Security Now! #837 - 09-21-21
Cobalt Strike

This week on Security Now!

This week we examine a devastating and still ongoing DDoS attack against the latest in a series of VoIP service providers. We checkout the once again mixed blessing of last Tuesday's Microsoft patches, and we examine a welcome feature of Android 11 that's being back-ported through Android 6. We catch-up with Chrome's patching of two more new 0-day vulnerabilities and attacks, then we look at a "Pwnage" eMail I received from Troy Hunt's Have I Been Pwned site — was GRC Pwned? I then have a quick Sci-Fi reminder for the end of the week, a SpinRite update and a fun related YouTube posting. Then we'll wrap up by introducing the latest weapon in the malign perpetrator's arsenal, the powerful commercial tool known as Cobalt Strike.

“Give that person a raise!”

Security News

The DDoS attack on VoIP.ms

After last week's podcast discussing the history, evolution, technology and countermeasures of the Internet's denial of service attacks, many of our listeners tweeted to let me know of the ongoing multi-day attack against voip.ms, a voice over IP provider whose website states has about 80,000 customers. And several of those who tweeted to me were among their customers.

https://twitter.com/voipms

Until last Thursday, VoIP.ms's Twitter feedback had occasionally posted an advertisement to discuss and promote the various benefits of voice over IP services. For example, Monday before last on the 13th they happily tweeted the rhetorical question: "Are you looking for a proven growth lever for your business? Are you making full use of cloud communications? Read this article and discover all the advantages offered by this technology: https://hubs.li/H0S-4zP0" And prior to that, it was all similar. That was on the 13th.

But on the 16th they tweeted: "We're aware of an issue that's keeping our customers to properly reach our website. There's a team actively working on the issue as we speak and we expect it is fixed shortly, thanks for your patience!"

And a bit later: "We continue to work with our provider as our top priority and we'll be ready to provide a post-mortem of this event once the service is reestablished. Thanks for your kind understanding and please stay tuned!"

And then: "After further investigation, our service provider confirmed they are currently facing a network attack resulting in a denial of services. For all practical purposes this is making our Domain Name https://VoIP.ms unreachable at the moment,"

And then "For those of you who kindly switched from hostname to IP address and are within the US, we have made changes in the configuration with our upstream carriers to mitigate the issue for incoming calls as well."

"If you need to access your customer portal, you may be able to do so by adding a DNS entry to your localhost. Our Public website and customer portal IP address is: 173.231.187.61.
See detailed steps below on how to do this for a Windows Computer:" ... and they provide instructions.

This attack continued without letup throughout the weekend. Then Monday morning they tweeted: "We want to assure you that all our energy and resources are being put into fighting this ransom DDoS attack."

This was the first mention of a ransom. Then, at a little after 11:30am, yesterday, BleepingComputer posted some news of this:

> *Threat actors are targeting voice-over-Internet provider VoIP.ms with a DDoS attack and extorting the company to stop the assault that's severely disrupting the company's operation.*

