

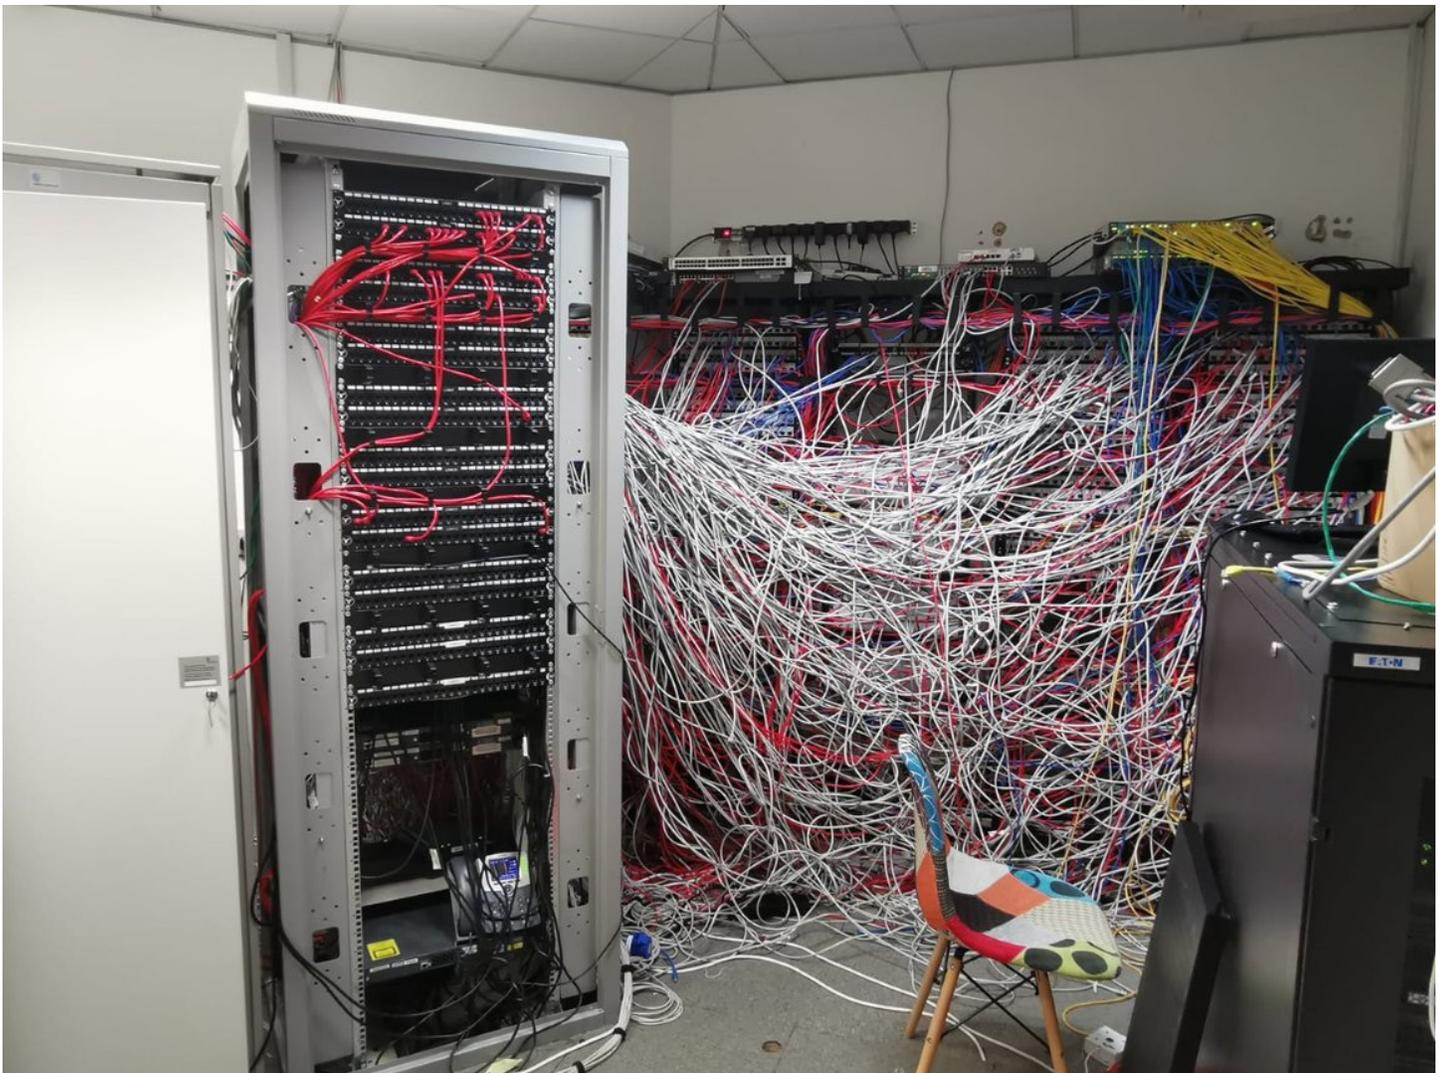
Security Now! #830 - 08-03-21

The BlackMatter Interview

This week on Security Now!

This week we look at FireFox's declining active user count, at the evolution of the Initial Network Access Broker world, at several different ransomware group renamings and revivals and we encounter a well-informed Active Directory security researcher who feels about Microsoft's July pretty much as we do. I want to turn our listeners onto a very interesting looking Hamachi'esque overlay for WireGuard and share a fun diagnostic anecdote that cost me a day of work last Friday. We have a bit of closing the loop feedback from a couple of our listeners, then we're going to share an interview with a member of the "maybe new or maybe rebranded" ransomware group BlackMatter which Recorded Future posted yesterday.

