be The Onion Router, T-O-R. And the idea was that you the user would choose a series of Tor nodes, typically three. The first one you connect to. Then one in the middle. And then an exit node at the end. And from each of those nodes, you would obtain their public key. You would then use the first node's public key to encrypt your traffic. And after that, that you would take that first encrypted traffic, and you would use the - oh, wait, I got it backwards. You'd first use the last node's public key to encrypt your traffic. Then you would use, to that, you would then use the middle node's public key to encrypt that, sort of like shells of an onion. And then you would use the first node's public key for the final encryption.

So what you've got now is this triple-encrypted thing. Think of it like an onion with successive layers. So now you send this to the first node. It's been encrypted with its public key, so it has the matching private key that it uses to take the encryption wrapper off. And of course the reason you do this is that nothing that went between you and that first node can be seen by anybody, your ISP and so forth, because it's been encrypted. So that node, that first node is able to take off the outer wrapper. Now it's looking at a thing with two layers of encryption. It doesn't know how to take off another layer because the layer that it's now got on the outer surface was encrypted using the middle Tor node's public key. So all it can do is send it on to the Middle Tor node.

So it does that. Middle Tor node knows its private key so it can take off the wrapper that nobody else can take off, which it does, which gives it the address of the exit node. But it can't go any further because it doesn't have the exit node's public key. So it sends it to the exit node. The exit node does have its private key that matches the public key that you originally got. So it's able to remove the final innermost wrapper of encryption. And now that thing you wanted to send through this Tor network is back in plaintext, and out it goes onto the Internet. And that's the problem is that that exit node that removed the final layer of encryption has decrypted fully after three bounces, the original plaintext that you put onto the Tor network. What is it doing with it?

Okay. So he's been tracking abuses of Tor exit nodes. Two days ago, this is why it popped up on my radar, he posted his most recent update to his earlier work which began in August of 2020 titled "Tracking One Year of Malicious Tor Exit Relay Activities Part II." And in his posting on Medium two days ago, Nusenu - maybe that's his name, I don't know. Anyway, he says: "In August of 2020 I reported about 'How Malicious Tor Relays are Exploiting Users in 2020.'" That was Part I.

He said: "Back then I made the hypothesis that the entity behind these malicious Tor relays" - and, okay, just to get everyone's attention, as many as one quarter of all Tor exit nodes are malicious. Okay? So not a couple. But your chances of hitting one are high, especially because you typically rotate among different nodes as you go. So the opportunity of your traffic exiting from a malicious node, depending upon when you're using Tor, is as high as 25%, and it rises as you use it over the course of its use. So anyway, my point is this is a big deal.

He made the hypothesis that the entity behind these malicious Tor relays is not going to stop its activities anytime soon. He said: "Unfortunately, this turned out to be true." In this follow-up post of his earlier - and by the way, in the show notes I have all three links: his very first one, this middle one, and then the one from two days ago. He says: "I will give you an update, share what additional information we learned about the attacker since August 2020, and to what extent they were and still are active on the Tor network."

So again, before I go any further, I'll share the extent of the trouble that Nusenu has uncovered. In August 2020 posting he explained: "What is this attacker actually exploiting, and how does it affect Tor users?" He said: "The full extent of their operations is unknown, but one motivation appears to be plain and simple: profit. They perform person-in-the-middle attacks" - and I guess we're no longer calling that man-in-the-middle, it's person-in-the-middle to be gender neutral - "person-in-the-middle attacks on Tor users by manipulating traffic as it flows through their exit relays." As I said, the exit relay has it back in the clear. He said: "They selectively remove HTTP-to-HTTPS redirects to gain full access to plain unencrypted HTTP traffic without causing TLS certificate warnings."

Okay. So of course we know all about this; right? You can't muck with TLS or you're going to break the authentication which is protected by the certificate, and you'll get bogus certificates. Also, it's encrypted if it's over SSL/TLS. So you really can't get anything done. But if the initial traffic is HTTP, and the far site returns a redirect to HTTPS, what these guys are doing is they're saying, oops, nope, we're not going to have the user moved over HTTPS. And we've spoken about this many times. GRC, for example, redirects anyone coming in over HTTP to HTTPS. It's not possible to access GRC without HTTPS, though it is possible to begin with HTTP and then be moved over to HTTPS to continue. And while web browsers all assumed HTTP, remember we've also talked about this, that's finally beginning to change. No idea what took them so long.

Until the assumption was being made, this moving people from HTTP to HTTPS was a necessary step since everyone entering just by typing GRC.com would default to http://GRC.com. And I should note, as our listeners will recall, GRC was among the first domains to be added to Chrome's permanent HSTS list, which Mozilla duplicates, and that explicitly gives Chrome and Firefox permission to always silently promote any and all HTTP queries to HTTPS, and it makes it quicker because it saves the HTTP to HTTPS redirect roundtrip and so forth.

