Security Now! #730 - 09-03-19

The Ransomware Epidemic

This week on Security Now!

Rather than looking at many small bits of news, this week we take longer looks a few larger topics. We'll examine several pieces of welcome news from the bug bounty front. We also take a look at Google's Project Zero revelation of a comprehensive multi-year campaign aimed at iOS visitors to specific websites. And then we conclude with a distressingly large array of news from the ransomware front. We figure out how to pronounce Sodinokibi (so-dee'-no-kee-bee) and ponder the future of ransomware.

2019 Ransomware Attack Victims

An Epidemic with no Patient Zero

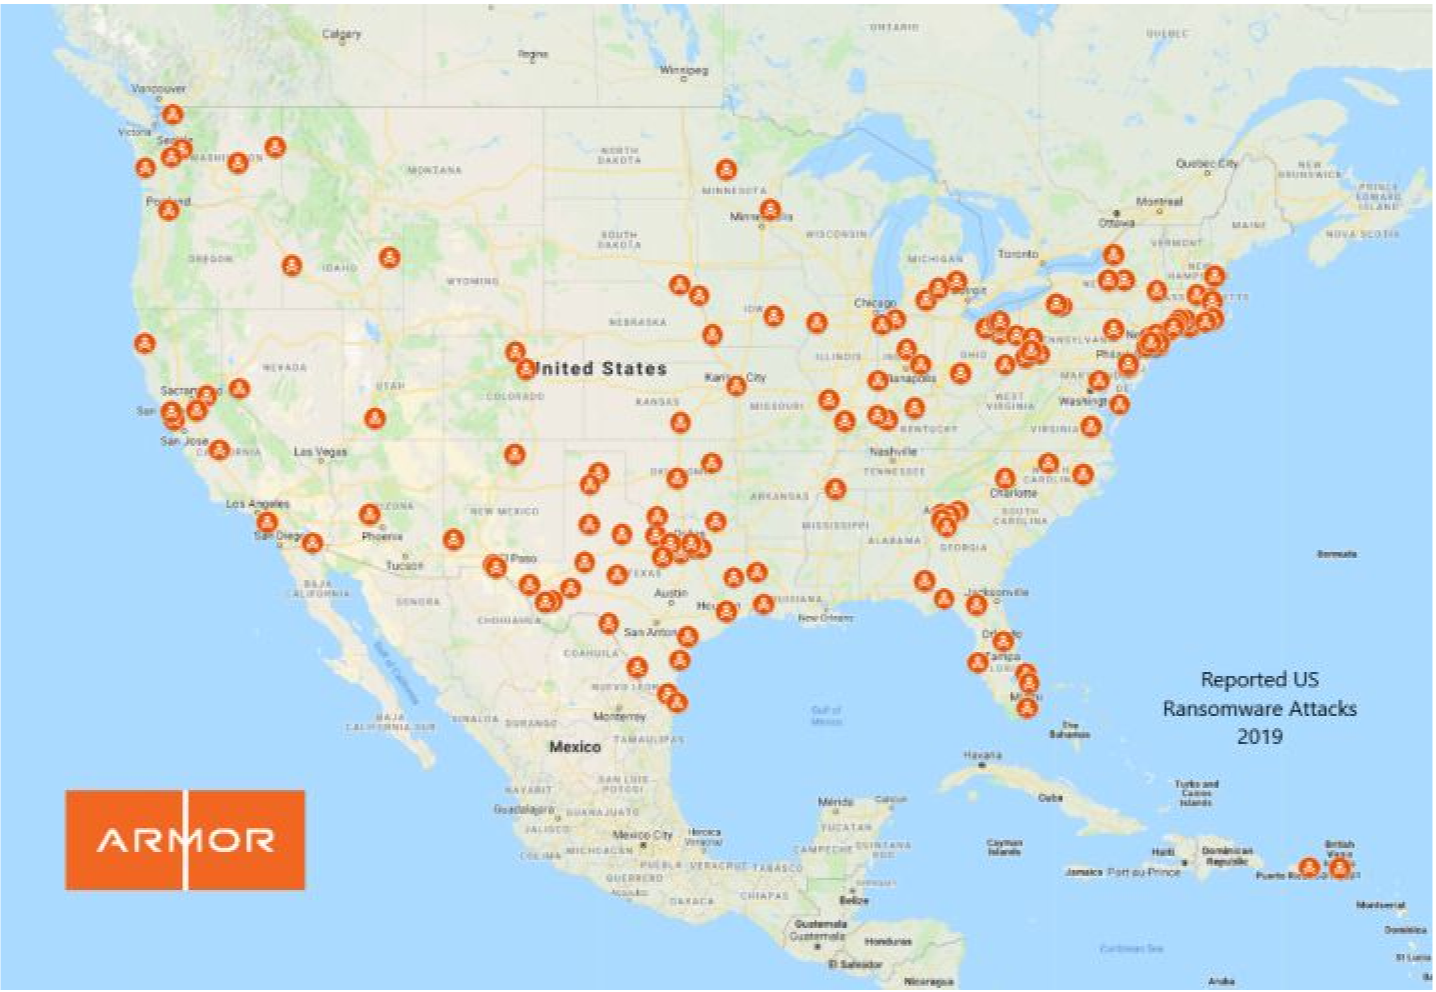

Security News

Android Developer's Blog last Thursday was titled: "Expanding bug bounties on Google Play"
https://android-developers.googleblog.com/2019/08/expanding-bug-bounties-on-google-play.html

[After some introductory stuff they announce...]