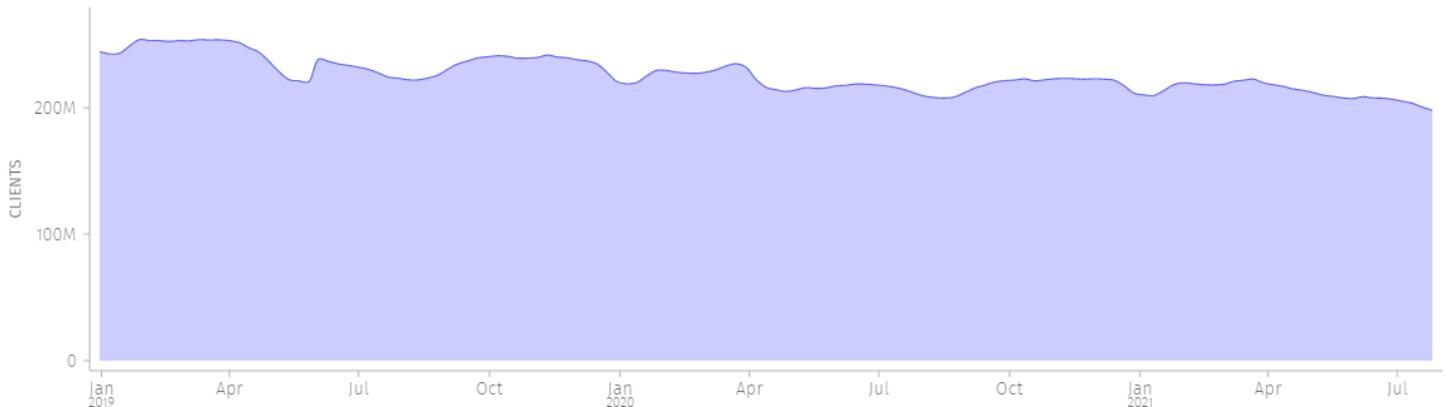
How does this happen?



Browser News

Mozilla's Firefox Monthly Active Users (MAU) slowly but steadily drops

<https://data.firefox.com/dashboard/user-activity>



The chart above depicts the Monthly Active Users, a measurement of the number of Firefox Desktop clients active during the previous 4 weeks (28 days). As seen above, the number peaked on Jan 27th, 2019 at **253,877,800** active users and its lowest point, seen in recent history, was where the charting ends on July 25th, 2021 with **197,874,100** Firefox users being active over the previous 4 weeks. This is a measured drop of **56,003,700** users.

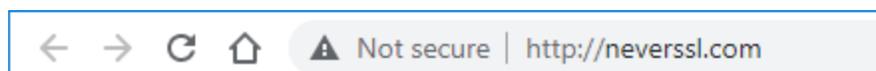
What does this mean? What is this telling us? Leo?

Google to finally assume HTTPS

We're currently at release 92 of Google's Chrome browser. As you surf around in these post-Snowden days, you'll mostly be seeing a little black padlock to the immediate left of the browser's current URL.



But if you manage to land on one of the increasingly rare non-TLS sites you'll see Chrome's explicit warning that the site you're on is "Not Secure":



That declaration has always annoyed me a bit, since there are completely legitimate sites that have absolutely no need for TLS security that Google is inherently smearing with scary insecurity warnings for without cause. But I get it that Google is wishing to warn and protect users when TLS is not present, and that sites which truly have no need for any of TLS's many protections are becoming few and far between. I had to track down the <http://neverssl.com/> site to capture that screenshot for the show notes. And just a tip that if you're wanting to see what a browser shows over anon-HTTPS plain HTTP connection, the <http://neverssl.com/> site will always do the job. We're talking about this again, because Chrome's default URL display will be changing again,

probably to its final form, with its next major release, 93. Google has decided that since HTTPS has finally become the universal norm, it need not say anything about those connections. So don't be surprised when that little featureless black padlock icon disappears completely.

Ransomware News

The evolution of "Initial Access Brokers"

As we know, unfortunately "Ransomware" has rapidly grown into an established and entrenched form of cybercrime that's not going away anytime soon. The first aspect of specialization that emerged was the idea of separating the developers of the ransomware from its use to attack victims. Some evil genius conceived of the notion of ransomware as a service. This created the so-called affiliates who would perform the attacking using rented ransomware. And it wasn't long before the affiliate role also split and further specialized into so-called Initial Access Brokers and pared down affiliates who purchased access to valuable illegal enterprise networks from those 3rd-party brokers.

Yesterday, the cybersecurity firm KELA (<https://ke-la.com/>) published a report documenting their year long exploration into the dark web and, specifically, the nature of the market that's forming around Initial Access Brokers.

<https://ke-la.com/all-access-pass-five-trends-with-initial-access-brokers/>

Their report provided so much interesting detail that it was originally going to be the main focus and topic of today's podcast. But that plan was preempted by what did become today's main focus and topic. So, instead of digging deeply into every detail of what Kela found and reported I'm going to summarize what they discovered about the nature of today's Initial Access Broker market. Yesterday, Kela framed their research by explaining:

For more than a year, KELA has been tracking Initial Access Brokers and the initial network access listings that they publish for sale on various cybercrime underground forums. Initial Network Access refers to remote access to a computer in a compromised organization. Threat actors selling these accesses are referred to as Initial Access Brokers. Initial Access Brokers play a crucial role in the ransomware-as-a-service (RaaS) economy, as they significantly facilitate network intrusions by selling remote access to a computer in a compromised organization and linking opportunistic campaigns with targeted attackers, often ransomware operators.

This research includes an in-depth analysis of Initial Access Brokers and their activity for a full year from July 1, 2020 to June, 30 2021. KELA analyzed IABs' activities over the last year (when their role became increasingly more popular in the cybercrime underground) and summarized 5 major trends that were observed throughout our analysis.

Kela's primary research takeaways include:

- KELA explored over 1000 access listings offered for sale over the last year. The average

price for network access during this period was 5,400 USD, while the median price was 1,000 USD. The top affected countries included the US, France, UK, Australia, Canada, Italy, Brazil, Spain, Germany, and UAE.