Anyway, Nusenu in his posting continues. He said: "It is hard to detect for Tor Browser users that do not specifically look for the https:// in the URL bar. This is a well-known attack called 'SSL stripping' that exploits the fact that users rarely type in the full domain starting with https://." He says: "There are established countermeasures, namely HSTS

Preloading and HTTPS Everywhere. But in practice, many website operators do not implement them, and leave their users vulnerable to this kind of attack."

He says: "This kind of attack is not specific to Tor Browser. Malicious relays are just used to gain access to user traffic. To make detection harder, the malicious entity did not attack all websites equally. It appears that they're primarily after cryptocurrency-related websites, namely multiple bitcoin mixer services," which we talked about last week. He says: "They replaced bitcoin addresses in HTTP traffic to redirect transactions to their wallets instead of user-provided bitcoin addresses. Bitcoin address rewriting attacks are not new, but the scale of their operations is. It is not possible to determine whether they engage in other types of attacks."

He said: "I've reached out to some of the known affected bitcoin sites, so they can mitigate this on a technical level using HSTS preloading. Someone else submitted HTTPS Everywhere rules for the known affected domains." And he notes that HTTPS Everywhere is installed by default in Tor Browser. "Unfortunately," he says, "none of these sites had HSTS preloading enabled at the time. At least one affected bitcoin website deployed HSTS preloading after learning about these events." Okay. So I have to say I am astonished that any sort of bitcoin transaction site might be lacking in such basic security awareness and provision. But since bitcoin is unregulated, it's user beware. And if this is the state of cryptocurrency security, I guess I'm less surprised that we keep hearing about this or that cryptocurrency exchange being hacked.

Elsewhere, Nusenu notes that SSL stripping and person-in-the-middle attacks are only one of many potential problems with Tor's inadvertent hosting of malicious exit nodes. As an example, he considers the instances where a new remote vulnerability is discovered in Firefox and thus in the Tor version of Firefox. Running a large network of exit nodes would allow attackers to immediately reach back down their end-node connection to exploit such newly discovered vulnerabilities before the Tor users' browser had a chance to update.

So just how big is the problem? Is it a couple of nodes that users are likely to exit from? Well, as I said, no. The graph above in the show notes shows just how big the problem is. The graph's scale on the left is difficult to read. But the uppermost number is 26%. Nusenu's caption for that graph reads: "Figure 1: Malicious Tor exit fraction measured in % of the entire available Tor exit node capacity over time by this particular malicious entity between July of 2020 and last month, April of 2021." He said: "Peak value: The attacker did manage approximately 27.5% of the Tor network's exit capacity on January 2nd of 2021."

Okay. And it's interesting, the graph sort of shows a rising percentage, then a sudden drop. And then it'll rise again and drop. And then it'll rise again and drop. And then it'll rise again. And in the case of the largest and longest one, it rose, and it kind of slowed down, and then dropped. Well, okay. What's happening is that the bad guys are being found. I mean, there is active combating of malicious exit nodes by Tor network administrators. But this is all sort of volunteer exit node; right? I mean, we talked about how, you know, anyone who wants to can contribute to the Tor network by setting up their own exit node, where they allow users' traffic to come encrypted into their system, get decrypted by this exit node that they run on their network, and then out it goes onto the Internet. I don't want to run one, but good Samaritans do.

It turns out that bad Samaritans do, as well. And that because they are set up quickly, and due to the nature of the way they're set up, it is possible to track their aggregation over time, which is what Nusenu has figured out how to do. And so what we see is a large population of malicious nodes built up. While they are active, as many as, actually more than, one out of four connections over Tor is exiting through a node controlled by malicious parties who are hoping you're going to do something without TLS encryption. And god help you if you do because these people are not working in your favor.

He did also note, though, that they're not mucking with all traffic. They are being selective about what traffic they mess with. And of course that does make their detection more difficult. So I guess that's good. So he said that there's better than, as a consequence of 27.5%, better than a one in four probability, which as I noted rises over time since exit nodes are being randomly chosen and rotated. So the chance that a user not using some form of encryption will have traffic exiting through a malicious node - now, of course, it also is dependent upon where in this weird sawtooth cycle of malicious activity, node activity growing and then being suddenly cut off, like where in that cycle you happen to be using the Tor network, well, that matters, too. But it demonstrates that you just can't take it at face value that the use of Tor is going to be secure.

So the bottom line here is there's no free lunch. Tor provides, as we know, some valuable services. But it's not a panacea. Any user of Tor must assume - and by the way, it's gotten way worse in the last couple years. This was not true when we first talked about the Tor network in the beginning, and even over the course of the last few years, while this guy Nusenu has been tracking this, although his tools are getting better, so maybe he's better at finding the problem, he's concluding that it is really getting worse, and way worse in 2021 than it had been before. So any user of Tor should assume, must assume that the exit nodes they're emerging from may be under the control of malicious entities who will take any and every opportunity to interfere with and subvert the user's traffic if they can.