> Google Play Security Reward Program Scope Increases
>
> We are increasing the scope of GPSRP to include all apps in Google Play with 100 million or more installs. These apps are now eligible for rewards, even if the app developers don't have their own vulnerability disclosure or bug bounty program. In these scenarios, Google helps responsibly disclose identified vulnerabilities to the affected app developer. This opens the door for security researchers to help hundreds of organizations identify and fix vulnerabilities in their apps. If the developers already have their own programs, researchers can collect rewards directly from them on top of the rewards from Google. We encourage app developers to start their own vulnerability disclosure or bug bounty program to work directly with the security researcher community.
>
> Vulnerability data from GPSRP helps Google create automated checks that scan all apps available in Google Play for similar vulnerabilities. Affected app developers are notified through the Play Console as part of the App Security Improvement (ASI) program, which provides information on the vulnerability and how to fix it. Over its lifetime, ASI has helped more than 300,000 developers fix more than 1,000,000 apps on Google Play. In 2018 alone, the program helped over 30,000 developers fix over 75,000 apps. The downstream effect means that those 75,000 vulnerable apps are not distributed to users until the issue is fixed.
>
> To date, GPSRP has paid out over $265,000 in bounties. Recent scope and reward increases have resulted in $75,500 in rewards across July & August alone. With these changes, we anticipate even further engagement from the security research community to bolster the success of the program.
>
> Introducing the Developer Data Protection Reward Program
>
> Today, we are also launching the Developer Data Protection Reward Program. DDPRP is a bounty program, in collaboration with HackerOne, meant to identify and mitigate data abuse issues in Android apps, OAuth projects, and Chrome extensions. It recognizes the contributions of individuals who help report apps that are violating Google Play, Google API, or Google Chrome Web Store Extensions program policies.
>
> The program aims to reward anyone who can provide verifiably and unambiguous evidence of data abuse, in a similar model as Google's other vulnerability reward programs. In particular, the program aims to identify situations where user data is being used or sold unexpectedly, or repurposed in an illegitimate way without user consent. If data abuse is identified related to an app or Chrome extension, that app or extension will accordingly be removed from Google Play or Google Chrome Web Store. In the case of an app developer abusing access to Gmail

> restricted scopes, their API access will be removed. While no reward table or maximum reward is listed at this time, depending on impact, a single report could net as large as a $50,000 bounty.
>
> As 2019 continues, we look forward to seeing what researchers find next. Thank you to the entire community for contributing to keeping our platforms and ecosystems safe. Happy bug hunting!

So Google has expanded the scope of their bug bounty program, not only by adding many more apps to qualify for bounties, and by encouraging app developers, who can, to create their own bounty programs... and also by expanding the range of what qualifies to include not only bugs but also the abuse of information.

And speaking of the bug bounty industry...
As we discussed in March, a 19-year-old white hat hacker, Santiago Lopez, known as "@try_to_hack" was the first to surpass the target of $1 million in earnings on HackerOne. Santiago earned himself more than $1 million in earnings by identifying vulnerabilities in the software or systems belonging to Twitter, HackerOne, Automattic, Verizon, various private companies, the U.S. government and others.

Now, HackerOne has announced that five additional hackers have joined Santiago to become Bug Bounty Millionaires by finding and reporting security vulnerabilities through HackerOne's vulnerability coordination and bug bounty platform.

- Mark Litchfield (@mlitchfield) from the U.K.
- Nathaniel Wakelam (@nnwakelam) from Australia
- FransRosen (@fransrosen) from Sweden
- Ron Chan (@ngalog) from Hong Kong, and
- Tommy DeVoss (@dawgyg) from the U.S.

… all now members of the $1M hacking bounty club.

I should also note my thanks to HackerOne for their very kind shoutout about me in their most recent blog posting last Friday, which they titled: “HackerOne Praised By An Original Hacker”
https://www.hackerone.com/blog/hackerone-praised-original-hacker

They were referring to our podcast #720 ten weeks ago which I titled “Bug Bounty Business.” As our listeners know well, I consider bug hunting for profit to be a 100% legitimate and very intriguing career. I have the feeling that the six HackerOne millionaires would likely agree. :)

Meanwhile, also last Thursday, Google's Project Zero group dropped a bomb on iOS...
by revealing that for a period of at least several years, a small group of websites was successfully infecting anyone visiting them with RAM-based monitoring malware, through the use of a quite sophisticated, iOS-version-based multi-stage exploit chain.

As far as everyone knows, this all finally ended with the emergency jump to iOS version 12.1.4. But as Project Zero's Ian Beer explained, the story before then is not only troubling due to the exposure created for those who were infected, but more so due to the conclusions Ian draws based upon the exact nature of what he found.

Ian’s Project Zero blog posting was titled: “A very deep dive into iOS Exploit chains found in the wild” Ian wrote:

> Project Zero’s mission is to make 0-day hard. We often work with other companies to find and report security vulnerabilities, with the ultimate goal of advocating for structural security improvements in popular systems to help protect people everywhere.
>
> Earlier this year Google's Threat Analysis Group (TAG) discovered a small collection of hacked websites. The hacked sites were being used in indiscriminate watering hole attacks against their visitors, using iPhone 0-day.
>
> There was no target discrimination; simply visiting the hacked site was enough for the exploit server to attack your device, and if it was successful, install a monitoring implant. We estimate that these sites receive thousands of visitors per week.
>
> TAG was able to collect five separate, complete and unique iPhone exploit chains, covering almost every version from iOS 10 through to the latest version of iOS 12. This indicated a group making a sustained effort to hack the users of iPhones in certain communities over a period of at least two years.