> *VoIP.ms is an Internet phone service company that provides affordable voice-over-IP service to businesses around the world. On September 16th, 2021, VoIP.ms became the victim of a distributed denial-of-service attack targeting their infrastructure, including DNS name servers.*
>
> *As customers configured their VoIP equipment to connect to the company's domain name, the DDoS attack disrupted telephony services, preventing them from receiving or making phone calls.*
>
> *As DNS was no longer working, the company advised customers to modify their HOSTS file to point the domain at their IP address to bypass DNS resolution. However, this just led the threat actors to perform DDoS attacks directly at that IP address as well.*
>
> *To mitigate the attacks, VoIP.ms moved their website and DNS servers to Cloudflare, and while they reported some success, the company's site and VoIP infrastructure still have issues due to the continued denial-of-service attack.*
>
> *An announcement posted to the VoIP.ms website says:"A Distributed Denial of Service (DDoS) attack continues to be targeted at our Websites and POP servers. Our team is deploying continuous efforts to stop this however the service is being intermittently affected. We apologize for all the inconveniences."*
>
> *At the time of this writing, the site is bouncing back and forth between being accessible and displaying a 500 Internal Server Error. Today, customers continue to experience issues with their telephone service, including loss of service, dropped calls, poor performance, and the inability to forward lines.*
>
> *On September 18th, a threat actor using the name 'REvil' claimed responsibility for the attack and posted a link to a ransom note posted to Pastebin. This ransom note has since been removed from Pastebin, but BleepingComputer was told it asked for one bitcoin, or approximately $45,000, to stop the DDoS attacks. Soon after their original tweet, the threat actors raised their extortion demand to 100 bitcoins, or approximately $4.3 million.*
>
> *The customers' responses to the attack against VoIP.ms have been mixed.*
>
> *Some feel that VoIP.ms should pay the ransom to restore services before they themselves do not lose customers. At the same time, other VoIP.ms customers are vowing to stick with them and telling the company not to give in to the ransom demand.*
>
> *BleepingComputer has contacted VoIP.ms with questions regarding the attack but has not received a reply.*

And it strongly appears that this is only the most recent component of a larger campaign. Earlier this month multiple VoIP providers in the UK were very similarly targeted.

The site “Unified Communications” (https://www.uctoday.com) posted the news of the DDoS attacks in the UK two weeks ago, writing:

> *At least three UK VoIP providers have been hit by a DDoS attack, according to the Cloud Communications Alliance (CCA). In an email sent this morning, CCA said it has learned of a “sophisticated, specific and ongoing attack”, believed to be from Russian cybercriminal organisation REvil. Two providers revealed they were victims of the attack last week, but the third company was not named by CCA.*
>
> *Poole-based “Voip Unlimited” said the attack on its core network started on August 31St and was continuous for 75 hours. An update from this firm this morning said it did not observe any further attacks over the weekend.*
>
> *London-based Voipfone reported its attack at the same time and said this morning that its services are fully operational, with traffic being closely monitored.*
>
> *Both firms saw their services disrupted over the three-day period.*
>
> *CCA said that the culprits were demanding ransom, starting at 1 Bitcoin but quickly increasing.*
>
> *The attack involves hammering a company’s network with traffic of between 100-450 gigabits per second, often for up to 24 hours on multiple occasions. It starts with an attack on IP addresses used for SIP ingress and egress but then migrates to other services.*
>
> *CCA said the attack is capable of evading some typical DDoS prevention measures.*
>
> *Voip Unlimited Managing Director, Mark Pillow, told The Register: “At 2pm 31st August, Voip Unlimited’s network was the victim of an alarmingly large and sophisticated DDoS attack attached to a colossal ransom demand. UK Comms Council have communicated to us that other UK SIP (Session Initiation Protocol) providers are affected and identified them as a criminal hacking organisation called REvil who appear to be undertaking planned and organised DDoS attacks against VoIP companies in the UK.”*

Here's what's interesting about this: The RPS — Requests Per Second — attacks and their mitigations which we covered last week, which Cloudflare and others are able to provide, are extremely specific to filtering web requests being made by web browser clients. They can offer very good protection there. But only for that specific application. There is currently no provider of large pipe VoIP protocol DDoS protection. So it apparently dawned upon some miscreants somewhere that attacking the non-web servers of the Internet's global VoIP providers would be a new revenue source for extortion demands.

And they're probably not wrong. I agree with everyone that nothing about this sounds like the true REvil gang. The first ask for one bitcoin then for 100 seems quite ham-fisted and unlikely to result in payment which is, as we know, all the ransomware gangs claim to want.

It's going to be interesting to see how this evolves. There are myriad botnets on the Internet and there are many currently unprotected and hard-to-defend mission critical protocols other than standard web traffic. We may be at the beginning of a new era of cryptocurrency-enabled DDoS-driven extortion.