He wrote: "We know about mitmproxy, sslstrip, bitcoin address rewrites, and" - get this - "download modification attacks. But," he said, "it's not possible to rule out other types of attacks. Imagine an attacker runs 27% of the Tor network's exit capacity and a Firefox exploit affecting Tor Browser gets published before all users got their auto-updates." And, wow. A download modification attack? Talk about chilling.

You use Tor to go get something that you want to keep very private. That's the reason you're using Tor. But the website that offers whatever it is doesn't support HTTPS. And apparently there are a lot that still don't. Okay, you know, they just say, hey, we're not going to do that. Still, you want it badly. So you download it over Tor. Even if the site in question was 100% legitimate, who knows what you actually downloaded? HTTP offers zero authentication of the other end's identity.

It was noted that a Tor HTTPS-only browser would be one solution. And about that, Nusenu wrote, he said: "The HTTPS-only mode, which might land in Tor Browser based on Firefox 91 ESR, would be a strong protection. But there are still some uncertainties with that as well," he says, "as a Tor Browser developer points out on a Tor mailing list. When Tor Browser migrates to Firefox 91 ESR," he wrote, "we will look at enabling HTTPS-only mode for everyone. But there remains a significant concern that there are many sites that do not support HTTPS," he said, "especially more region-specific sites, and the question of what messaging Tor Browser should use in that case."

In other words, unfortunately, it's still not practical to force HTTPS. Yet arguably it's not safe not to have HTTPS if you're using Tor, without some other kind of protection. So I think our takeaway here should be that Tor needs to be used with a full awareness of its inherent dangers. While it can significantly obscure its users' real-world location and identity, many entities, both malicious and, Leo, as you noted, law-enforcing also, may be closely monitoring everything they can about a user's activities or about Tor's users' activities. And even in some cases, if they're malicious, actively modifying and subverting any traffic that's available to them in the clear. So whenever using Tor, keep in mind the danger of HTTP and the real need for some other privacy and security protecting tunnel such as a trustworthy VPN. At this point, knowing what I know, I wouldn't consider using Tor without the added protection of a VPN. I just, you know, I don't think you can.

Leo: Hey, did you skip the Exim story?

Steve: Oh, my goodness. How did I? Thank you.

Leo: I mean, you might have on purpose because, you know. But I just thought I'd mention it.

Steve: No. Thank you, thank you, thank you.

Leo: Well, it wasn't me. The chatroom and Jason and everybody went, "Hey, what about our Exim story?" You did tease it.

Steve: I sure did. And here it is. So, okay, 21 nails are not going to kill Exim. Nothing will kill Exim. But it does mean that, if you or your organization is using the extremely popular, and we'll talk about just how popular in a second, Exim email transfer agent, which is the default email transfer agent provided by many Linux distros including Debian to send and receive email, you will definitely want to be sure - I mean, like this is one of those, okay, like pause the podcast and go update - you've got to be sure that you're running the most recently patched version.

Two months ago in March, E-Soft performed an Internet-wide study, probably due to the Microsoft Exchange Server debacle, studying the Internet's email servers. They approximated that 60, six zero, percent of the publicly reachable mail servers on the Internet were running Exim - 60%. So that obviously makes it, without any further computation, the most popular email server on the Internet, period. Unfortunately, Exim, E-X-I-M, is short for "EXperimental Internet Mailer." And after 17 years of its presence on Git, it might be nice if, today, it was a bit less experimental.

In response to Qualys's most recent security research, which we'll get to in a minute, all of the most widely used Linuxes CentOS, Red Hat Enterprise, SUSE have rolled out fixes. Debian's "oldstable," codename Stretch; its "stable," codename Buster; and its "Still-in-development," thus Sid versions, they're all updated. But the "unstable," which is codenamed Bullseye, remains vulnerable. The problem is that there are hundreds of also-ran distributions, and it's of course up to each individual distribution to update their own packages and to then work to get those updated and replaced online, old instances updated and online.

So, okay. Since most of - and of course 21 nails is 21 vulnerabilities. Most of the 21 serious vulnerabilities Qualys uncovered date back to Exim's emergence 17 years ago, in 2004. That is to say, all versions of Exim on the Internet are vulnerable. So we're back in the all-too-familiar position of having publicly known and remotely exploitable vulnerabilities in email software that may not be receiving regular maintenance. And a great many Internet-connected appliances may be based upon a build of Linux with a publicly exposed email agent running Exim.