And here comes the very troubling part:

> I’ll investigate what I assess to be the root causes of the vulnerabilities and discuss some insights we can gain into Apple's software development lifecycle. The root causes I highlight here are not novel and are often overlooked: we'll see cases of code which seems to have never worked, code that likely skipped QA or likely had little testing or review before being shipped to users.

Now, we know that iOS and Android are inherently competitive, that Ian is on Google’s team, and that what he wrote there sounds harsh. But when you dig deeper into this work -- and it’s truly stunning work on Ian’s part -- you see that he’s been working for years to make Apple’s iOS better and more secure for everyone, and that comparatively few people are ever going to be able to appreciate what he has done. So, despite the way that sounds, he’s really not attacking Apple. He’s looking at the code that he is seeing being exploited and drawing the only conclusions that can be drawn from what’s in front of him.

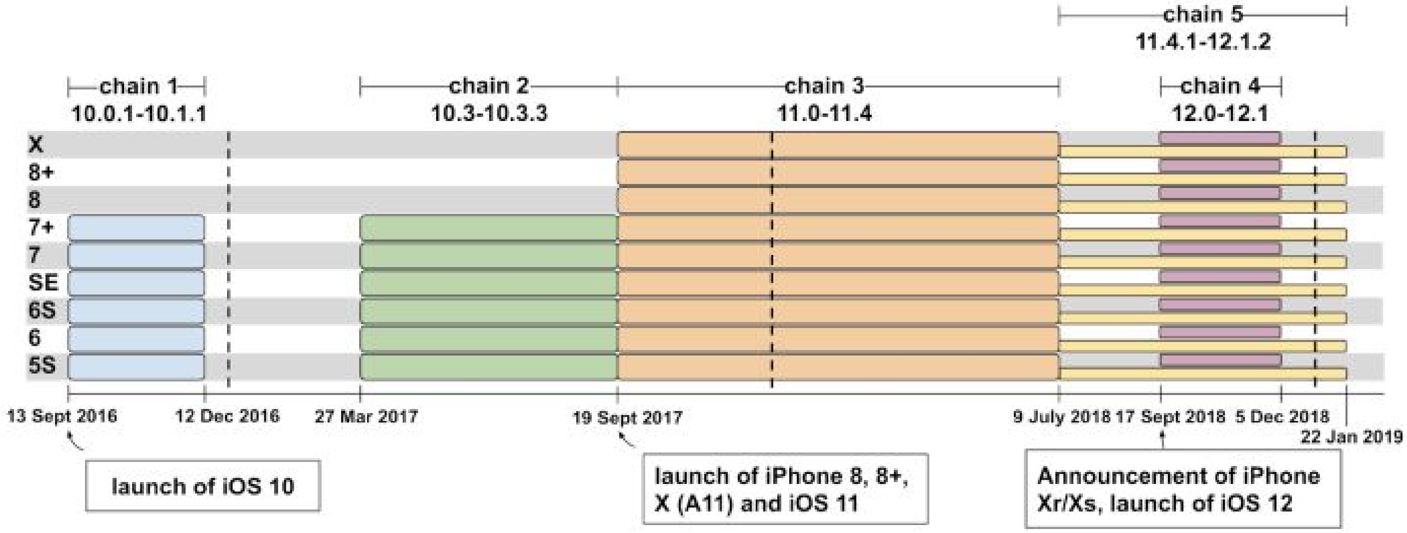

launch of iOS 10

launch of iPhone 8, 8+, X (A11) and iOS 11

Announcement of iPhone Xr/Xs, launch of iOS 12

Working with TAG (Google's Threat Analysis Group), we discovered exploits for a total of fourteen vulnerabilities across the five exploit chains: seven for the iPhone's web browser, five for the kernel and two separate sandbox escapes. Initial analysis indicated that at least one of the privilege escalation chains was still 0-day and unpatched at the time of discovery (CVE-2019-7287 & CVE-2019-7286). We reported these issues to Apple with a 7-day deadline on 1 Feb 2019, which resulted in the out-of-band release of iOS 12.1.4 on 7 Feb 2019. We also shared the complete details with Apple, which were disclosed publicly on 7 Feb 2019.

Now, after several months of careful analysis of almost every byte of every one of the exploit chains, I'm ready to share these insights into the real-world workings of a campaign exploiting iPhones en masse.

This post will include:

- detailed write-ups of all five privilege escalation exploit chains;
- a teardown of the implant used, including a demo of the implant running on my own devices, talking to a reverse-engineered command and control server and demonstrating the capabilities of the implant to steal private data like iMessages, photos and GPS location in real-time, and
- analysis by fellow team member Samuel Groß on the browser exploits used as initial entry points.

Let's also keep in mind that this was a failure case for the attacker: for this one campaign that we've seen, there are almost certainly others that are yet to be seen.

[In other words, they first discovered a website that was doing this, then followed that trail back into iOS. But they very reasonably assume that there were likely other websites that were also exploiting the same suite of vulnerabilities.]

Real users make risk decisions based on the public perception of the security of these devices. The reality remains that security protections will never eliminate the risk of attack if you're being targeted. To be targeted might mean simply being born in a certain geographic region or

> being part of a certain ethnic group. All that users can do is be conscious of the fact that mass exploitation still exists and behave accordingly; treating their mobile devices as both integral to their modern lives, yet also as devices which when compromised, can upload their every action into a database to potentially be used against them.
>
> I hope to guide the general discussion around exploitation away from a focus on the million dollar dissident and towards discussion of the marginal cost for monitoring the n+1'th potential future dissident. I shan't get into a discussion of whether these exploits cost $1 million, $2 million, or $20 million. I will instead suggest that all of those price tags seem low for the capability to target and monitor the private activities of entire populations in real time.
>
> I recommend that these posts are read in the following order: ...

