

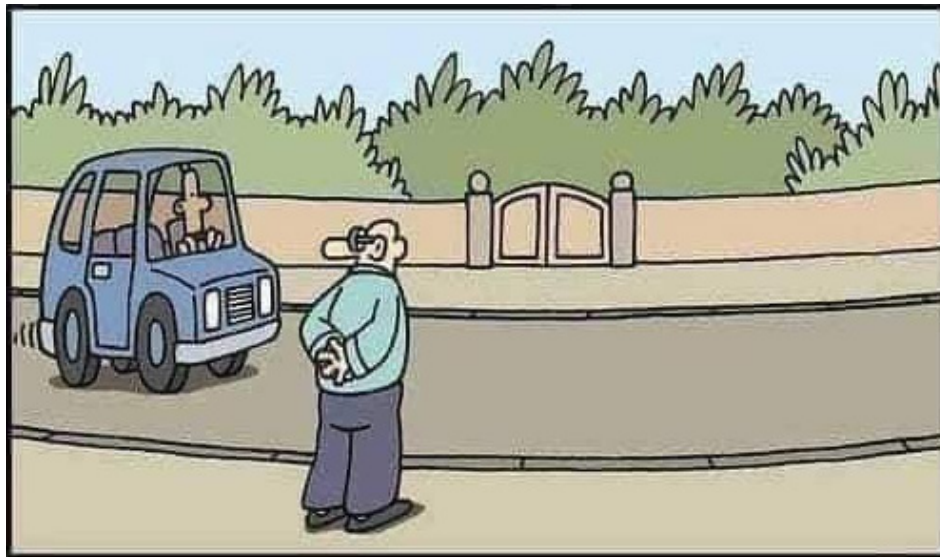
Security Now! #818 - 05-11-21

News from the Darkside

This week on Security Now!

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 this latest very high profile ransomware attack from a previously low-key attacker.

When you get it... Shhhh!



Security News

TsuNAME - "DNS Configuration Flaw Lets Attackers Take Down DNS Servers"

This is one of those click-bait stories where its name "TsuNAME" is the best part. But it's still interesting and 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 Kevin 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 naïve recursive DNS resolvers are asked to resolve one such misconfigured domain, that recursive server will get itself into a name resolution loop, causing 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 has 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, published way back in October of 1993, was titled: "*Common DNS Implementation Errors and Suggested Fixes.*" That's right. Things that can go wrong and how to fix'em. Section 2 of RFC 1536 bears the title "Recursion Bugs." 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 for the queried name. Some servers reprocess this name to try and get the desired record type.

However, in certain cases, this chain of events 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, more difficult to detect, a set of servers might form a loop wherein A refers to B and B refers to A. This loop might involve more than two servers.

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:

ABSTRACT

The Internet's Domain Name System (DNS) is one of the core services on the Internet. Every web page visit requires a series of DNS queries, and large DNS failures may have cascading consequences, leading to unreachability of major websites and services. 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 (we observe up to 5.6k queries/s). Using production data from .nz, the country-code top-level domain (ccTLD) 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 dependent 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.

We used CycleHunter to evaluate roughly 184 million domain names in 7 large, top-level domains (TLDs) and discovered 44 cyclic dependent NS records (likely from configuration errors) used by 1400 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, we helped them to mitigate their vulnerability to TsuNAME.

Later in their paper, they discuss their use of their CycleHunter tool and show that they found a total of 3,696 DNS resolvers which were not protecting their queries from cyclic DNS misconfigurations.

They manually tested the DNS resolvers Unbound, BIND, KnotDNS, 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. And, interestingly, DNS developers do need to always be, and generally are, on the lookout for DNS looping errors. They noted that the change log for the Unbound DNS resolver contains 28 entries related to looping.

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 that, sure enough, there were DNS definitions that would cause those servers to become stuck. So this research identified those servers and got them patched up. And the researchers also developed their CycleHunter tool to allow administrators of DNS to check their own DNS Zone definitions for any cyclic lookup trouble.