So what did Qualys find? The security researchers at Qualys dubbed their report "21 Nails" because from a source code audit - they just read the source. From a source code audit they found 10 vulnerabilities that can be remotely exploited. And most of the entire

21 can be exploited either in Exim's default configuration or in what they said was a very common configuration. And, as I mentioned before, most of them affect all versions of Exim, all the way back 17 years to 2004.

There are 11 local vulnerabilities. And I'll just give you a sense for that. Link attack in Exim's log directory. Assorted attacks in Exim's spool directory. Arbitrary file creation and clobbering. Arbitrary file deletion. Heap buffer overflow in queue_run. Blah blah blah. Those are local. So those are not remote. We're mostly worried about the remote ones because that's where the attacks are going to come from, largely.

So we have, in all versions of Exim, 60% of the servers on the Internet, right: Integer overflow in receive_add recipient. Integer overflow in receive_msg. Out-of-bounds read in smtp_setup_msg. New line injection into spool header file. Heap out-of-bounds read and write in extract_option. Line truncation and injection in spool_read_header. Failure to reset function pointer after BDAT error. Heap buffer underflow in smtp_ungetc. User-after-free in tls-openssl.c. And Heap out-of-bounds read in pdkim_finish_bodyhash.

Okay. So those all sounds tricky and techie. Qualys has published a detailed write-up, I've got the link in the show notes, showing step-by-step code mistakes in the source and exploitation mechanisms. But they stopped short of working exploits. However, since Exim is open source and published under the GNU GPL, there's no point in attempting to obfuscate any of this. So we can expect to be seeing still more trouble downstream as remote attackers use any older and not-just-updated Exim instances as their means of gaining entry to internal enterprise and government networks. We already know what's going to happen. I mean, this story has already been written. I'm not going to go into the blow-by-blow detail here. It's all available, as I said, on Qualys's excellent vulnerability disclosure. But here's how they introduced their research.

They said: "We recently audited central parts of the Exim mail server and discovered 21 vulnerabilities, 11 local and 10 remote. Unless otherwise noted, all versions of Exim are affected since at least the beginning of its Git history, in 2004. We have not tried to exploit all of these vulnerabilities, but we successfully exploited four Local Privilege Escalations and three Remote Code Executions." They have four bullet points: "We will not publish our exploits for now. Instead, we encourage other security researchers to write and publish their own exploits." Oh, yeah. What could possibly go wrong with that? They said: "This advisory contains sufficient information" - and indeed it does - "to develop reliable exploits for these vulnerabilities. In fact, we believe that better exploitation methods exist." Sure. Why not try some?

They said: "We hope that more security researchers will look into Exim's code and report their findings. Indeed, we discovered several of these vulnerabilities while working on our own exploits." Oh, Jesus, they're cascading. And, finally, they said: "We will answer to the best of our abilities any questions regarding these vulnerabilities and exploits on the public 'oss-security' list." And then there's a link in the notes. And they said: "Last-minute note. As explained in the timeline, we developed a minimal set of patches for these vulnerabilities. For reference and comparison, it is attached to this advisory and is also available at" - and then we have the link.

So in their disclosure, as opposed to the vulnerability disclosure in their announcement, basically, they wrote: "Once exploited, they could modify sensitive email settings on the email servers, allow adversaries to create new accounts on the target mail servers." And it's worth noting that Exim already has a history of trouble. Back in June of 2019, Microsoft warned of an active Linux worm targeting an earlier Exim remote code execution bug. And a month later, attackers started exploiting vulnerable Exim servers to install the Watchbog Linux trojan, which as a consequence added them into a Monero cryptomining botnet. We know that's not going to happen now. Now what's going to happen is ransomware.

And the U.S. NSA, the National Security Agency, said last May of 2020, a year ago, that the Sandworm Russian military hackers have been exploiting that same critical Exim remote code execution since at least August of 2019. In other words, we already have evidence of an older remote code execution vulnerability, known, published, and patched years before, still being leveraged by bad guys a year later. Now Qualys has just dropped another goodie bag of these vulnerabilities in the email servers running 60% of the Internet's domains into the public discourse. Of course, the Microsoft Exchange Server catastrophe showed us just how vulnerable an exploitable email server can be. Now the whole world knows that Exim, the most widely deployed email server, can now be remotely exploited.

As Qualys themselves wrote: "This advisory contains sufficient information to develop reliable exploits for these vulnerabilities. In fact, we believe that better exploitation methods exist." Oh, joy. And if we thought that updating and cleaning up the big mess created by Exchange Server was a problem, just try doing that with the Internet's Exim servers, especially all those that are embedded into firmware-based appliances and long-forgotten dusty closets. Yes, we will be talking about this, I'm afraid, in coming months.

Leo: Those dusty closets are full of bad stuff, I'll tell you.