Ian then has seven postings, one for each of the five individual exploit chains, each of which he dissects in exquisite mind-blowing detail, which I cannot possibly summarize here except to note that it was amazing work.

In his sixth posting he examine the Webkit exploits used to attain an initial foothold in iOS (and he notes for the record that Chrome on iOS would have also been vulnerable even though Safari was attacked.)

And in his final 7th posting he examines the operation of the implant that is finally dropped into the device's RAM…

> In the earlier posts we examined how the attackers gained unsandboxed code execution as root on iPhones. At the end of each chain we saw the attackers calling `posix_spawn`, passing the path to their implant binary which they dropped in `/tmp`. This starts the implant running in the background as root. There is no visual indicator on the device that the implant is running. There's no way for a user on iOS to view a process listing, so the implant binary makes no attempt to hide its execution from the system.
>
> The implant is primarily focused on stealing files and uploading live location data. The implant requests commands from a command and control server every 60 seconds.
>
> Before diving into the code let's take a look at some sample data from a test phone running the implant and communicating with a custom command and control server I developed. To be clear, I created this test specifically for the purposes of demonstrating what the implant enabled the attacker to do and the screenshots are from my device. The device here is an iPhone 8 running iOS 12.
>
> The implant has access to all the database files (on the victim's phone) used by popular end-to-end encryption apps like Whatsapp, Telegram and iMessage. We can see here screenshots of the apps on the left, and on the right the contents of the database files stolen by the implant which contain the unencrypted, plain-text of the messages sent and received using the apps:

He then shows access to WhatsApp's, Telegram's and iMessage's data. Writing about Hangouts, he notes:

> ## Hangouts
> Here's a conversation in Google Hangouts for iOS and the corresponding database file uploaded by the implant. With some basic SQL we can easily see the plain text of the messages, and even the URL of the images shared.
>
> The implant can upload private files used by all apps on the device; here's an example of the plaintext contents of emails sent via Gmail, which are uploaded to the attacker's server:
>
> ## Contacts
> The implant also takes copies of the user's complete contacts database:
>
> ## Photos
> And takes copies of all their photos:
>
> ## Real-time GPS tracking
> The implant can also upload the user's location in real time, up to once per minute, if the device is online. Here's a real sample of live location data collected by the implant when I took a trip to Amsterdam with the implant running on a phone in my pocket:
>
> The implant uploads the device's keychain, which contains a huge number of credentials and certificates used on and by the device. For example, the SSIDs and passwords for all saved wifi access points:
>
> The keychain also contains the long-lived tokens used by services such as Google's iOS Single-Sign-On to enable Google apps to access the user's account. These will be uploaded to the attackers and can then be used to maintain access to the user's Google account, even once the implant is no longer running. Here's an example using the Google OAuth token stored as `com.google.sso.optional.1.accessToken` in the keychain being used to log in to the Gmail web interface on a separate machine:

After showing those examples Ian gets back down into the weeds to looks at the details decompilation of the implant and its interaction with the system. He finally concludes this posting with a work about its impact:

> ## Impact
>
> The implant has access to almost all of the personal information available on the device, which it is able to upload, unencrypted, to the attacker's server. The implant binary does not persist on the device; if the phone is rebooted then the implant will not run until the device is re-exploited when the user visits a compromised

> site again. Given the breadth of information stolen, the attackers may nevertheless be able to maintain persistent access to various accounts and services by using the stolen authentication tokens from the keychain, even after they lose access to the device.

The Ransomware Epidemic

Still nothing more definitive from Texas, though now nine of the reported 22 affected state municipalities have been identified. I'm still struck by the surprisingly different "feel" that these attacks have. And I continue to think that an attack on a common services provider is the most likely explanation for everything we're seeing... though evidence to support that opinion is decidedly scant.

Meanwhile, we have a bunch of (actually 13!) new ransomware attack victims. While most of them are school districts, we also have a county in Indiana, a hospice in California, and a newspaper in Watertown, New York.

Armor, the cloud security firm whose data generated our picture of the week, has tracked the following new ransomware infections:

- Lake County, Indiana
- Rockville Center School District in Rockville Center, New York
- Moses Lake School District in Moses Lake, Washington
  (the attack occurred back in July but was only reported to be ransomware this month)
- Mineola Public Schools in Mineola, New York
- The Stevens Institute of Technology in Hoboken, New Jersey
- New Kent County Public Schools in New Kent, Virginia
- Nampa Idaho School District, Nampa, Idaho
- Middletown School District, Middletown, Connecticut
- Wolcott Public Schools, Wolcott, Connecticut
- Wallingford School District, Wallingford, Connecticut
- New Haven Public Schools, New Haven, Connecticut
- The Watertown Daily Times in Watertown, New York
- Hospice of San Joaquin, San Joaquin, California

Although we still don't know what's going on in Texas, our old friend Ryuk (ree-ook) has been identified as the culprit in at least three of these additional recent attacks.

Newsday reported that the Rockville Center School District in New York initially received a ransom demand of $176,000. The district's insurance company negotiated with the ransomware operator, reducing the payout to $88,000. The school district paid a deductible of $10,000.