- IABs built a pricing model for initial access. The most valuable offers include domain admin privileges on a computer within a company with hundreds of millions in revenue.
- With RDP and VPN-based access being the most common offer, IABs find new attack vectors and accommodate the changing software targets of ransomware gangs, including network management software and virtual servers.
- Successful IABs find regular customers, some of which are ransomware affiliates, and move most of their operations to private conversations. However, new actors continually enter the scene.
- Some IABs adopted ethics that were introduced by some ransomware gangs. Namely, there is a certain criticism against actors trading access to healthcare companies, though it's still an initiative of a few actors and not a typical attitude.
- IABs are eager to monetize their access and are using all means to do so. Some IABs were seen stealing data from the affected company to gain profits even if the access is not bought.
- IABs have become professional participants of the RaaS economy. They constantly find new initial access vectors, expanding the attack surface, and follow their customers' demands. It requires network defenders to track IABs activities and all other actors who have formed around ransomware.

KELA found that, not surprisingly, an important metric for setting the access price is the level of privilege that the access enables with domain admin access being the most expensive, costing at least 10 times more than access to a machine with standard user rights. The priciest offers from reputable threat actors KELA observed included:

Access to an Australian company with 500 million USD in revenue that enables an attacker with "admin" level of privileges (most likely domain admin) offered for 12 BTC (\$465,168)

Access through ConnectWise to a US IT company offered for 5 BTC (\$193,820)

Access to a Mexican government body offered for 100,000 USD and used for the LockBit ransomware attack

KELA is also seeing a diversification of access as the "market" (I so much dislike using terms which tend to legitimize this illegal behavior.) But a market it is. That's the entire point of spilling off the IAB role. So, as I was saying. KELA is seeing a diversification of access being demanded by a growing market. In this marketplace the term "network access" is loosely defined; threat actors use it to describe multiple vectors, permission levels, and entry points. Over a year, KELA observed that the most commonly offered access is RDP- and VPN-based access, usually provided in the form of valid logon credentials. Remote access can also be supplied through the ConnectWise and TeamViewer software, which can provide actors with RDP-like capabilities. VPN access can be gained and sold through various software, such as Citrix, Fortinet, and Pulse Secure products, all of which are popular among those in the cyber underworld.

In addition, IABs are finding new attack vectors and ways to supply access to buyers, meaning that the overall attack surface is expanding. For example, access to VMWare's ESXi servers which have recently become quite popular among ransomware attackers — REvil and Darkside both had versions of their malware specifically targeting ESXi.

Tracking and counting IAB transactions is difficult because once a IAB's advertisement is responded to by a credible and interested buyer, nothing further appears after they take their communication private. And it also appears that after the parties have first met through a public advertisement, and a relationship is formed where the buyer may say "The quality of the access you're selling is good. We'll buy anything else you have to sell for a fair price. So let us know first since we may pay top dollar ...or... top ruble."

KELA also confirmed the trend we've seen reported about the (so-called) "professional ethics" — at least with regard to who is an acceptable victim. Before their disbanding and disappearance we saw the DarkSide gang promise not to target certain sectors. This trend is spreading and solidifying, though it's not yet fully established and it varies depending upon the specific gang. But there have been bans on attacking healthcare, government, education, and non-profit sectors so as to not cause damage to patients, students, citizens, and other categories of people. The ransomware gangs appear to be passing the message that they will only hunt companies and aim only for financial gain.

In line with this, IABs have been seen posting access ads for victims from the healthcare sector then later deleting the offers after receiving public criticism from other users. However, there are still no hard and fast rules on this matter with most brokers being glad to sell all the accesses they're able to gain. And as always, IABs trading in Russian-speaking forums do not attack Russia, following the rules of the forums.

DarkSide Returns (that didn't take long)

As we know, an affiliate of the Ransomware as a Service (Raas) group known as DarkSide was (and "was" is the operative word here) made what turned out to be a big mistake by attacking critical US infrastructure in the now famous Colonial Pipeline attack. This shut down the US's largest fuel pipeline, causing fuel shortages across the Eastern seaboard of the US. That's what's known as "gaining unwanted attention." And shortly afterward, DarkSide's ransomware operation suddenly shut down after they lost access to their servers, and at least some of their ill gotten cryptocurrency was seized. We later learned that the FBI had somehow recovered 63.7 Bitcoins of the approximately 75 Bitcoin (\$4 million) ransom payment made by Colonial Pipeline.

Since ransomware—when it's done right—can generate so much money, no one thought for a moment that the culprits had learned their lesson and had decided to update their Linked-In resumes and start interviewing for jobs in the Russian IT sector. What we all thought was that they would return under another name, probably with at least the intention of not again stepping in such a big pile of trouble.

And sure enough, a recent detailed forensic analysis of the cryptographic algorithms being employed by an apparent newcomer named "BlackMatter" suggests that BlackMatter is actually

DarkSide 2.0.

However, since this new group is soliciting Initial Access Brokers directly, it may be that they've scrapped their previous affiliate model, at least for now, probably in the interest of maintaining more control and thus preventing a recurrence of the disaster that shut them down last time.

This new "BlackMatter" group is actively attacking victims and purchasing network access from other threat actors to launch new attacks. Over this past weekend, BleepingComputer reported that multiple victims have been targeted by BlackMatter with ransom demands ranging from \$3 to \$4 million and that one victim had already paid \$4 million to delete stolen data and receive both a Windows and Linux ESXi decryptor.

So, why do we believe this is the work of DarkSide rebranded? Over this past very busy weekend, EMSIsoft's Fabian Wosar tweeted:

Fabian Wosar / @fwosar

After looking into a leaked BlackMatter decryptor binary I am convinced that we are dealing with a Darkside rebrand here. Crypto routines are an exact copy pretty much for both their RSA and Salsa20 implementation including their usage of a custom matrix.

Okay, what does he mean by that?

Salsa20 is a stream cipher which XOR's any data it's given with the output of its cryptographically strong pseudorandom bit generator. As we know from our many previous discussions about the operation of cryptographic routines, As counter-intuitive as it may seem, simply inverting random bits of a plaintext—XORing the plaintext with noise—totally scrambles that plaintext to yield a cryptographically strong ciphertext. And when that ciphertext is later re-XORed with the same random bitstream, the original plaintext is restored. That's Salsa20.