Steve: Oh, Leo. It's not just dust bunnies. It's bad guys. And they're going to use this to get into corporate networks and to launch more ransomware. Because now botnets are considered quaint, as is Monero mining. Why do that when you can extort millions of dollars from a juicy target?

Leo: Well, we're going to talk about that in a little bit, too, yeah.

Steve: We are.

Leo: Steve, let's go with some extra stuff here. Come on.

Steve: Indeed. Yeah. We have a novel.

Leo: I'm so excited.

Steve: Yeah. When I checked it out over on Amazon, I was told that I could have it for free as part of my Audible free trial. So when we next talk...

Leo: Wait a minute. You're going to listen to it?

Steve: No, no, no. I'm not. But I know many - I just wanted to mention that it...

Leo: I thought you were going to go Audible. I was going to just fall off my ball.

Steve: I was going to mention you could have somebody read it to you, if you would like.

Leo: Actually, Andy Weir uses a really good reader. Lisa and I listened to "The Martian" together driving on the road to Hana in Hawaii.

Steve: Nice.

Leo: We'll never forget it. It was like a life experience that we shared that we'll always remember really, really well.

Steve: Okay. So what we have for our listeners who don't yet know...

Leo: I haven't said yet, yes.

Steve: Andy Weir, who is famous for having written "The Martian," has a new novel which the reviewers are just falling all over themselves for. It's called "Project Hail Mary." It's a solid five stars. I looked at the demographic breakdown of stars, and it's like 84% are fives, and the balance are fours, with only a couple threes. For example, Nick, who reviewed this on the fourth, the novel just came out a week ago, he's a verified purchaser. He said: "I don't even remember pre-ordering this book. It just showed up in my Kindle app this morning." He said: "So I decided to read the first chapter before starting work. Four hours later, I can finally put the book down since I'm done."

Leo: Wow.

Steve: Now, I don't like to read that way because...

Leo: It's gulping, not chewing and tasting.

Steve: You look around, and it's like, wait, what happened?

Leo: What happened? Where am I? It is a 16-hour book. So that's a good amount of reading in four hours.

Steve: Yeah, he went fast. So he says: "'The Martian' was a great story. 'Artemis'" - that's another one that Andy Weir wrote. "'Artemis,' he says, "was a great story. This one is better than either of those." He says: "If you like science fiction with actual science, this is for you. If you like stories with interesting, well-developed characters, this also has that. If you want excitement and a thrilling plot, here you go. If you want romance and sex, well, there you're completely out of luck. But if that was the kind of book you wanted, I doubt you'd be reading this review anyway. Speaking of, why ARE you still reading this review? Go read the book. It's way better than this."

Somebody else said: "Andy does it again." He said: "A spiritual sequel to 'The Martian' that had me grinning throughout the entire book. Made my inner nerd squeal with delight on many occasions. Has everything I ever wanted in a sci-fi book, just didn't realize it until now. Read it. That is all."

And I'll share one more, another five out of five. I mean, they virtually all were. This one's subject was "Stop reading this review. Read 'Project Hail Mary.'" He said: "A previous reviewer said: '"The Martian" was a great story. "Artemis" was a great story. This one is better than either of those.' Wrong. This one is MUCH better than either of those." He said: "Instant classic." He said: "If you mixed Asimov's 'The Gods Themselves' and Heinlein's 'Citizen of the Galaxy,' and added in a few gallons of Clarke and Niven, it would be like this. I'd write more, but I'm off to re-read the novel."

Leo: Oh, my goodness. I want to get this now.

Steve: It sounds really good.

Leo: Actually, you know, I want to get Andy - I interviewed Andy Weir of course after "Artemis." It's interesting because "Artemis" was the beginning of a new series for him, and this book does not continue that. Maybe he's planning to down the road.

Steve: So this is not a spoiler, and I have not read the book. But this is something about a team of three go off on some distant mission to save the Earth. And only one guy is left to solve, like, to figure this out. So again, as I said, that's not a spoiler. I've not read the book. I don't know anything about it. But wow.

Leo: I'm going to try to get Andy in and do a special interview because...

Steve: Given that he apparently has really outdone himself.

Leo: Yeah. We interviewed him after "The Martian," and I think I interviewed him again after "Artemis." So we should really get him for this. All right. We don't have the show anymore, but maybe we'll put it in Club TWiT or something like that. Very cool. Very cool. I can't wait to read it.

Steve: Sounds like a win for our listeners. Paul Babiak, he posted in the grc.securitynow newsgroup under the subject "One possible solution to QNAP vulnerability." Actually, he found what I was maybe suggesting as a solution, a walkthrough of an installation of OpenMediaVault for the QNAP hardware. I've got a YouTube link and a link to OpenMediaVault.org. You can install non-QNAP firmware onto your QNAP NAS in order to get something that, I mean, it could - I was going to say, I was going to hedge my bets here and say, wait a minute, can I really assert that it's more secure? Yes, because it could be not be less secure than what you're getting from QNAP. So yes, thank you, Paul.