There's no word on whether other victims have paid any ransoms yet. Back in June, Brian Krebs did some reporting on the Baltimore, Maryland attack which took Baltimore's servers down on May 7th. After communicating with Armor's Joe Stewart, Brian reported that Baltimore was the victim of a ransomware strain known as “RobbinHood.” Now, here we are four months later and

Baltimore's leadership recently revealed that $6 million of the money needed to cover the city's more than $10 million ransomware cleanup operation would be pulled from funds earmarked for upkeep of city parks and public facilities. So far, the RobbinHood ransomware cost the city over $8 million in lost revenue and interest on deferred revenue. And as a result of all of this, Baltimore is now considering a contract for a $20 million "cyber liability" insurance plan. A vote on the proposed $835,000 contract was deferred until next week.
So these 13 new attacks brings the total known, publicly reported ransomware attacks this year to a total of 149, and of those, 30 have involved educational institutions. Chris Hinkley, the head of Armor's Threat Resistance Unit security team told ArsTechnica: "Just like municipalities, which rely on critical systems to manage records and revenue in a community, school districts host data and systems critical to their community and its students. Thus, hackers know that schools cannot afford to shut down, that budgets are typically stretched thin, and that they often have few security protections in place, all aspects which make them a viable target." And I'll add that we now know that insurance carriers indirectly provide many of these institutions with virtual deep pockets.

Last May Ars observed that school districts are a particularly easy target for ransomware operators because of their generally low budget for IT and their limited security resources. According to data collected by Armor, schools have become the second-largest pool of ransomware victims—slightly behind local governments and closely followed by healthcare organizations.

But we should note that these numbers are only for known and publicly reported incidents.

CNN also reported last week that Percsoft and Digital Dental Record, two companies that handle online services in the dental industry, told roughly 400 customers that the software they use to connect to individual offices had been infected with ransomware earlier in the week. Dental administrators told CNN they couldn't access basic information such as patient charts, X-ray data, or payment services. Digital Dental provided a statement saying 100 of the affected practices had been restored this week.

These more private incidents highlight the fact that many ransomware attacks may go unreported, especially those against small and mid-sized companies who largely turn to insurers to help them pay off attackers quietly. And this distinction of privacy might help explain why a report by the US Federal Bureau of Investigation (FBI) reported 1,493 ransomware cases last year, which was many more than are publicly known.

Other than New York's Rockville school district which was insured and negotiated that $88,000 ransom payment, it's still too soon to know whether any of the other newly attacked public organizations have cyber insurance in place or plans to pay ransoms. The payouts which ransomware operators have recently received from similar targets, such as Riviera City, Florida's $600,000 ransom and Lake City, Florida's $500,000 ransom have, as we have noted here, very clearly signaled to the hackers that attacking entire communities can be very lucrative.

A recent ProPublica investigation revealed that insurance companies are fueling the rise of ransomware threats by covering the cost minus a deductible — which is usually far less than the ransom demanded by attackers. Hackers targeting entities that they know are likely to have cyber insurance has led to the advent of incident response firms that provide **"cyber extortion**

negotiation services” and help companies recover data post infection. ProPublica’s report noted:

“By rewarding hackers, it encourages more ransomware attacks, which in turn frighten more businesses and government agencies into buying policies.”

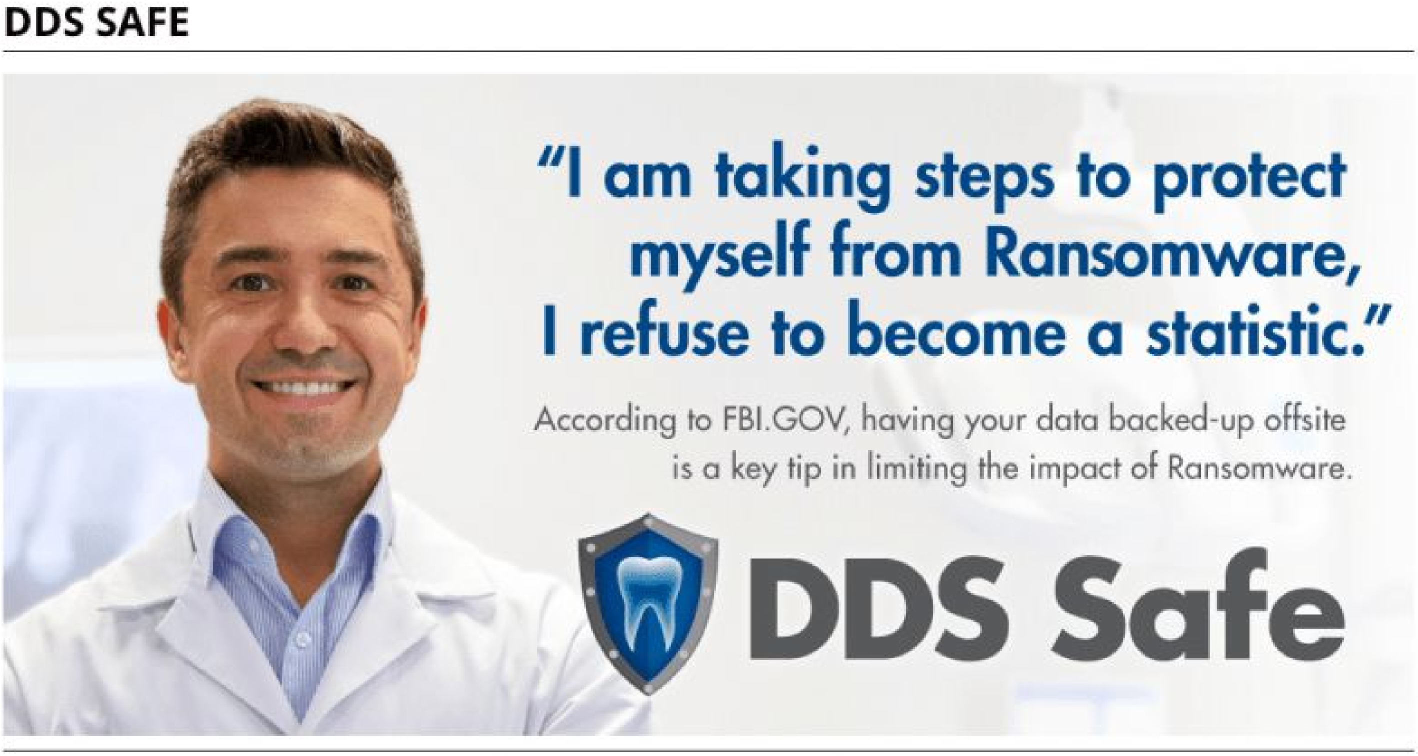
DDS SAFE
“I am taking steps to protect myself from Ransomware, I refuse to become a statistic.”
According to FBI.GOV, having your data backed-up offsite is a key tip in limiting the impact of Ransomware.
DDS Safe