Salsa20's internal state is held in 16, 32-bit quantities which are conceptually arranged in a 4x4 matrix. The formal Salsa20 spec. specifies how those 16 values are to be initialized, even though it's at least somewhat arbitrary. Interlaced through four of the 32-bit values is the 16 character string "expand 32-byte k". Two of the sixteen 32-bit values are used to specify stream position, another two are nonces, and the remaining eight 32-bit values are used as the key. That's the formal spec. But Darkside blows all that off and simply initializes those words with random data then encrypts that Salsa20 state matrix with a public RSA key which is appended to the end of the encrypted file. Fabian Wosar said that this implementation of Salsa20 implementation was previously **only ever used** by DarkSide... and now this approach has resurfaced, being used by BlackMatter. Additionally, DarkSide's implementation of 1024-bit RSA was also unique to their encryptor, and BlackMatter also uses **the same** unique implementation.

When we also consider that both groups use similar language and color themes on their public and dark websites, evidence a similar lust for media attention, and that the "new" (in quotes) BlackMatter group is going to great pains to note that it will not target the "Oil and Gas industry (pipelines, oil refineries)" it seems about as certain as it could be that DarkSide has returned. And at the end of this podcast we're going to hear from the BlackMatter group, themselves.

And... "DoppelPaymer" becomes "Grief"

While we're on the topic of ransomware group rebranding I should also note that the previously

named "DoppelPaymer" group has apparently renamed itself "Grief". There's no big news about this new group, beyond the observation of their renaming for the record.

The evidence is similar to what we just looked at with DarkSide and BlackMatter. The technical crypto details of the two groups are unique and identical. Fabian Wosar of Emsisoft told BleepingComputer that the two shared the same encrypted file format and used the same distribution channel, the Dridex botnet. And despite Grief's efforts to appear as a separate RaaS, the similarities to DoppelPaymer are impossible to ignore. In fact, researchers at the cloud security company Zscaler analyzed an early Grief ransomware sample and noted that the ransom note dropped on infected systems pointed to the DoppelPaymer portal. Whoopsie! And both use very similar code that implements identical encryption algorithms (2048-bit RSA and 256-bit AES), import hashing, and code entry point offset calculations.

Finally, both Grief and DoppelPaymer point out that the European Union General Data Protection Regulation (GDPR) could impact their non-paying victims who might still face legal penalties due to information disclosure following the breach.

So, it seems very clear that we are seeing, and doubtless will be seeing, criminal ransomware groups rebranding and renaming themselves as needed to suit whatever perceived need they may have to shed their previous skins.

And, "Avaddon" has become "Haron"

One last note is that the ransomware group formally known as "Avaddon" is now calling itself "Haron." Similar evidence, nearly identical before and after web sites and text wording. So it has become quite clear that these groups place very little stock in their names. And being criminal enterprises that tend to keep getting themselves into trouble, I suppose that's not too surprising.

Security News

"A Microsoft July 2021 Recap"

<https://github.com/cfalta/MicrosoftWontFixList>

I wanted to begin with a brief look back at the month we've just survived. I found a nice set piece for this in the form of a little corner of GitHub which belongs to a guy named Christoph Falta of Vienna, Austria. He describes himself as "Random infosec guy. Rainbow-teamer. Focusing on windows security" and his past work clearly shows that he has an interest in Microsoft's Windows Active Directory. He doesn't follow me on Twitter, but I have many fewer followers than we have podcast listeners. So perhaps he's a Security Now! listener. I mention that because it sure sounds as though Christoph is channeling me and the podcast on this latest page of his, which is titled: "Microsoft Wont-Fix-List (July 2021 Edition)"

He updated the page yesterday, posting with yesterday's date:

02.08.2021 - Update: thank you all for your feedback :-)

summary of what happened in July 2021 and I decided I'll keep it that way, because I honestly think I don't have the energy to maintain an up-to-date list of ALL won't fixes Microsoft has to offer. So I'll keep this remark here for clarity and change the description.

A list of vulnerabilities or design flaws Microsoft does not intend to fix. Since the number is growing, I decided to make a list.

LPE = Local Privilege Escalation

DPE = Domain-wide Privilege Escalation

RCE = Remote Code Execution

The page consists primarily of a table with columns labeled: vulnerabilities, associated/assigned CVE's, Attack Type descriptions, "It's NTLM again, right?" and "How it works, in a nutshell". And, essentially, it's everything that we've been talking about throughout this past July.

Under vulnerabilities he lists: SpoolSample, PetitPotam, RemotePotato0, SeriousSAM, PrintNightmare, ADCS-ESC8. A couple of the vulnerabilities we've talked about have been given names "SpoolSample" and "RemotePotato0", and they are issues we've discussed previously.

What was heartening to me was to see that at least I'm not alone in the overall sense I came away from July with, which I've been conveying through this podcast, that not only was July 2021 an unusually rough month for Microsoft, but that what appears to be emerging is a longer-term problem with Microsoft's legacy protocols which are all still enabled by default "just in case" and that some of these problems "work as designed" (as Christoph terms it in his CVE column where there is no CVE because it's not a bug, it's a feature.

It's probably the case that the IT staff in many enterprises have become accustomed to assuming that whatever's wrong with Windows will be auto-patched as soon as Microsoft can get around to it. And while that might not be soon enough, in the case of the ProxyLogon debacle with Microsoft's Exchange server early this year, Microsoft does eventually get their machines patched. So applying updates is clearly crucial.

But the shift that July's revelations of significant problems which all "work as designed" creates means that simply applying Microsoft's monthly patch extravaganza will no longer mean that an enterprise's network is being kept safe and secure. When Microsoft wrote "vulnerable by design" they meant it. And that means that there's no patch forthcoming, late or ever.

So I wanted to make sure that our listeners fully realized that we've entered a bit of a different world with the discovery, publication and multiple proofs-of-concepts of these vulnerabilities that Microsoft says they have no plans to fix—because nothing's broken. What's going to be needed is for those who are responsible for their enterprise's network security to go beyond pressing Microsoft's "update our systems" button and looking carefully at these multiple edge and corner cases to determine how to set enterprise wide policies to disable these dangerous features on a case by case basis.