And also Jon S. sent by DM, he said: "First hack that hits close to home. Sitting in the ER of Scripps Health with my wife." This was on Sunday. He said: "They were hacked a few weeks ago and are still doing all charts and orders via paper records. The process is taking about 4 to 6 hours longer than normal for doctors to get lab work back. Nurses are making notes on square sticky note pads. I'm an IT sysadmin and security guy." And obviously a listener to Security Now!. And he DM'd me. He says: "This upsets me to no end. Thought I'd share a few pictures for observations." And he did include some photos of some screens of computers that are down at Scripps. So we talked about the Scripps Health attack last week, and here he is. I also told him that I hoped everything was okay with for whatever reason he was in Scripps ER with his wife. But it really is having real-world consequences. These things do.

Okay. And I'll just mention that I have nothing huge to report on the SpinRite front. I am unglamorously working my way through the code, line by line, changing the sizes of the registers and the variables used to manage drives, to accommodate today's larger than 2.2TB drives, containing any partitioning and any file system. And also, since we'll be living with and using this codebase after it's converted from 16-bit real mode segmented code to 32-bit protected mode flat model, and also booting under UEFI and BIOS, and also to host native operation for USB and NVME mass storage, I am taking some time to clean things up a bit while I'm there, as I'm moving through it, to get it a little bit more ready for its future, which seems bright.

Now that I have access to upper memory, which I have never had before, I'm able to move some of the things that SpinRite had been cramming into lower memory up into upper memory to ease the pressure on the use of lower memory, which eliminates some jumping through hoops. So anyway, I'm at work on it, and I am posting updates to the newsgroup. And when I have them, new code to test, as I mentioned before. And that's all been going well.

Okay. News from the DarkSide. Because this latest high-profile ransomware attack has been extensively covered by the popular press, I assume that our listeners already know that the largest fuel pipeline in the United States, run by a company called Colonial Pipeline, and actually the pipeline is also called Colonial Pipeline, it was shut down late Friday when they were forced to terminate all of their network operations in an effort to contain a ransomware attack. And I assumed that there wasn't much more to know. But in doing my due diligence for the podcast, I discovered that was not the case.

So Colonial Pipeline is keeping rather quiet about specifics, likely following advice coming at them from many sides. But the FBI has confirmed that this was a ransomware attack conducted by DarkSide, a new Ransomware as a Service group, and remember we talked about Ransomware as a Service, how that's like the new way to do this. And what we're developing is essentially a ransomware economy and sort of an ecosystem where we're getting specialization among the players that then form a chain. So there are what do with the money specialists, bitcoin mixing and so forth. There are the software development specialists, and there are the hack-into-the-system specialists. And they're actually, I think they call them "access agents" or something. I saw that the other day, it's like, oh, goodness.

Anyway, these guys are new. They first appeared on the scene in August last summer, 2020. So just to set the stage for anyone who may have been out hiking through the wilderness over the weekend and offline ever since, incredibly, Colonial Pipeline is responsible for transporting refined petroleum products - gasoline - between refineries located down in the Gulf Coast to markets throughout the southern and eastern U.S. When its pipeline is up and running, as it always is, it transports 2.5 million barrels per day through the 5,500 miles of pipeline to provide an astonishing 45% of all fuel consumed by the East Coast.

So when the East Coast's petrochemical fuel supply suddenly and unexpectedly drops by nearly half, markets are upset, and states of emergency are declared, as has happened, by the Biden administration for Washington, D.C. and the seven states that the pipeline runs through. This was temporarily done to lift restrictions on fuel transport by road in an endeavor to keep at least some fuel moving. But good luck with that. At 42 gallons per barrel, tanker trucks are not going to match a continuous flow of the 105 million gallons of refined fuel which normally flows through that pipeline every day.

The Governor of Virginia today, just today, Tuesday, declared their own state of emergency. Their declaration begins: "On this date, May 11th, 2021, I declare that a state of emergency exists in the Commonwealth of Virginia to prepare and coordinate our response to the voluntary shutdown of the Colonial Pipeline due to a cyberattack on its business systems' informational technology infrastructure on May 7th. If prolonged, the pipeline closure will result in gasoline supply disruptions to various retailers throughout the Commonwealth, since the pipeline is the primary source of gasoline to many Virginia retailers." And yesterday North Carolina declared a similar emergency, and gas station pump rationing has been instituted there.