I was curious about where they are this Tuesday morning. So I checked in with VoIP.ms's Twitter feed and found the following:

> *We want to assure you that all our energy and resources are being put into fighting this ransom DDoS attack.*
>
> *We do feel grateful to count on your patience and understanding on these unfortunate events, please stay tuned on our social media feeds for further updates, thank you.*
>
> *Our entire team are working 24/7 on implementing all the required measures in order to have the service back up and running. We definitely understand your current level of frustration but please due to the current circumstances, we can't disclose our action plan at this time, don't doubt we are definitely using all internal and external resources that can be used in situations like these and that as soon as this event is over, your service will be just like you used it before the attacks.*
>
> [Then…]
>
> *All the team at http://VoIP.ms continues to work hard on recovering all services as soon as possible. With the help of internal and external specialists, all efforts and resources are being put into stabilizing our website and voice servers as well as protecting our network and mitigating this attack. Multiple changes have been applied in order to improve connectivity and reliability and we have observed positive changes in certain areas of the service.*
>
> *The following services have been recovered and are fully functional at the moment of this post: SMS / MMS / Recordings / Call Recordings / Conference Recordings*
>
> *Please continue following our social media channels and issue tracker to not miss any important updates.*

So it sounds as though actual VoIP calling remains offline due to this attack. They have messaging & playback of previous recorded audio working, but not real time interactive calling.

The great danger to them is that VoIP services don't come with a great deal of customer lock-in. So a percentage of their previous 80,000 customers, all of whom will have been without mission critical VoIP service since last Thursday, may be unable to tolerate this outage and will port their numbers to another VoIP carrier. And once they've done that, why would they return? There is a lot of conversation online about exactly how to leave VoIP.ms for other providers.

Patch Tuesday's Mixed Blessing

This being the third Tuesday of the month we're able to get some perspective on what last Tuesday's monthly patch Tuesday wrought. An "wrought" likely describes the mental state of at least a few of the industry's IT professionals following last Tuesday. It occurred to me that the well worn term "side effects" is apropos for labeling things that go wrong after one of Microsoft's monthly patch batches are installed. Like a powerful pharmaceutical drug that a doctor prescribes, the application of the month's patch updates attempts to fix things that really **do** need to be fixed. But also like so many powerful pharmaceuticals, there can often be unwanted side effects to make the patient—or the IT administrator—think twice before swallowing the pill.

On the “you really do need to take this pill” side, last week's batch of updates fixed a total of 86 vulnerabilities, two of them being 0-days being actively exploited ITW (which, you may recall, is our recently coined abbreviation for "In The Wild). The demographic breakdown of those was:

- 1 Denial of Service Vulnerability
- 2 Security Feature Bypass Vulnerabilities
- 8 Spoofing Vulnerabilities
- 11 Information Disclosure Vulnerabilities
- 16 Remote Code Execution Vulnerabilities
- 27 Elevation of Privilege Vulnerabilities

The biggest news was that, as we hoped would be true, Microsoft was, indeed, able to get that nasty MSHTML 0-day patched. That was the one that the hacker forums were having a field day with. The one where an ActiveX control embedded in various Office docs could cause IE's still-present engine core to be invoked to download a remote malicious site's content and compromise the system. There was no terrific workaround because the weak workarounds that had been proposed had all been worked around. So my advice and hope was that Microsoft would be able to deal with that one swiftly. And indeed they did.

Since then, we learned that the flaw was being used to download and execute a malicious DLL that, in turn, would install a Cobalt Strike beacon onto the victim's computer. The Cobalt Strike beacon allows bad guys to obtain remote access to the macchine, allowing them to exfiltrate files and also to spread laterally throughout the network.

So all that's the good news. And it provides sufficient reason for swallowing last week's update pill. But as I noted, this month's treatment was not without its side effects which are causing, believe it or not, many networked printer users to gag.