<https://tsuname.io/>
https://tsuname.io/tech_report.pdf

Huh Google?

Last Thursday, Google's Mark Risher, 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." Unfortunately, that's not what his blog posting addressed. And no one seems sure about what exactly his blog posting DID address since, it led to many confusing and misleading tech press headlines. I saw: "*Google wants to enable multi-factor authentication by default*" and "*Google is turning on two-factor authentication by default*" and "*Google Will Start Automatically Enrolling Users in Two-Step-Verification (2SV) Soon.*"

I saw many confused users who read this to mean that Google would be requiring the use of two-factor authentication. And I can see how one might get that 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, typically abbreviated 2FA. But now, we have Google's 2SV for two-step verification. 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 up to three steps for the verification of our identities? And if you need to go get your phone, arrange to unlock it with your identity, then respond to a prompt or text message or one-time password, aren't we up to six or step steps by that point? I've lost count.

Anyway, I've read through Mark Risher's blog posting and here's the problematic paragraph that no one is quite sure how to interpret:

"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. Soon we'll start automatically enrolling users in 2SV if their accounts are appropriately configured."

Uhhhhh what????!! I have no idea what he means when he says: "we'll start automatically enrolling users in 2SV if their accounts are appropriately configured" What does "appropriately configured" mean, exactly? And that's the problem. It apparently means something to Mark, but it's gobbledygook 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. As we all know, you either have 2nd factor authentication enabled for authentication to your Google account (as I do), or you don't. There's no third setting labelled "I'm open to the idea, hit me up when you want."

The only thing I can figure is that Mark Risher 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! And I'm the Director of Product Management for Identity and User Security. I'd better think of something to say today." So he banged out something that confused the whole world.

I think what we need to take away from Mark's aberrant posting, is that Google is a fan of using more than just our eMail address and password for authentication—as we know they are. 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 choose to put it, "another step", to their logons.

But as for what Mark actually wrote last Thursday for World Password Day, I have no idea what Google could possibly mean by "automatic enrollment in 2SV" nor does anyone else. Maybe they don't know. But 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 2SV... notwithstanding the fact that it's actually 3 or more SV's.

21 Nails in Exim's coffin

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 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 that you're running the most recently patched version.

Two months ago in March, E-Soft performed an Internet wide study which approximated that 60% of the publicly reachable mail-servers on the Internet were running Exim. 60%!! That obviously makes it the single most popular mail server on the Internet. Unfortunately, "Exim" is short for "EXperimental Internet Mailer" ... and after 17 years of its presence on Git, it might be nice if, by today, it was a bit less "experimental."

In response to Qualys 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 "Still-in-development" (Sid) versions are updated. But the "unstable" (codename Bullseye) remains vulnerable. The problem is that there are hundreds of "also ran" distributions and it's up to each individual distribution to update their own packages and to then work to get those updated to replace the old instances that may be online. And since most of the 21 serious vulnerabilities Qualys uncovered date back to Exim's emergence 17 years ago, in 2004, 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. A great many Internet-connected appliances may be based upon a build of Linux with a publicly exposed eMail agent running Exim.

What did Qualys find? The security researchers at Qualys dubbed their report "21 Nails" because from a source code audit they found 10 vulnerabilities that can be exploited remotely. And most of the entire 21 can be exploited in either Exim's default configuration or in what Qualys said was "a very common configuration." And, as I mentioned before, most of them affect all versions of Exim going back at least 17 years.