Miscellany

Tailscale

As we know, WireGuard is widely, and I think accurately, regarded as the logical and overdue successor to the venerable OpenVPN... and even to some degree, to IPSec. OpenVPN, like OpenSSL, is suffering from its age and from the fact that it has been, for decades, the test bed for many experiments as we've been learning the right way to do things only after first doing many of those things wrong. For example, TCP cannot be the protocol used by a VPNs tunnel. It's not what it was designed for and doing so is just wrong. Yet OpenVPN offers the option. So once all of the wrong solutions have been tried and the right solutions have been found, it's really best to lighten one's load and just start over again from scratch with a blank slate that can host an entirely new design. That's what Jason Donenfeld set out to create when he launched what has turned out to be the incredibly successful WireGuard project.

Exactly three years ago from yesterday, Linus Torvalds posted the following:
<https://lists.openwall.net/netdev/2018/08/02/124>

Date: Thu, 2 Aug 2018 10:15:40 -0700
From: Linus Torvalds <torvalds@...ux-foundation.org>
Subject: Re: [GIT] Networking

Btw, on an unrelated issue: I see that Jason actually made the pull request to have wireguard included in the kernel.

Can I just once again state my love for it and hope it gets merged soon? Maybe the code isn't perfect, but I've skimmed it, and compared to the horrors that are OpenVPN and IPSec, it's a work of art.

Linus

Wikipedia reminds us that:

WireGuard is a communication protocol and free and open-source software that implements encrypted virtual private networks (VPNs), and was designed with the goals of ease of use, high speed performance, and low attack surface. It aims for better performance and more power-saving than the IPsec and OpenVPN tunneling protocols. The WireGuard protocol passes traffic over UDP.

In March 2020, the Linux version of the software reached a stable production release and was incorporated into the Linux 5.6 kernel, and backported to earlier Linux kernels in some Linux distributions. The Linux kernel components are licensed under the GNU General Public License (GPL) version 2; other implementations are under GPLv2 or other free/open-source licenses.

WireGuard utilizes:

- Curve25519 for key exchange
- ChaCha20 for symmetric encryption
- Poly1305 for message authentication codes

- SipHash for hashtable keys
- BLAKE2s for cryptographic hash function
- UDP-based only

So, now, with the reminder of what WireGuard is, and why we should love it, I want to introduce everyone to TailScale which two of our listeners, Ben Hutton and Jack Hayter, recently turned me on to. A search on the phrase "WireGuard versus TailScale" brought me to a page asking and answering exactly that question...

Should I use TailScale or WireGuard® to secure my network? The answer is yes!

[TailScale's founders—who are, by the way, some highly credentialed ex-Google and Alphabet developers—then explain...]

TailScale is built on top of WireGuard; we think very highly of it.

We designed TailScale to make it easier to use WireGuard to secure your network connections. You might decide to use WireGuard directly, without TailScale. This is a guide to using TailScale vs. configuring and running WireGuard directly.

[So we sort of have TailScale providing a connectivity layer with WireGuard providing the super-secure packet level transport.]

Configuration

WireGuard is typically configured using the wg-quick tool. To connect two devices, you install WireGuard on each device, generate keys for each device, and then write a text configuration for each device. The configuration includes information about the device (port to listen on, private IP address, private key) and information about the peer device (public key, endpoint where the peer device can be reached, private IPs associated with the peer device). It's straightforward, particularly for a VPN. Every pair of devices requires a configuration entry, so the total number of configuration entries grows quadratically in the number of devices if they are fully connected to each other.

To connect devices using TailScale, you install and log in to TailScale on each device. TailScale manages key distribution and all configurations for you. This can be particularly useful if some of the devices belong to non-technical users.

Connectivity

WireGuard ensures that all traffic flowing between two devices is secure. It does not ensure that those devices can connect; that is up to you. WireGuard has a persistent keepalive option, which can keep the tunnel open through NAT devices. But in some cases to ensure that your devices can communicate, you may need to open a hole in your firewall or configure port forwarding on your router. WireGuard can detect and adapt to changing IP addresses as long as a connection remains open and both ends do not change addresses simultaneously. Establishing a connection or re-establishing a broken connection requires updating configuration files.

TailScale takes care of on-demand NAT traversal so that devices can talk to each other directly in most circumstances, without manual configuration. When NAT traversal fails, TailScale relays encrypted traffic, so that devices can always talk to each other, albeit with higher

latency in that case. There is no need to modify firewalls or routers; any devices that can reach the internet can reach each other. (Tailscale traffic between two devices on the same LAN does not leave that LAN.)

Security

Tailscale and WireGuard offer identical point-to-point traffic encryption.

Using Tailscale introduces a dependency on Tailscale's security. Using WireGuard directly does not. It is important to note that a device's private key never leaves the device and thus Tailscale cannot decrypt network traffic. Our client code is open source, so you can confirm that yourself.

With the Team and Business plans, Tailscale adds an ACL layer on top of WireGuard, so that you can further control network traffic. You can do some of this directly with WireGuard by not setting up tunnels between devices that should not communicate or by using the operating system firewall to control traffic flow. Tailscale ACLs allow you to express ACLs for everything in a single place using users, groups, and tags, which are easier to maintain than a list of which device pairs may communicate

Even without the Team or Business plan, Tailscale offers some basic, unidirectional ACL controls. For example, any node may turn on "Shields Up" mode, which prevents all incoming connections.

[I got a kick out of that. And I should mention that TailScale is completely free for personal use with a single user, Single-Sign On, multi-factor authentication and linking 20 devices. Multiple users, access control lists, advanced network segmentation and other group features are billed per user month.]

Performance

Using WireGuard directly offers better performance than using Tailscale. Tailscale does more than WireGuard, so that will always be true. We aim to minimize that gap, and Tailscale generally offers good bandwidth and excellent latency, particularly compared to non-WireGuard VPNs.

The most significant performance difference is on Linux. On Linux, WireGuard is available as a kernel module. Tailscale currently uses the userspace WireGuard implementation, which has more overhead.