The newly introduced trouble appears to be the result of Microsoft's attempt to resolve the last outstanding known vulnerability with their very troubled networked printing management. The so-called PrintNightmare vulnerability is/was, and perhaps is still being tracked as CVE-2021-36958. After applying last week's patches, which included a fix for 36958, Windows admins began reporting wide-scale network printing problems.

https://docs.microsoft.com/en-us/answers/questions/517533/pint-server-and-print-nightmare-update.html

For example, from Microsoft's forum under the subject: “Print server and Print Nightmare update”

> *This just hit us this morning too. 9/15/2021. No one can print to the network printers.*
> *I removed KB5005613 from our server and rebooted the server and that fixed it. Had to do that at all 8 of our branch offices too. Microsoft updates seem to be more like hackers. Not professional.*

> *Today, 16 September 2021, I got the same problem, cannot print to printer on the server. Fortunately, I read this article and then I can assume what was happen to me, is caused by BAD Windows update. Then, I check Updates history, and find one update installed on 15 September 2021 (Security updates KB5005565). So, I uninstall it, and reboot. And, YES, the printer works normally,*

> *I Can CONFIRM we had the same Problem and nothing would work , even our Tec couldn't figure it out so i got on this Forum and YES The Above answer Solved our Prob , Deleted the Security updates KB5005565 and restarted and bingo , Printer can connect again..*

> Someone over on Reddit: *"Installed the Windows Updates on our small network today. Now none of our computers can print over the network, they can only print to local printers. Attempting to reinstall the printers gives the error "Windows cannot connect to the printer", and the print queue is showing a status of "Access Denied" for the pending jobs."*

So, the trouble is rather widespread and uninstalling that security update KB5005565 appears to be the only way to bring networked printing back online. But uninstalling the fix restores the vulnerability that was allowing bad guys to perform remote code execution and gain local SYSTEM privileges. The ransomware gangs, Vice Society, Magniber, and Conti, are all known to have jumped on this flaw and are now using it as long as it lasts to obtain elevated privileges on compromised machines.

And how long will it last? There's apparently no telling. Remember that this remaining PrintNightmare vulnerability was originally reported privately to Microsoft back in December of 2020 by Victor Mata of FusionX, Accenture Security. One hopes that we might see a permanent and working fix for this serious flaw before we get to THIS December, by which time Microsoft will have known of it for a full year.

A bitter pill to swallow, indeed.

And I keep coming back to the still unanswered question: Given Microsoft's virtually unlimited resources, how is it that they allow this to go on month after month after month?

Android to auto-reset app permissions on many more devices

Last Friday, Google announced that later this year, support for Android 11's privacy protection feature — which debuted one year ago with the release of Android 11 on September 8, 2020 — which proactively resets unneeded app permissions for applications that haven't been used in months, would be made available to billions of devices running older Android versions. Or at least to any of those which are able to update their OS.

This is great news. Very much like having old and unneeded database information self-deleting, auto-removing unneeded permissions is terrific security. It's unfortunate that so many older Android devices may never be able to obtain this. Google boasts of 3 Billion devices are running Android, but the timeline demographics suggest that older versions are not rushing to update:

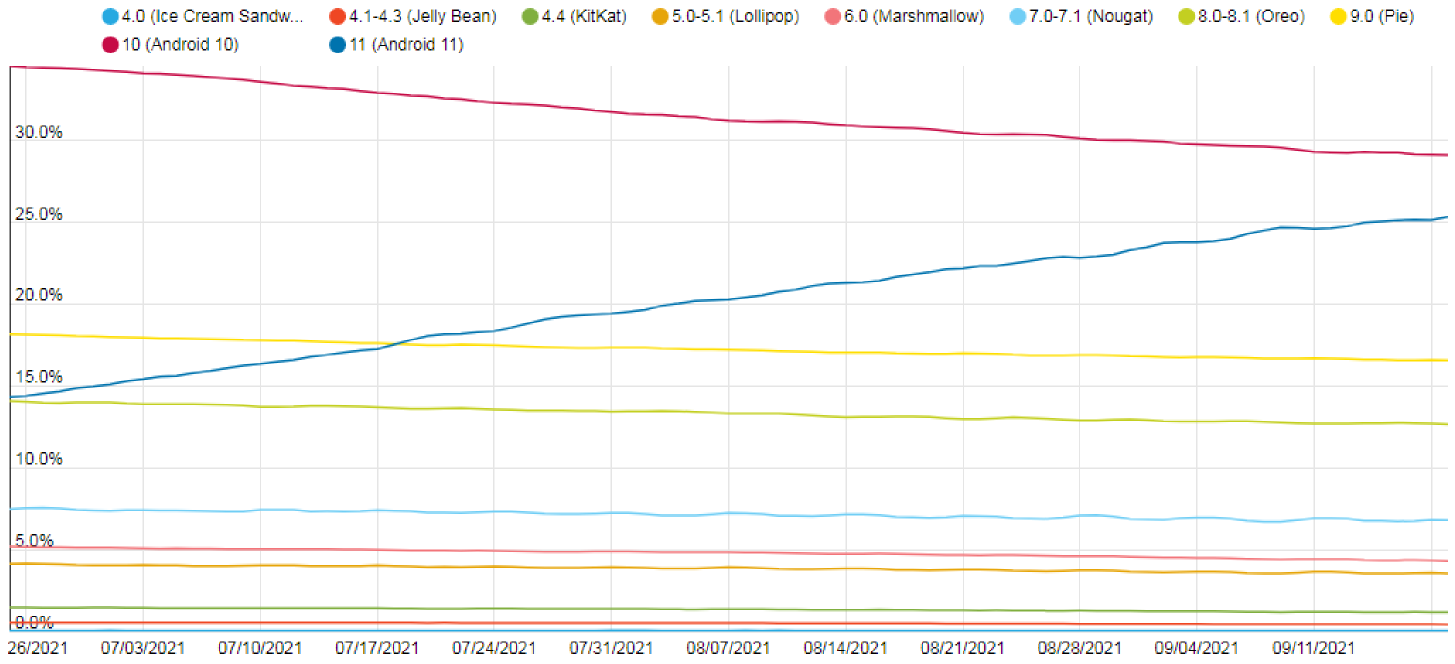

This chart makes the trend clear: Android 11's adoption is growing almost linearly, but slowing down slightly. But nearly all of its growth is coming from Android 10. The lines for Android 8 & 9 are drooping a little. But the older Androids are holding pretty steady. And this is what we'd expect, right? In time, the batteries of those increasingly decrepit Android devices will fail, or their radios will become obsolete as cellular technologies advance. So they will eventually be tossed into the recycle bin and their components and precious metals will hopefully be reprocessed to make new gadgets.

But for those devices which can upgrade, this new and very useful feature will become available on all devices with Google Play services running Android from 6.0 (API level 23) through Android 10 (API level 29).

In their posting, Google said: *"Starting in December 2021, we are expanding this to billions more devices. This feature will automatically be enabled on devices with Google Play services that are running Android 6.0 (API level 23) or higher. On these devices, users can now go to the auto-reset settings page and enable/disable auto-reset for specific apps. The system will start to automatically reset the permissions of unused apps a few weeks after the feature launches on a device."*

I studied Google's announcement and came away unsure whether this would be enabled by default on older devices. As we know, thanks for the tyranny of the default, the **only** thing that matters is whether this is enabled by default. The reason this is confusing is that Google's announcement says:

"This feature will automatically be enabled on devices with Google Play services that are running Android 6.0 (API level 23) or higher. The feature will be enabled by default for apps targeting Android 11 (API level 30) or higher. However, users can enable permission auto-reset manually for apps targeting API levels 23 to 29."

So this suggests that even when older devices are updated, their older apps which might remain unaware of the new support offered by Android 11, would not be automatically enabled. Presumably, Google was unwilling to change this behavior without the user being aware. But, if I'm understanding this correctly, I think that was a mistake. Just ask once for ALL older apps. That would not be unduly burdensome and it's the older devices that are most in need of shoring up.