Local vulnerabilities

- CVE-2020-28007: Link attack in Exim's log directory
- CVE-2020-28008: Assorted attacks in Exim's spool directory
- CVE-2020-28014: Arbitrary file creation and clobbering
- CVE-2021-27216: Arbitrary file deletion
- CVE-2020-28011: Heap buffer overflow in queue_run()
- CVE-2020-28010: Heap out-of-bounds write in main()
- CVE-2020-28013: Heap buffer overflow in parse_fix_phrase()
- CVE-2020-28016: Heap out-of-bounds write in parse_fix_phrase()
- CVE-2020-28015: New-line injection into spool header file (local)
- CVE-2020-28012: Missing close-on-exec flag for privileged pipe
- CVE-2020-28009: Integer overflow in get_stdinput()

Remote vulnerabilities

- CVE-2020-28017: Integer overflow in receive_add_recipient()
- CVE-2020-28020: Integer overflow in receive_msg()
- CVE-2020-28023: Out-of-bounds read in smtp_setup_msg()
- CVE-2020-28021: New-line injection into spool header file (remote)
- CVE-2020-28022: Heap out-of-bounds read and write in extract_option()
- CVE-2020-28026: Line truncation and injection in spool_read_header()
- CVE-2020-28019: Failure to reset function pointer after BDAT error
- CVE-2020-28024: Heap buffer underflow in smtp_ungetc()
- CVE-2020-28018: Use-after-free in tls-openssl.c
- CVE-2020-28025: Heap out-of-bounds read in pdkim_finish_bodyhash()

Qualys has published a detailed write up showing step-by-step code mistakes and exploitation mechanisms — but not working exploits. However, since Exim is open source and published under the GNU GPL, there's no point in attempting to obfuscate. 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. You know it's going to happen.

I'm not going to get into blow-by-blow detail here. It's all available on Qualys' excellent vulnerability disclosure which I've linked to here in the show notes:

<https://www.qualys.com/2021/05/04/21nails/21nails.txt>

But here's how they introduced their research:

We recently audited central parts of the Exim mail server and discovered 21 vulnerabilities: 11 local vulnerabilities, and 10 remote vulnerabilities. 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 4 LPEs (Local Privilege Escalations) and 3 RCEs (Remote Code Executions):

- We will not publish our exploits for now; instead, we encourage other security researchers to write and publish their own exploits:
- This advisory contains sufficient information to develop reliable exploits for these vulnerabilities; in fact, we believe that better exploitation methods exist.
- 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 exploits.
- We will answer (to the best of our abilities) any questions regarding these vulnerabilities and exploits on the public "oss-security" list (<https://oss-security.openwall.org/wiki/mailling-lists/oss-security>).

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 <https://www.qualys.com/research/security-advisories/>.

<https://blog.qualys.com/vulnerabilities-research/2021/05/04/21nails-multiple-vulnerabilities-in-exim-mail-server>

In their disclosure, Qualys wrote: *"Once exploited, they could modify sensitive email settings on the mail servers, allow adversaries to create new accounts on the target mail servers."*

Exim has a history of trouble. Back in June of 2019, Microsoft warned of an active Linux worm targeting an earlier Exim RCE bug. And one month later, attackers started exploiting vulnerable Exim servers to install the Watchdog Linux trojan, adding them to a Monero cryptomining botnet. And the US National Security Agency said last May 2020 that the "Sandworm" Russian military hackers have been exploiting that same critical Exim RCE since at least August 2019.

The Microsoft Exchange Server catastrophe showed us just how vulnerable an exploitable eMail server can be. Now, the whole world knows that Exim, the world's most widely deployed eMail server can also be remotely exploited. And 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.

Tor's Exit Nodes

Since 2015, a Tor network researcher who uses the moniker "nusenu" has been tracking the deliberate abuse of the Tor network by quite determined — and lately quite increasingly determined — attackers. 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, Leo often cites the dangers inherent in Tor exit nodes. Once you hear what this researcher has been tracking, I doubt that anyone will or should feel comfortable using Tor without the added protection of a VPN.

Because this was fascinating to me I've read most of what this researcher has been tracking and discovering. "Nusenu's" most recent posting to medium, two days ago, was titled "Tracking One Year of Malicious Tor Exit Relay Activities (Part II)." He starts out explaining:

In August 2020 I reported about "How Malicious Tor Relays are Exploiting Users in 2020 (Part I)". Back then I made the hypothesis that the entity behind these malicious tor relays is not going to stop its activities anytime soon. Unfortunately this turned out to be true. In this follow-up post, 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.