The most common scenario in which Tailscale users notice bandwidth or latency issues is when Tailscale is relaying network traffic, which is unavoidably slower. In that case, the devices would be unable to connect at all using WireGuard directly, so no direct comparison is available.

Bonus features

By design, WireGuard provides secure point to point communication. It is intended to be a building block.

Tailscale has a broader set of features. For example, we offer MagicDNS to make it easier to reach other devices on your VPN. We have out of the box support for subnet routing to allow employees access to an office network via an exit node running Tailscale. And more features are in the works.

IT/network administration

When using WireGuard directly, you may use any tools desired to administer your network. There is an active community that can answer questions on IRC or a mailing list.

Tailscale's focus on convenience makes many IT requests self-service. Tailscale has an admin panel on our website. As of Dec 2020, Tailscale's admin API is in beta and available by request. Tailscale offers community support for our free pricing tiers and direct support for all paid plans.

The bottom line

We suspect that using WireGuard directly will be most appealing if you have a small, stable number of Linux servers whose connections you want to secure. Using Tailscale will make the most sense if you want things to Just Work, you are administering a VPN for many different users, or if you want the extra features or centralized ACLs Tailscale offers.

But everyone's network and needs are different. And we've helped debug a lot of networks; when we say everyone's network is different, we know whereof we speak, and we mean it!

Using WireGuard directly is a very reasonable choice, and if you're thinking about doing it, we encourage you to give it a try. If you later decide that you want the convenience and extra features that Tailscale offers, it's easy to switch.

Last updated Jul 29, 2021

It turns out that this paired combination is almost exactly what I had planned for my own commercial CryptoLink project. But, as we know, I chickened out due to my concern that governments were eventually going to require what we refer to in the US as warrant-compatible encryption. I would never have agreed to weaken CryptoLink, and I wasn't willing to invest years of work in something that I hoped to package up as a commercial product only to have its use made illegal.

And, as it turns out, I'm quite happy with that decision. I'm SO happy to be back in the saddle with SpinRite, whose future is brighter than ever, since it promises to be able to offer early warning of solid state mass storage failure, with the ability to refresh slow-reading aging data before it's too late. And I plan to follow that with "Beyond Recall" to make super-secure data wiping fast, easy and very reliable for non-techies.

While I love the idea of having multisite networks statically glued to each other into a single big privately routable network, the downside is the threat presented by today's ransomware. I bring inter-network connections up when and as I need them and then right down afterward. Perhaps I've been listening to my own podcast for too long, but I fully realize that it's just too easy to make a mistake and let something in. So, convenient as it would be for me to have static network linkages to GRC's servers, I never allow myself that luxury. The convenience would be nice, but for me it's not worth the pain that would ensue from a mistake. I've offset that by making it quick and easy to bring links up and the way I have set things up, I don't have any automated background processes that need to have inter-network communication. I wish that weren't the world we're living in today, but it is.

Closing the Loop

Håkan Lindqvist / @HakanLindqvist

Re: SN 829

Regarding the previous slowness of Chrome's rather fascinating color-based phishing site detection in relation to auto-filling passwords, that does not seem like it would be an actual thing? If you go to a phishing site, the browser or the password manager would not have any saved passwords matching the site URL, so there would be no auto-fill available (a warning sign in itself to a perceptive user). The actual problem for password forms would seem to be limited to the user manually entering or copy/pasting the password, which is less likely to win the race.

awk @adrianteri

Good day Steve, from this weeks episode 829 at around 24 mins when talking about Chrome's 92 phishing detection you mention something about automated form fields. I take this to be form fills by password managers? I thought one of the selling points of passwords managers was to prevent you from this as the site(s) doesn't match any URL in your passwords database. In fact I get wary of a site I am logging into if I am forced to manually copy and fill the form from searching my vault entries.

They are both right, of course. That's an advantage of using any of the forms of credential autofill that I fail to mention often enough. When you visit a site with a spoofed lookalike domain name that doesn't trigger your browser's or addon's auto-fill, that's an immediate tip off that something's wrong and not to proceed. So as Hakan reminded me, the worry about hitting "Login" before Chrome has finished its known spoofed site detection would not be a problem.

Clay Seale / @TexasLazyK

Steve: Any recommendations after #ProjectHailMary? It was an outstanding read!!

"The Bobiverse" trilogy is highly and often recommended by our listeners as being fun and a bit whimsical, so in some ways reminiscent of Project Hail Mary. The teaser about the book reads:

Bob Johansson has just sold his software company and is looking forward to a life of leisure. There are places to go, books to read, and movies to watch. So it's a little unfair when he gets himself killed crossing the street.

Bob wakes up a century later to find that corpsicles have been declared to be without rights, and he is now the property of the state. He has been uploaded into computer hardware and is slated to be the controlling AI in an interstellar probe looking for habitable planets. The stakes are high: no less than the first claim to entire worlds. If he declines the honor, he'll be switched off, and they'll try again with someone else. If he accepts, he becomes a prime target. There are at least three other countries trying to get their own probes launched first, and they play dirty.

The safest place for Bob is in space, heading away from Earth at top speed. Or so he thinks. Because the universe is full of nasties, and trespassers make them mad - very mad.

I do not yet have any firsthand knowledge or recommendation, myself. I'm currently on book #19 of my incredibly enjoyable re-read of Ryk Brown's 30-books-so-far-out-of-his-planned-75 Frontiers Saga. I LOVE the Frontiers Saga. As we know, I LOVE reading science fiction. So perhaps the bar isn't that high, and it does feel as though it's time to move on. So I doubt I'll reread them again until I'm in my dotage, fully resigned from software development and R&D in areas of human health and wellness. And I think that from now on I'll hold off until Ryk finishes each subsequent 15-book arc, since it's too frustrating to be waiting month after month for the next book to drop. But, for what it's worth, I have also received a LOT of positive feedback about The Frontiers Saga, so it's not just me.

In any event, once I finish my current reread, the Bobiverse trilogy, beginning with "We Are Legion" will be up next.