Okay. So now the famous SolarWinds attack, as we all know, made the news in March, loudly, because it was labeled "the most significant cyberattack ever." So, okay, whoo, big headlines. And people could be upset by the idea of that, especially since the attacks were credited to Russia-linked cybercriminals. But the idea of that was the attack's only real effect on most people. This time, of course, this is an effective attack against critical American infrastructure, forcing declarations of emergency. When you cause the shutdown of nearly half the supply of gasoline to a large and influential portion of the U.S., the problem is no longer theoretical or superficial.

Okay. So what about DarkSide? I found a copy of their extortion demand note. Actually I found many of them over time because this has been on the cybersecurity industry's radar since, as I mentioned, last summer. And I have a - I'm getting close to the screen so that I can read this. Maybe I can zoom in. Although zooming in it's so fuzzy.

Leo: Fine print, yeah.

Steve: Yeah. It doesn't really help very much. So they said: "Welcome to DarkSide. Your computers and servers are encrypted. Backups are deleted. We use strong encryption algorithms so you cannot decrypt your data. But you can restore everything by purchasing a special program from us, Universal Decryptor."

Leo: How thoughtful.

Steve: Yeah, isn't that nice they make that available, Leo, for the low, low price of several million dollars. Anyway, they said: "This program will restore all your network. Follow our instructions below, and you will recover all your data." And then they said: "What guarantees? We value our reputation. If we do not do our work and liabilities, nobody will pay us. This is not in our interests. All our decryption software is perfectly tested and will decrypt your data. We will also provide support in case of problems. We guarantee to decrypt one file for free. Go to the site and contact us."

Then they say: "How to get access on website. Using a Tor Browser, download and install Tor Browser from this site." And they point you to TorProject.org. "Open our website." And then they give us an onion domain, http, okay, no "s," http://darksid, and it's just "sid," and then fqzquhtk2.onion/ and then a big crypto-looking thing, looks like Base64 all caps. Then they said: "When you open our website, put the following data in the input form," and then they give a key. Then they said: "!!!DANGER!!!," three exclamation points on either side. "Do not modify or try to recover any files yourself. We will not be able to restore them." And then "!!!DANGER!!!"

So, okay. That's these guys. In addition to the ransom note, victims of a DarkSide attack receive an information pack informing them that their computers and servers are encrypted. The info pack lists all of the types of data that were stolen, and provides the URL of a "personal leak page" where the data is already loaded, waiting to be automatically published, should the company or organization being extorted from choose not to pay up before the deadline expires. DarkSide also tells victims it will provide proof of the data it has obtained, and is prepared to delete all of it from their own storage once payment has been received. I did also see, although this may have been earlier, I didn't see it in this particular attack, the doubling of the ransom demand in equivalent dollars if negotiation isn't concluded by a certain date.

Now, what's weird about these people is they appear to imagine that they're running a business more than a crime ring.

Leo: They act like that, don't they. It's like, we're a serious enterprise.

Steve: Yeah. They really do. Yeah. Well, and when they released a new version of their software two months ago which could encrypt data faster than before, they issued a press release...

Leo: Now 20% faster. Geez. Oh my god.

Steve: Yes, we'll mangle and tangle your network in half the time as previously. They invited journalists to interview them.

Leo: Oh, yeah, sure.

Steve: And their website on the dark web lists all the companies they have attacked and hacked and what was stolen from them. And, get this Leo, they have an ethics page.

Leo: Oh.

Steve: Listing which types of organizations they will not attack. They've stated that they will not attack hospitals, hospices, schools, universities, non-profit organizations, or government agencies. And I suppose after this, what they've just stepped in, they'll be adding "critical infrastructure" to that list.

Leo: I have to figure Seal Team 6 is about to jump in their window. This is not going to go well.

Steve: Exactly. So anyway, that's something different about these guys. They said they intend to cause no harm, they just want money. On the website they wrote: "Our goal is to make money and not create problems for society. We do not participate in geopolitics, do not need to tie us with a defined government and look for our motives." And in this case they realize they have probably painted a huge bull's-eye on themselves. They indicated that they had not been aware that Colonial Pipeline was being targeted by one of their affiliates. They wrote: "From today, we introduce moderation" - a little late, but okay. "We introduce moderation and check each company that our partners want to encrypt to avoid social consequences in the future."

So we know that they used the Salsa20 symmetric cipher with a custom matrix - and actually switching to that would have been responsible for the speed increase because it's very quick - and RSA-1024 for their public key operations. So from a tech standpoint, their crypto appears to be well designed. And that's been the consensus of the security industry since they appeared back last August. Their ransoms have generally ranged from 200,000 to two million, so not nutty 50 million requests.

And traditionally much of this podcast is focused upon developing an understanding of just exactly how porous most of our network and, well, our computer and network security is today. We look at the details, and we attempt to determine why these problems happen and what might be done to prevent such trouble in the future. And unfortunately we've reached the conclusion that we're not ready for the world. Most, if not all, of our existing IT infrastructure is not ready to stand up to determined attack. Look what just happened with Exim. This is going to be a catastrophe.