Before I proceed to share the extent of the trouble that nusenu has uncovered, I want to explain the mischief that the bad guys are getting up to with Tor exit nodes. In his August 2020 posting nusenu explains:

What is this attacker actually exploiting and how does it affect Tor users?

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 on Tor users by manipulating traffic as it flows through their exit relays. They (selectively) remove HTTP-to-HTTPS redirects to gain full access to plain unencrypted HTTP traffic without causing TLS certificate warnings.

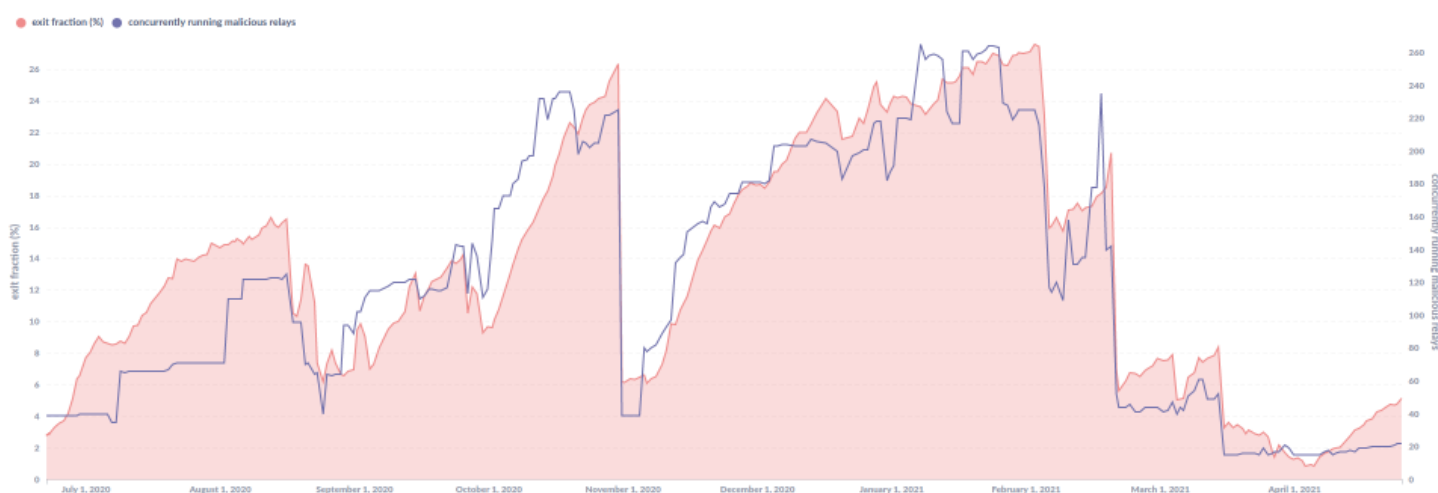
[We've spoken of this many times. GRC, for example, redirects anyone coming in over http to https. It is not possible to access GRC without https, though it is possible to begin with http and then be moved over to https to continue. While web browsers all assumed http, this was a

necessary step since everyone entering grc.com would default to http://grc.com. I should note that, as our listeners will recall, GRC was among the first domains to be added to Chrome's permanent HSTS list which Mozilla duplicates. It explicitly gave Chrome and Firefox permission to always silently promote any and all http queries to https. Anyway, nusenu continues...]

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://". 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. 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 are primarily after cryptocurrency related websites — namely multiple bitcoin mixer services. They replaced bitcoin addresses in HTTP traffic to redirect transactions to their wallets instead of the user-provided bitcoin address. 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.

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 (HTTPS Everywhere is installed by default in Tor Browser). Unfortunately 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.