Browser News

Google patched the 9th & 10th ITW 0-days in Chrome this year.

Everyone running Chrome desktop for Windows, Mac, and Linux will find themselves now running v93.0.4577.82. And that's good news, since the bad guys are scouring Chrome for way in, and so far they've been discovering and exploiting them all this year at a rate slightly faster than one every month.

Chrome's update to v93.0.4577.82 fixes a total of 11 security vulnerabilities, two of them being 0-days exploited in the wild. CVE-2021-30632 is an out-of-bounds write in the V8 JavaScript engine, and the CVE-2021-30633 bug is a use-after-free bug in the Indexed DB API. So they are both memory bugs. The release notes commented that: "Google is aware that exploits for CVE-2021-30632 and CVE-2021-30633 exist in the wild."

The two 0-day vulnerabilities fixed last week were disclosed to Google on September 8th. So Chromium gets fixed in less than a week. Five days, actually. This makes me even more glad that Microsoft decided that they were incapable of managing the ongoing development and maintenance of a start of the art web browser. They adopted the Chromium core. As we know, our web browsers are today's primary attack surface. We push them out onto the Internet and actively solicit unknown remote web servers to download and run their code on our machine. It sounds insane. And it really is. So Lord knows that we cannot be waiting nine months for Microsoft to get around to fixing critical problems that they've been informed of.

Miscellany

Was GRC Pwned?

The last time we talked about Troy Hunt's excellent "Have I Been Pwned?" web service, I noted that not only is it possible for individuals to subscribe to the site's autonomous notification via eMail in the event of that eMail address appearing in any new breach which HIBP adds to their massive and growing breach database, but that anyone who runs their own eMail domain can similarly submit their entire domain for notification.

And I did that quite a while ago. So I was interested to receive a notification on Sunday from HIBP informing me that one or more eMail accounts belonging to GRC.COM had just popped up in a new data breach. And, sure enough, while gathering the cybernews of the past week to share with everyone here, I encountered the mention that the site which had been mentioned by HIBP had been breached.

The site was EPIK.COM, which I had no memory of ever giving any eMail credentials to. And Wikipedia knows about them, saying: *“Epik is an American domain registrar and web hosting company known for providing services to websites that host far-right, neo-Nazi, and other extremist content.”* Uhhh... what?? And a breach of that site had leaked some GRC credentials?

Security industry coverage of this breach notes that EPIK is the host of sites including Gab, Parler, and “The Donald”. And the reputable site “The Record”, which was tipped off of the breach on Monday, first received a small subset of samples which was later followed by a full copy of the entire leak from an individual who claimed to be loosely associated with the Anonymous group — the group which was proudly claiming responsibility for the breach and extensive exfiltration. When The Record then reached out for comment last Tuesday, Epik denied the breach and the hackers’ claims in an email reply. Epik wrote:

“We are not aware of any breach. We take the security of our clients’ data extremely seriously, and we are investigating the allegation. — [signed] Epik spokesperson”

Despite Epik’s denial, the data The Record received in full and reviewed confirms the hackers’ claims. In a 32 GB torrent file hosted through the DDoSecrets portal, the hackers included several SQL database dumps containing gigabytes of sensitive information such as domain ownership details, domain transactions, account details, and troves of personal data points.

Okay. So presumably a subset of this massive trove of information had been provided to Troy and was added to his Have I Been Pwned database, whereupon any of those who had previously signed up for notification would receive an eMail, just as I had. But I still had no idea how GRC could possibly have been part of the breach of such a domain registrar and hosting provider — of which I still had no memory. So I went over to HIBP to see what was up. It's possible to query Troy's terrific site for any matches on an eMail address or, if one owns an entire domain, any matches against that domain. So I did that and I found the two eMail addresses that had apparently been “leaked” by the breach of EPIK.COM:

network-solutions-public-whois@grc.com
whois2011-1@grc.com

Whew! So, it was neither I nor Sue nor Greg’s nor any corporate eMail. Returning to the eMailed notification I had received from HIBP, I discovered that the notification had provided an interesting description of the breached site along with the notification. HIBP says:

> *In September 2021, the domain registrar and web host Epik suffered a significant data breach, allegedly in retaliation for hosting alt-right websites. The breach exposed a huge volume of data not just of Epik customers, but also scraped WHOIS records belonging to individuals and organisations who were not Epik customers. The data included over 15 million unique email addresses (including anonymised versions for domain privacy), names, phone numbers, physical addresses, purchases and passwords stored in various formats.*

So, whew. That all fits. Those two grc.com eMail addresses would have, indeed, been scraped from some of GRC's public ICANN domain registrations. And who knows why this distasteful seeming domain registrar might have had them? Perhaps to use is spamming the others of other domains with promotional “come on over here” registration solicitations?

In any event, I thought it was cool that HIBP's proactive service works, since this was the first time it had ever had the occasion to trigger for me. And I wanted to take the opportunity to remind our listeners of Troy’s useful and 100% free service.

Sci-Fi

I wanted to remind our listeners that the first three installments of Apple's 10-part miniseries, based upon Isaac Asimov's Foundation trilogy, becomes available this Friday the 24th. Lorrie and I just had some dinner party plans moved from Friday to Thursday, which worked for me since I’d love to be able to set aside this Friday night for those first three Foundation episodes.

We will need to wait another month for the remake of “Dune” which will be appearing on HBOMax, as well as for AppleTV's release of their 10-part miniseries “Invasion” which chronicles an invasion of the Earth by hostile extraterrestrials.

So, it's looking like the Sci-Fi lovers among us will finally be having some fun new visuals to enjoy.

SpinRite

My work on SpinRite is progressing. As I last noted, I was becoming uncomfortable with having written so much completely untested code. So I decided to stop writing, to exercise and debug everything that I've written. SpinRite has always incorporated drive benchmarks. They perform deliberately repetitive and non-repetitive reads to measure various aspects of a drive's physical data read and cached data performance. The benchmarks also have the benefit of giving SpinRite's new IO abstraction system a rather thorough workout, which is precisely what it needs. So I have updated the benchmarking code to work with the new system and I'm nearly ready to make another test release available to the SpinRite testing gang. Once the dust has settled from that, I'll finish updating the data recovery code, which was where I was when I stopped because I was feeling that too much had not been tested. Then I'm pretty sure that the logging system will need adjusting. I've already been in that code updating things, but nothing was running back then, so I wasn't able to confirm that the various new formatting things I had written were working. And while I'm sure that other stuff will come up... it really does feel at though we're sort of getting to the point where we can see the light at the end of the very long tunnel.

But in the meantime, I have something VERY FUN to share:

A guy by the name of Adrian Black has a retro computing YouTube channel called “Adrian's Digital Basement” to which a bunch of our listeners subscribe. I know that because I began receiving tweets informing me of Adrian’s latest posting. A couple of Saturday's ago, Adrian was playing with a sort of amazing emulator for the very first MFM (modified frequency modulation) hard drives of the sort that the first IBM XT could be equipped with. Anyone who recalls the iconic golden-shelled 10 megabyte Seagate ST-225 drive will see one sitting there on Adrian's workspace.