And it's a sad fact that we have to somehow deal with. Much more focus, time, and attention is going to have to be put into the security side of our technology. It's going to burn a bunch of time, effort, and money just to prepare. But there's just no way around it. It's expensive. It's a waste of resources or consumption of those. But it has to happen.

Leo: Well, that's what we talked about last week with this governmental task force.

Steve: Right.

Leo: Right? The timing was interesting because of course...

Steve: Oh, isn't that weird?

Leo: ...then there's a massive infrastructure attack shortly after that.

Steve: Yup. Yup.

Leo: We clearly have to do something. This has gotten out of hand. And, you know, it's only a matter of time before something really serious gets hacked.

Steve: Well, yes. And I had that same thought. We've talked about how bad as COVID-19 has been, there are previous viruses like the Spanish flu of 1812 or whenever it was, where actually, if that one had happened today, the consequences would have been far worse. My point is we get wakeup calls. Remember the old expression, "Fool me once, shame on me?" Wait, no, wait. "Fool me once, shame on you. Fool me twice, shame on me."

Leo: Fool me three times, George Bush. No, no, that's something else.

Steve: So, you know, we like having the lights on. Lights are handy. And having power for refrigeration and all the things that we use electricity for now. Whenever we have a brief outage of our electric supply, often scheduled by our local supplier, you walk around flipping switches on rooms when you walk into them and think, oh, shoot, I forgot, we don't have any power. We really, really, really are vulnerable. And so again, this is inconvenient. And I won't say in any way am I glad for this. I am certainly not. Except, as I said, the SolarWinds attack was arguably, ooh, bad headlines. Bad Russians. But now we don't have gas on the Eastern seaboard.

Leo: Not good.

Steve: You couldn't get a better wakeup call. You couldn't get a better, you know, a declaration of emergency by the administration to allow for more tanker trucks to run north and south. Good luck. That's 105 million gallons a day, that monster pipeline. And Leo, we've also talked about a monoculture. How about a mono pipeline? That's just, you know, this whole, quietly in the background, everything is getting consolidated. So we end up with many fewer, much less redundancy, and it becomes much more critical. The lack of redundancy becomes protected.

Leo: Apparently it blew up in 2016 and was shut down for two weeks. So it's not the first time this pipeline has failed. It just does seem like it's a very, very vulnerable setup.

Steve: It's fragile, yeah.

Leo: Yeah. Wow. Boy, it just - it feels like we're hanging by a thread at all times. I'll be honest with you. Modern civilization is so interdependent and so unredundant. That's why the Internet is such a miracle. It's designed to survive catastrophe. But apparently nothing else is.

Steve: We hung onto it.

Leo: Yeah, yeah. Oh, man. That's scary.

Steve: And the Internet's dodged a few bullets. We talked about Dan and discovering the danger that DNS was in.

Leo: That's right, the DNS, yeah.

Steve: Okay, so here's Exim. I guarantee you, you know, how much did we have to talk about the Exchange Server problem?

Leo: Yeah, yeah. It's not over. Just beginning. Well, Steve, we've come to the end of this grim edition of Security Now!. As always, we thank you for elucidating these difficult topics and giving us at least some hope that something can be done about it.

Steve Gibson's at GRC.com, that's his website. That's where you'll find of course SpinRite, the world's finest mass storage maintenance and recovery utility.

Steve: It's starting to roll off your tongue, Leo.

Leo: Comes just right off, just like that. No longer just hard drives. Anything you store your data on. You'll find it there. 6.0 is the current version. Work proceeds apace, as Steve mentioned, on 6.1. You can participate in that development and of course get a free copy of 6.1 if you buy 6.0 now. You really need it. While you're there, you can get a copy of this show, too. Steve has the only 16Kb audio version of this show, for good reason. But if you're bandwidth-impaired, you'll be glad. He also has transcripts written by an actual human being. Elaine Farris does such a nice job with those. That's at GRC.com.

And of course, as always, it's free. 64Kb audio versions, as well. We have audio and video at our website, TWiT.tv/sn. So you can download it there. If you want to watch us do it live, it's every Tuesday at about 1:30 p.m. Pacific, 4:30 Eastern, 20:30 UTC. The livestreams are at TWiT.tv/live, audio and video. And if you're watching live, you should chat with us live at irc.twit.tv. Steve, I hope you have a wonderful week, and we'll see you next week.

Steve: Will do, my friend. Ciao.

Copyright (c) 2014 by Steve Gibson and Leo Laporte. SOME RIGHTS RESERVED

This work is licensed for the good of the Internet Community under the Creative Commons License v2.5. See the following Web page for details: http://creativecommons.org/licenses/by-nc-sa/2.5/