(Uhhhh ... stand and be counted.)

In a bit of irony, the Digital Dental Record company advertises its "DDS Safe" service on its website as a way to safeguard files from ransomware attacks.

In his reporting on this, Brian Krebs was unable, as we know I also would have been, to resist titling this coverage: *“Ransomware **Bites** Dental Data Backup Firm”*

Brian explained that PerCSoft, based in West Allis Wisconsin, is a company that manages a remote data backup service relied upon by hundreds of dental offices across the country through its provision of cloud services to Digital Dental Record (DDR), which offers an online dental data backup service called DDS Safe that archives medical records, charts, insurance documents and other personal information for dental offices across the United States.

The ransomware attack hit PerCSoft on the morning of Monday, Aug. 26, and encrypted dental records for some — but not all — of the practices that rely on DDS Safe. Brenna Sadler, the director of communications for the Wisconsin Dental Association, said the ransomware encrypted files for approximate 400 dental practices, and that somewhere between 80-100 of those clients have now had their files restored. Sadler said she did not know whether PerCSoft and/or DDR had paid the ransom demand, what ransomware strain was involved, or how much the attackers had demanded.

But updates to PerCSoft's Facebook page and statements published by both PerCSoft and Dental Data Record suggest someone may have paid up: The statements note that both companies worked with a third party software company and were able to obtain a decryptor to help clients regain access to files that were locked by the ransomware.

In Brian's reporting, several sources were reporting that PerCSoft did pay the ransom, although it is not clear how much was paid (Bleeping Computer has some updated data on that which we'll get to in a moment). One member of a private Facebook group dedicated to IT professionals serving the dental industry shared a screenshot, which is purportedly from a conversation between PerCSoft and an affected dental office, indicating the cloud provider was planning to pay the ransom.

What's still very unclear due to conflicting accounts using annoyingly fuzzy reporting -- probably because they either don't know or don't know that there's a difference -- is whether the infection of PerCSoft's cloud infrastructure allowed the infection to spread down the connection into the individual dental offices which relied upon DDR and PerCSoft. If, as it seems, the cloud was only being used for backup, and also as it appears more than the backups were affected and the operating records of these 400-some dental offices were impacted, then it does sound as though the malware was able to climb down into the machines of individual dental offices.

Brian reported that some affected dental offices have reported that the decryptor did not work to unlock at least some of the files encrypted by the ransomware and several affected dentistry practices said they feared they might be unable to process payroll payments this week as a result of the attack.

I obtained a PDF of the media statement put out by DentalRecord.com
https://www.dentalrecord.com/assets/images/MediaStatement2.pdf

In their initial disclosure they state:

> "At 8:44 a.m. on Monday, Aug. 26, we learned that ransomware had been deployed on the remote management software our product uses to back up client data. Immediate action was taken to investigate and contain the threat. Our investigation and remediation efforts continue. Unfortunately, a number of practices have been and continue to be impacted by this attack."

The key phrase may be "we learned that ransomware had been deployed on the remote management software our product uses to back up client data." As we know, the use of remote management software suggests that PercSoft may have remote management access to the networks and systems of their client dental offices... And that may be the way this relatively new "REvil" ransomware.

Not surprisingly, cloud data and backup services are a prime target of cybercriminals who deploy ransomware since ransomware encrypts data and a lot of data is stored in the cloud. Last month, attackers hit the QuickBooks cloud hosting firm iNSYNQ and held the data of many of the company's clients hostage. Earlier this year, last February, the cloud payroll data provider Apex Human Capital Management was taken down for three days following a ransomware infestation. And on Christmas Eve at the end of last year, the cloud hosting provider Dataresolution.net took its systems offline in response to a ransomware outbreak on its internal networks. The company was adamant that it would not pay the ransom demand, but it ended up taking several weeks for customers to fully regain access to their data.

As we know, the FBI and many security firms have advised victims not to pay any ransom demands, arguing that doing so encourages the attackers and maay not result in regaining access to encrypted files. But in reality, many cybersecurity consulting firms quietly note to their customers that paying up is the fastest route back to business-as-usual.

In a report published by the cybersecurity firm Fidelis last week, REvil emerged as the fourth most prevalent strain of ransomware, at 12.5%. Ryuk (ree-ook) holds the lead at 23.9%, followed by Phobos with 17% and Dharma with 13.6%. Sheesh.

Sodinokibi (so-dee’-no-kee-bee) aka REvil

Bleeping Computer, who were very early to shine a bright spotlight on the ransomware scene have a bunch of great information about this newest kid on the block whose name I have **no idea** how to pronounce.