I'm 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'm less surprised that we keep hearing of 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 the chance to update.

So, just how big is the problem? Is it a couple of nodes that users are unlikely to exit from? The graph above plots the percentage of all Tor exit nodes that are known to be malicious. 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 network exit capacity) over time by this particular malicious entity between July 2020 and April 2021. Peak value: The attacker did manage approx. 27.5% of the Tor network's exit capacity on February 2ns, 2021. Graph by nusenu (raw data source: Tor Project/onionoo)

And he notes that while that's better than one in four, over time, since exit nodes are randomly chosen and rotated, the chance that a user will emerge from a maliciously-controlled exit node increases to near certainty.

The chart above shows sudden dramatic downward drops periodically whenever the Tor admins become aware that a provider of exit nodes is misbehaving. But nusenu notes that, invariably, every drop is immediately followed by a gradual return of malicious node count.

The bottom line here is, there is no free lunch. Tor provides some valuable services. But it's not a panacea. Any user of Tor should **assume** that the exit nodes they are 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 possible. Nusenu observes:

We know about mitmproxy, sslstrip, bitcoin address rewrites and download modification attacks, but it is 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.

Download modification attacks? Talk about chilling. You use Tor to go get something that you want to keep very private. But the website that offers it doesn't support https:. 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 obtained? Http offers zero authentication of the other end's identity.

A Tor HTTPS-Only browser would be one solution, and about that nusenu writes:

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, as a Tor Browser developer points out on a Tor mailing list:

When Tor Browser migrates to Firefox 91esr 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 (especially more region specific sites) and the question of what messaging Tor Browser should use in that case.

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 law enforcing, may be closely monitoring everything they can about a user's activities through Tor — and even actively modifying and subverting any traffic that's available in the clear. So, whenever using Tor, keep in mind the danger of http and the true need for some other privacy and security protecting tunnel, such as a trustworthy VPN.

- <https://nusenu.medium.com/the-growing-problem-of-malicious-relays-on-the-tor-network-2f14198af548>
- <https://nusenu.medium.com/how-malicious-tor-relays-are-exploiting-users-in-2020-part-i-1097575c0cac>
- <https://nusenu.medium.com/tracking-one-year-of-malicious-tor-exit-relay-activities-part-ii-85c80875c5df>

Miscellany

Project Hail Mary: A Novel

by Andy Weir (famous for having written "The Martian")

Available from Audible. / 5 stars.

"Nick" / 5.0 out of 5 stars / Amazing. Wonderful. Excellent.

Reviewed in the United States on May 4, 2021

Verified Purchase

I don't even remember pre-ordering this book. It just showed up in my Kindle app this morning. 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.

"The Martian" was a great story. "Artemis" was a great story. This one is better than either of those. 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.

5.0 out of 5 stars / Andy does it again!

Reviewed in the United States on May 5, 2021

Verified Purchase

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.

5.0 out of 5 stars / Stop reading this review. Read "Project Hail Mary".

Reviewed in the United States on May 5, 2021

Verified Purchase

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! Instant classic.

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.

Closing The Loop

From: Paul Babiak

Newsgroups: grc.securitynow

Subject: One possible solution to QNAP vulnerability

A walk-through of an installation of OpenMediaVault on a QNAP:

https://www.youtube.com/watch?v=iTBxua_7D_k

<https://www.openmediavault.org/>

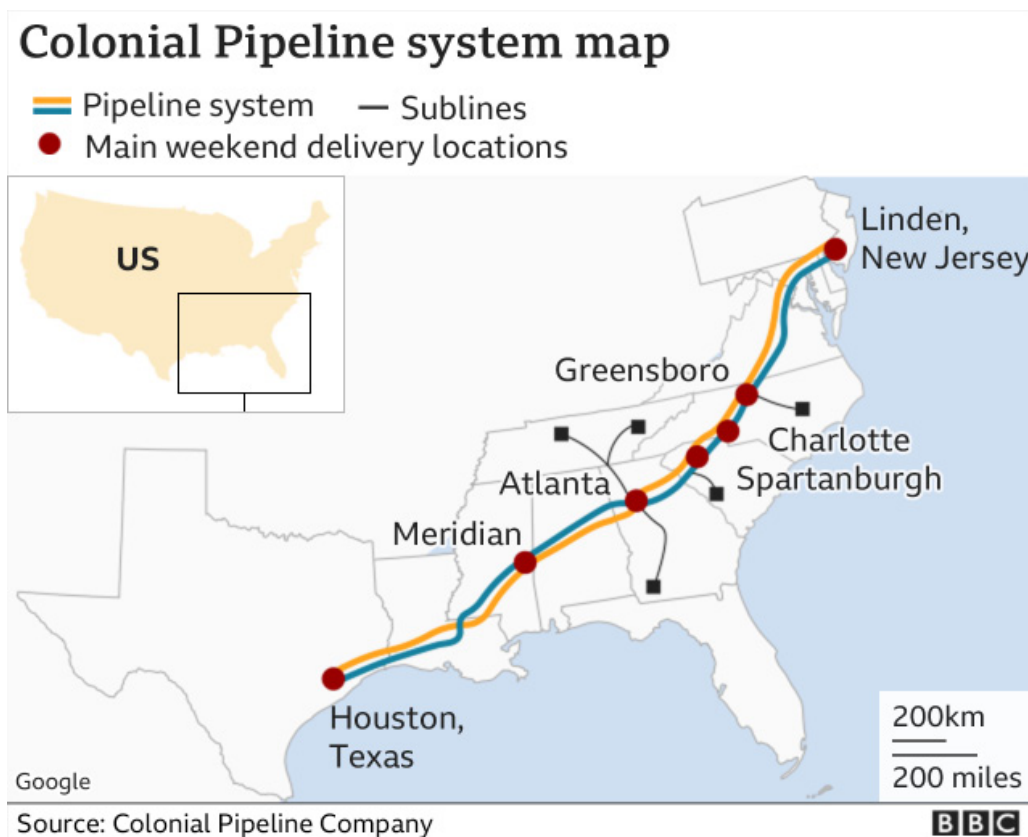
Jon S. @studerje

First hack that hits close to come. Sitting in the ER of Scripps health with my wife. They were "hacked" a few weeks ago and are still doing all charts and orders via paper records. The process is taking about 4-6 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. This upsets me to no end. Thought I'd share a few pictures of observations.

SpinRite

Nothing huge to report on the SpinRite front. I'm 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-2.2 terabyte drives containing any partitioning and any file system. Also, since we'll be living with, and using this code base after it's converted from 16-bit real mode segmented to 32-bit protected mode flat model, for booting under UEFI and BIOS and to host native operation on USB and NVMe mass storage, I'm also taking some time to clean things up a bit while I'm there to get it more ready for its future. Now that I have access to upper memory, I'm able to move some things up there to ease the pressure on lower memory to eliminate some jumping through hoops.

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,” was shutdown 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.

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 (RaaS) group which first appeared on the scene last August 2020.

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 between refineries located down in the Gulf Coast to markets throughout the southern and eastern US. When its pipeline is up and running, it transports 2.5 million barrels per day through 5500 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 have been declared by the Biden administration for Washington D.C. and 17 states. This was done to temporarily 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 105 million gallons of refined fuel per day.