SpinRite

Since I got a ton of feedback after sharing my story of the recovery of that inaccessible Bitlocker encrypted drive, I thought I'd share an engaging recent anecdote that I posted to the SpinRite.dev newsgroup last Saturday:

Gang,

I hadn't checked-in for a while, so I thought that after a lost day of work, yesterday, I'd do so before settling back down to it.

The day before yesterday (the 29th), around noon, the eCommerce system I wrote back in 2003 began failing and reporting timeout errors when attempting to connect to our backend credit card processing provider. I would normally have been informed of this immediately, but the monitoring system I have been using for years never recovered after a power outage a few months back, and I hadn't wanted to take the time away from work to fix it.

So yesterday, Sue let me know that a few would-be customers had reported that they'd been unable to purchase SpinRite. She sent me a text message which captured my attention.

The short version is that I spent nearly the entire day pulling out what (very) little hair I have trying to figure out WTF was going on. The error reports that my own code was logging was a 0x2EE2 from the WinInet API, which is "operation timed out" and Windows' own error logging was complaining of "TLS handshake errors" — which could have been more informative.

That sent me off on what turned out to be a wild goose chase, assuming that my provider had changed their TLS connection parameters in a way that was incompatible with my aging Win2008 R2 server. The fact that Digicert (their cert provider, too) had just revised some of their intermediate certs and GRC's server certs were reporting an invalid intermediate, didn't help.

Many hours later, the final clue came when a ping to the service to the IP address I received from nslookup worked, whereas a ping to the same service with 'ping' doing the IP lookup did not. When I looked more closely I saw that nslookup and ping were resolving different IPs. I use my network's own UNIX BIND instance as my network's recursive resolver, so I became suspicious of it. But additional testing showed that it wasn't at fault.

So I finally thought FINE!, I'll just force the resolution to an IP that I know works by adding an entry to the local HOSTS file!

And there I found the OVERRIDE to that domain's OLD IP, already IN the local HOSTS file! <sigh> For some reason, sometime in the distant past, I had hard-wired the backend provider's IP. And then, two days ago, they finally changed their server's IP, no doubt doing so with great forethought, running over all IPs while giving DNS caches times to expire and refresh... and my old hard-wiring didn't allow my system to follow.

I removed the entry from the HOST's file and everything worked again perfectly.

I love computers because they always do exactly what we tell them to.

On a happy note, in the same vein of loving computers because they do exactly what we tell them to, I wanted to report that the use of my built-to-suit virtualized I/O function, and the new way I'm handling errors occurring in a massively long block of sectors, has turned out to be somewhat jarringly correct. As can happen when everything is designed properly, everything has just fallen into place by itself.

When I was interrupted yesterday I was working on synchronizing the logging system with the new inner loop since the information that can be logged has changed significantly.

It took me a while to understand why I had originally built the logging system the way I had, since it seemed WAY over designed and overly complex. It uses short log entry trigger tokens which are accumulated into a queue, then later flushed, expanded into their full size log entries and written. I couldn't figure out why I went to so much work, until I remembered the challenge that I had taken up and accepted: The SpinRite could log onto the same FAT partition that it was operating on... without any compromise. This meant that the log file itself might be written to the same track and sectors that SpinRite was in the middle of working on at the time.

But I already had full track virtualization. I was intercepting DOS's writes to the drive through the BIOS or device drivers, or compression drivers <shudder> and rerouting any reads and writes to the drive to a virtual buffer of the current track. So that wasn't why I was deferring the logging with a queue. It turned out that I was deferring the logging with a queue of short event tokens since I was re-using the track buffer for token expansion to reduce SpinRite's memory footprint to the absolute minimum.

Anyway, the system I built is pretty slick. It uses macros to implement its own meta language to make the implementation and result visually clean and clear. Although I don't need any of

that anymore, it's all in place and it works, so I've left it alone. I just needed to remember and understand how it worked so that I could confidently modify and extend its operation.

I'm returning to work on SpinRite, now with the mystery of the dead eCommerce system nicely resolved. Sheesh! :)

The BlackMatter Interview

The security firm Recorded Future introduced their exclusive interview of a representative of the group which now calls itself BlackMatter by noting:

In July, a new ransomware gang started posting advertisements on various cybercrime forums announcing that it was seeking to recruit partners and claiming that it combined the features of notorious groups like REvil and DarkSide.

Named BlackMatter, the gang said it was specifically interested in targeting large companies with annual revenues of more than \$100 million. However, the group said some industries were off limits: It would not extort healthcare, critical infrastructure, oil and gas, defense, non-profit, and government organizations.

A representative from the group talked to a Recorded Future expert threat intelligence analyst recently about how BlackMatter is learning from the mistakes of other ransomware groups, what they look for when they recruit partners, and why they avoid certain targets. The interview was conducted in Russian and translated to English with the help of a professional translator, and has been edited for clarity.

So, as we're listening to this conversation remember that there's no honor among thieves and that we already know with virtual certainty that BlackMatter is DarkSide, sharing virtually identical and unique codebases.

Recorded Future's Dmitry: Your product appeared quite recently and as far as we know, there have been no public attacks using BlackMatter yet. How long ago did you start developing it?

BlackMatter: There haven't been any attacks yet if you are judging by the public blog. In fact, there have been, and the companies we attacked are already communicating with us. As long as the negotiations are successful we do not publish a blog post on the main page of the blog.

The product has been in development for the last six months. Perhaps it seems simple (judging by the blog or the communication page), but it is not—what users see publicly is the tip of the iceberg.

Before starting the project, we studied the following products in detail:

- LockBit has a good codebase, but a skimpy and non-functional panel (at the time we used their product). If you compare it to a car, you can say that this is a Japanese car production line with good engines but an empty and non-functional interior. You can ride one, but with little pleasure.*
- REvil is a good project on the whole, time-tested software (since GandCrab, they haven't made any significant edits since that time), a fairly functional panel, but focused more on the overall number of successful "loads" as opposed to specific targeted cryptography.*
- Darkside is a relatively new software with a good codebase (partly problematic, but the ideas themselves deserve notice) and an interesting web part compared to other RaaS.*
- The executable itself has incorporated the ideas of LockBit, REvil, and partly DarkSide. The web part has incorporated the technical approach of DarkSide since we consider it the most structurally correct (separate companies for each target, and so on).*

DS: How difficult is it to organize an affiliate program (also known as ransomware-as-a-service)?