Bleeping Computer notes that although it is relatively new on the ransomware scene, Sodinokibi has already made impressive profits for its administrators and affiliates, some victims paying as much as $240,000 and network infections netting an average of $150,000.

Our listeners may have noted that I used the term “affiliate” -- which should be a clue to this new terror’s distribution and exploitation model: Yep… it’s RaaS -- Ransomware as a Service!

As we covered back in April, the guys behind GandCrab, claiming to have made all the money they needed, and having laundered it and reinvested it in legitimate businesses, decided to shut down their operation. We wondered and speculated at the time what might surface to fill the void left behind by this diabolical network marketing model.

We need wonder no longer. It’s the thing that knocked those 22 towns off the map in Texas and then nuked those 400-some dental practices.

Yup… it’s truly REvil.

Bleeping Computer reports that the Texas extortion totaled $2.5 million and that the ransom to restore the 400-odd dental offices machines was $5,000 per computer… So somewhere around another $2 million.

Since its first appearance last April, Sodinokibi (a.k.a. REvil) has become prolific and quickly gained a reputation among cybercriminals in the ransomware business and security researchers. And, as we know, this is aided by the fact that it’s making OTHERS among the underworld quite a bit of money in turn. Unlike Ryuk (ree-ook), this ransomware is available under an affiliate model.
After its first month, back in mid-May, to put itself more squarely on the map, a Sodinokibi advertiser using the forum name UNKN deposited over $100,000 on underground forums to show that they meant business.

Then online underground advertisements for the new file-encrypting malware appeared in early July on at least two forums. UNKN said that they were looking to expand their activity and that it was a private operation with "limited number of seats" available for experienced individuals.

UNKN describes the malware as “private ransomware” flexible enough to adapt to the RaaS business model.

UNKN offered their affiliates 60% of the payments at the beginning and a 10% increase to 70% after the first three transactions. The actor also made it clear that they would not be working with English-speaking affiliates as part of this private program.

Bleeping Computer posted a screenshot monitoring the ransomware’s bitcoin transactions showing the money pouring in from its victims. Look at the second one showing that one victim paid 26.388 bitcoins, which converted to more than $240,143 USD at the time of the transaction:

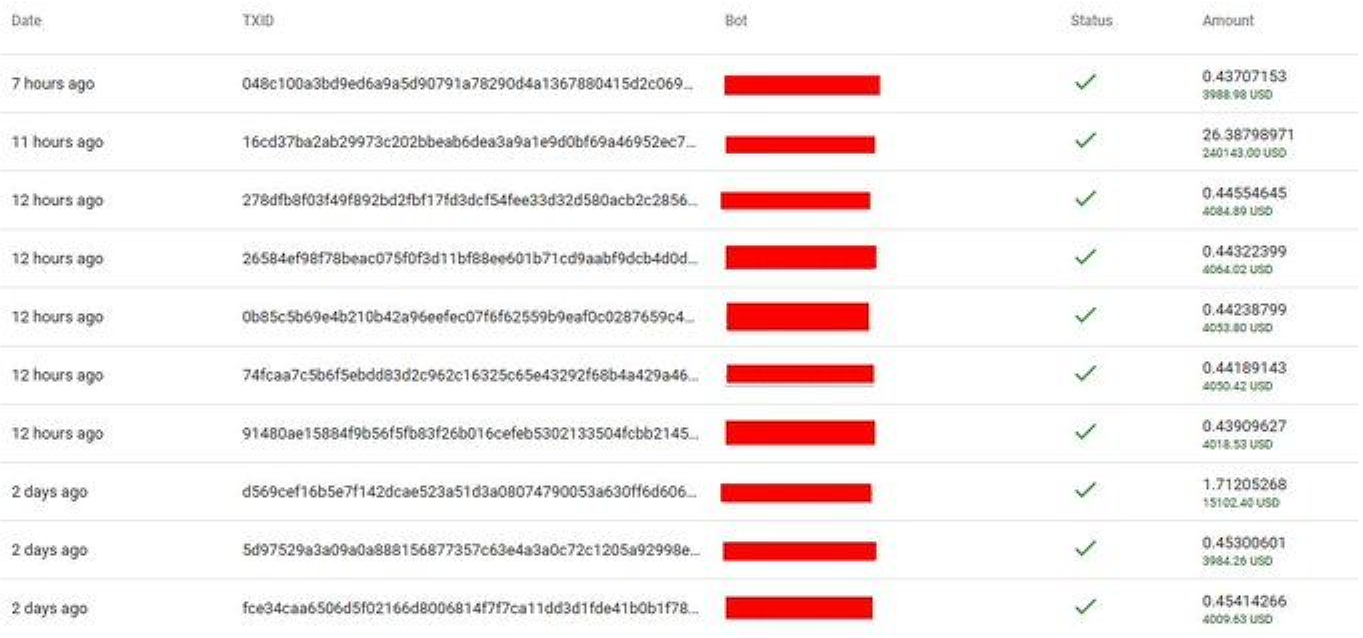

Date	TXID	Bot	Status	Amount
7 hours ago	048c100a3bd9ed6a9a5d90791a78290d4a1367880415d2c069...	[redacted]	✓	0.43707153 3988.98 USD
11 hours ago	16cd37ba2ab29973c202bbeab6dea3a9a1e9d0bf69a46952ec7...	[redacted]	✓	26.38798971 240143.00 USD
12 hours ago	278dfb8f03f49f892bd2fbf17fd3dcf54fee33d32d580acb2c2856...	[redacted]	✓	0.44554645 4084.89 USD
12 hours ago	26584ef98f78beac075f0f3d11bf88ee601b71cd9aabf9dcb4d0d...	[redacted]	✓	0.44322399 4064.02 USD
12 hours ago	0b85c5b69e4b210b42a96eefec07f6f62559b9eaf0c0287659c4...	[redacted]	✓	0.44238799 4053.80 USD
12 hours ago	74fcaa7c5b6f5ebdd83d2c962c16325c65e43292f68b4a429a46...	[redacted]	✓	0.44189143 4050.42 USD
12 hours ago	91480ae15884f9b56f5fb83f26b016cefeb5302133504fcbb2145...	[redacted]	✓	0.43909627 4018.53 USD
2 days ago	d569cef16b5e7f142dcae523a51d3a08074790053a630ff6d606...	[redacted]	✓	1.71205268 15102.40 USD
2 days ago	5d97529a3a09a0a888156877357c63e4a3a0c72c1205a92998e...	[redacted]	✓	0.45300601 3984.26 USD
2 days ago	fce34caa6506d5f02166d8006814f7f7ca11dd3d1fde41b0b1f78...	[redacted]	✓	0.45414266 4009.63 USD

For affiliates who arrange to infect an entire network, the REvil/Sodinokibi developers allow a victim to purchase a decryption tool for the entire fleet of infected computers. According to forum post shared with BleepingComputer, these network-wide decryptors have an average cost of $150,000.

With revenue flooding in, other malware distributors are attempting to gain access to the program, but UNKN stated that there are no available openings for affiliates at this time.

Bleeping Computer also reported that the operators behind the Sodinokibi Ransomware started searching for affiliates to distribute their software soon after GandCrab shut down. Underground reactions towards the new product suggest that there may be a connection with the administrators or the affiliates of the now defunct GandCrab operation.

Some malware analysts pointed to code-level similarities between the two ransomware strains, though many differences also exist between the two. One similarity is that administrators of both malware families refuse to carry out business in the Commonwealth of Independent States (CIS) area. This includes Russia, Ukraine, Moldova, Belarus, Armenia, Turkmenistan, Uzbekistan and several other "-stans" (Kyrgyzstan, Kazakhstan, Tajikistan.)
These breadcrumbs, along with the rapid ascension of this new malware, suggest the involvement of the previously well-established GandCrab crew or its affiliates. Already having connections on private forums may have allowed them to quickly promote Sodinokibi and be selective about their partners.

There is no clear, undeniable evidence that Sodinokibi is run by the same individuals that administered GandCrab, but they obviously know the ransomware game and are into the money-making business.

-#-

So, I'll conclude here by observing again that everything changed the moment we moved from viruses and malware for it own sake to viruses, botnets, malware and ransomware for profit. As soon as cryptocurrency exchanges existed, so that cyber currency could be used to pay the rent and buy a burger, stealing CPU cycles through cryptojacking jumped to the forefront.

But the emergence of cryptocurrencies also meant that there was now a means for ransomware perpetrators to safely receive extortion payments. Gone are the days where Western Union was used to "wire" monies that could be easily traced. So the rise of cryptocurrency was an enabling factor for the ransomware that soon followed.

Today, there isn't a malware author alive who isn't aware that it's now possible to live well by finding, infecting, and encrypting the data of the right targets. And really changes everything. The twisted brilliance of ransomware was that the victim's data was still there. They can still see their files, now with a ".crypto" or other extension appended to the end. The files weren't deleted, their contents were simply moved just out of reach. This allowed the carrot of full data recovery to be dangled in front of the victim. And, as a result, more often than not, though no one wants to knuckle under to extortion, the sanity of self-interest would prevail and money would flow into the bad guys' cryptocurrency wallets... thus further encouraging them to find their next victim.

One of my favorite analogies for this podcast is how "porous" our computer security really is. I think that "porosity" for most cyber security deployed in the field is exactly the right term. It does exist. It's kind of there. It's mostly useful. It pretty much protects us. But not really very well when put under pressure. And pressure is what the emergence of cryptocurrency exchanges has created.

So what's going to happen? What's the future look like?

When the world learned that the NSA had installed data sucking taps at many of the Internet's major network exchange points, our collective reaction was to treat encryption as more than an expensive option. It became important. And so today, to a far greater degree than previously, the entire world's communications are end-to-end encrypted.

I expect that we will soon see that something similar happens with system backup imaging. IT departments have not prioritized the maintenance of good hot and cold backup because they haven't had to. For the most part, things have been running along just fine. But that's no longer true. It will no longer be reasonable for a ransomware event to take a municipality offline for months or to close a school district. That won't fly in the future.
Storage has become super cheap. And where security concerns are paramount, there's no need to move backups offsite into the cloud. Just get a large rack of local storage and use it. I think we're going to soon see explicitly ransomware-oriented backup protection, where nightly snapshot images are rotated after a day where a computer has been used. And those snapshots will be firmly offline and inaccessible.

I have previously mentioned that this is something I'm already doing using the Windows "MountVol" command and some registry editing to completely hide unmounted volumes. This allows me to transiently bring a well-hidden offline drive online while an image is being made of my main working system, but to keep it hidden otherwise.

We should also note that at the operating system level, what ransomware is doing really does stand out. It's very unusual behavior. The wholesale encryption of a large number of files is something that the OS itself should be readily able to observe, detect, and put a stop to. Although whitelisting every programs that a system uses can be quite burdensome, whitelisting the few programs that might legitimately need to encrypt lost of files -- if there are any such programs -- would not be a big problem.

Another anti-ransomware approach is to perform autonomous deep file versioning. And, as we see in a few weeks, that sort of ransomware protection is one of my requirements for effective cross-system folder syncing.