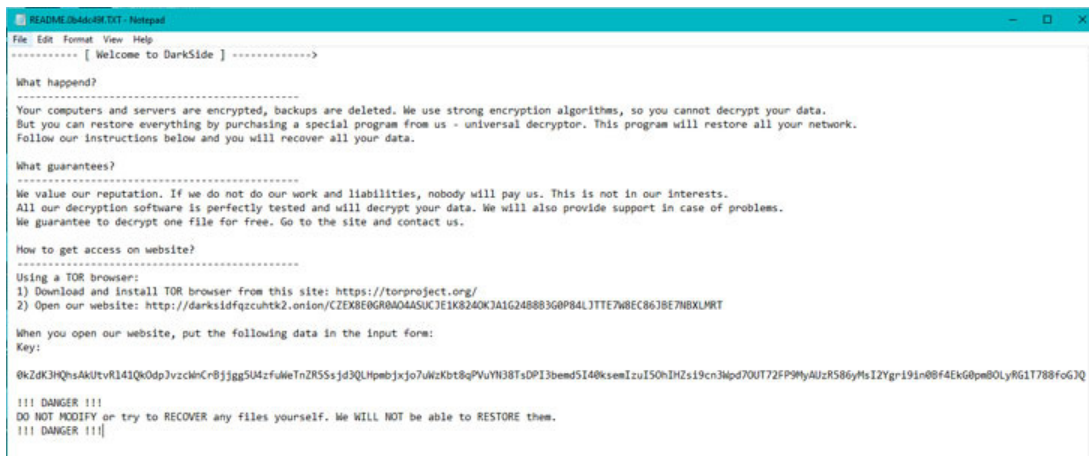
The Governor of Virginia today just declared their own state of emergency. Their declaration begins:

On this date, May 11, 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 cyber-attack on its business systems' informational technology infrastructure on May 7, 2021. 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.

The now famous "Solarwinds" attack made the news in March because it was labeled the most significant cyberattack ever. So, Woooo! — 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, is different. 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 US, the problem is no longer theoretical or superficial.

So what about "Darkside" ?? I found a copy of their extortion demand note:



```
----- [ Welcome to DarkSide ] ----->

What happend?
-----
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. This program will restore all your network.
Follow our instructions below and you will recover all your data.

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.

How to get access on website?
-----
Using a TOR browser:
1) Download and install TOR browser from this site: https://torproject.org/
2) Open our website: http://darksidfzcuhtk2.onion/CZEXBE0GR0A04ASUCJE1K8240KJA1G248883G0P84LJTTE7W8EC863BE7NBXLMTT

When you open our website, put the following data in the input form:
Key:
0kZdK3HqhsAKUtrR141Qk0dpJvzckncr8jgg5U4zfuWeTnZR5SsJd3QLHpebjxjo7uWzKbt8qPVuYN38TsDP13beed5I40ksemIzu150nIH2s19cn3Mpd70UT72FP9MyAUzRS86yHs12Ygr191n08f4EkG0pe80LyfG1T788foG7Q

!!! DANGER !!!
DO NOT MODIFY or try to RECOVER any files yourself. We WILL NOT be able to RESTORE them.
!!! DANGER !!!
```

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 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.

They appear to imagine that they are running a business more than a crime ring. When they released a new version of their software two months ago which could encrypt data faster than before, they issued a press release and invited journalists to interview them. And their website on the dark web lists all the companies they have attacked and hacked and what was stolen.

And, they have an “ethics” page listing which types of organisations they will not attack. They have stated that they will not attack hospitals, hospices, schools, universities, non-profit organizations, or government agencies. I suppose after this they’ll be adding “critical infrastructure” to that list.

So what’s somewhat different about these particular criminals is that they say they do not intend to cause harm. They just want money. On their 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 that they have probably painted a huge bullseye 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 and check each company that our partners want to encrypt to avoid social consequences in the future.”

We know that they use the Salsa20 symmetric cipher with a custom matrix and RSA-1024 for public key operations. And their ransoms have generally ranged from \$200,000 to \$2,000,000.

Much of this podcast is focused upon developing an understanding of just exactly how porous most of our computer and network security is today. We look at the details and attempt to determine why these problems happen and what might be done to prevent such trouble in the future.

One thing is very clear: Most, if not all, of our existing IT infrastructure is not ready to stand up to determined attack. It’s a sad fact that it needs to. What a bunch of time, effort and money that wastes.