BM: On the whole, less difficult than not. The level is important, RaaS can also be offline (when builds are issued via jabber/tox), but there is no market demand for this and current customers, after using REvil and DarkSide, are not ready to take such affiliate programs seriously. We created a project and brought it to the market exactly at a time when the niche is vacant and the project fully meets the market demands, therefore its success is inevitable.

DS: Most recently, the largest groups—DarkSide, REvil, Avaddon, BABUK—have disappeared from the scene. Many researchers believe that this was due to the attention of the top leadership of the United States and Russia to the situation with ransomware attacks. Is it true? Do you think your product will have the same fate?

BM: Yes, we believe that to a large extent their exit from the market was associated with the geopolitical situation on the world stage. First of all, this is the fear of the United States and its planning of offensive cyber operations, as well as a bilateral working group on cyber extortion. We are monitoring the political situation, as well as receiving information from other sources. When designing our infrastructure, we took into account all these factors and we can say that we can withstand the offensive cyber capabilities of the United States. For how long? Time will tell. For now, we are focusing on long-term work. We also moderate the targets and will not allow our project to be used to encrypt critical infrastructure, which will attract unwanted attention to us.

DS: You mentioned that your product brings together the very best of DarkSide, REvil, and LockBit. What are their strengths?

BM: Our project has incorporated the strengths of each of the partner programs:

From REvil—SafeMode, their implementation was weak and not well thought out, we developed the idea and thoroughly implemented it. We also implemented the PowerShell version of the ransomware variant given the REvil implementation.

From LockBit—an approach to the implementation of the codebase, we took some things from there, mostly little things.

From DarkSide—first of all, this is the idea of impersonation (the ability of the encryptor to use the domain administrator account to encrypt the shared drives with maximum rights), we also borrowed the structure of the admin panel from there.

DS: Based on the latest reports published this week, BlackMatter is visually very similar to DarkSide. Can you confirm that your infrastructure is based on DarkSide?

BM: We can confidently say that we are fans of dark mode in design, we are familiar with the DarkSide team from working together in the past but we are not them, although we are intimate with their ideas.

DS: LockBit 2.0 is considered the fastest locker at the moment. What is the encryption/decryption speed of your variant?

BM: This is not true. After reading the question – we decided to prepare ourselves by downloading the latest publicly available version of LockBit (end 06.21) and conducting tests, we can state the following:

BlackMatter: time required 2.22

LockBit: time required 02.59

The tests were carried out under the same conditions. Moreover, LockBit encrypts the first 256 kb of the file (which is pretty bad from the point of view of cryptographic strength). We, on the other hand, encrypt 1 MB. Essentially, that's the secret to their speed.

DS: Are you planning to add new features to the product, following the example of StealBit?

BM: Yes, the software is constantly being improved, in terms of the new functions that will appear in the near future—printing the text of the note on all available printers. We also watch our competitors and always implement what we consider promising and in demand by our clients.

DS: I have already seen several recruiting announcements for your team. How many penetration testers would you like to recruit? Is it easier to work with a small but strong team, or with an army of script kiddies?

BM: We are geared at strong, self-sufficient teams with experience, their own technical solutions, and a real desire to make money, not someone who wants to try the business out. We usually filter out script kiddies before they get access to our admin panel.

DS: Obviously, there are many talented professionals on your team. Why is it that this talent is aimed at destructive activities? Have you tried legal penetration testing?

BM: We do not deny that business is destructive, but if we look deeper—as a result of these problems new technologies are developed and created. If everything was good everywhere there would be no room for new development.

There is one life and we take everything from it, our business does not harm individuals and is aimed only at companies, and the company always has the ability to pay funds and restore all its data.

We have not been involved in legal pentesting and we believe that this could not bring the proper material reward.

DS: What do you think about the attacks carried out against Colonial Pipeline's infrastructure or JBS? Does it make sense to attack such large networks?

BM: We think that this was a key factor for the closure of REvil and DarkSide, we have forbidden that type of targeting and we see no sense in attacking them.

DS: The US Department of Justice said they were able to recover some of the bitcoins paid by Colonial. How do you think this has happened?

BM: We think that the DarkSide team or their partners transferred bitcoins to web wallets, which led to the seizure of private keys.

DS: You are actively buying access to the networks and declare that you are NOT interested in government and medical institutions. At the same time, you stated that you will not encrypt a wider range of industries, including critical infrastructure, defense, non-profit, and oil. Who has the last word to encrypt the network or not?

BM: The last word is ours. We check each target and decide if it has potential negative consequences for us. The discrepancy between the industries in the blog and on the forum is related to marketing. In personal correspondence we filter out those which we are not interested in.

DS: What type of primary network access is the easiest in 2021 in your opinion?

BM: We do not work with VPN and other time-consuming types of initial access but are focused on getting direct access to the network immediately.

DS: What carries more effect motivating the company to pay: The infrastructure being unavailable, or the fear of a data leak?

BM: It varies from company to company. For some it is important to maintain confidentiality, and for others it's restoring infrastructure. If the network is completely encrypted and there is also a risk of data being published, the company will most likely pay.

DS: "Unknown" [REvil's public spokesperson] spoke about a special outlook towards insurance companies. Do you think that if insurance companies abruptly stop covering ransomware incidents it will change your interest in ransomware?

BM: It will not change, the companies will continue to pay money regardless. It is possible that the amount being paid will decrease.

Now the insurance fees have increased, but fearing that they will be left alone in the situation everyone will continue buying the insurance.

DS: What's happened with Unknown? There are a lot of rumors, can you clarify the situation?

BM: We do not know. Most likely, after the last payment, he went on vacation or is preparing a rebranding of their project.

DS: Tell me a secret.

BM: There are no secrets, but we believe in our motherland, we love our families, and we earn money for our children.