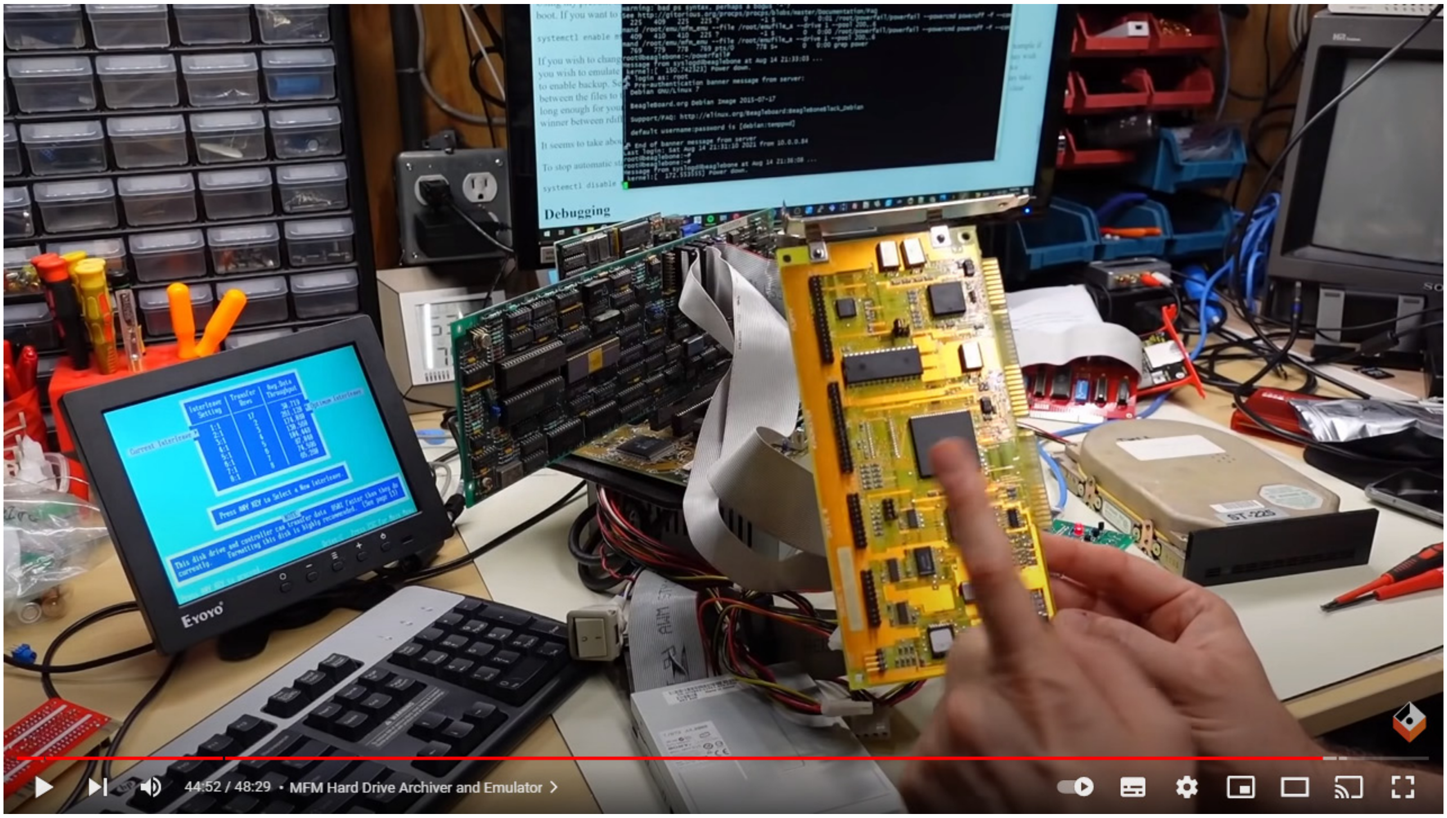

But the point is, this MFM hard drive emulator is not transferring data very quickly, and Adrian suspects that perhaps that's because the emulator was low-level formatted with a one-to-one sector interleave. Which is to say, with no sector interleaving at all. But the controller he has the emulator hooked up to needs some sector interleaving. It’s not quick enough to catch the immediate next sector after reading the previous one.

So... without batting an eye, Adrian he fires up an ancient copy of SpinRite and has SpinRite examine the current speed and interleave... then he has SpinRite optimize the drive's interleave, just like in the old days, by successively trying each interleave I turn and measuring the drive's data transfer performance at each interleave setting, displaying a correspondence table.

This week's GRC shortcut link will jump you 41 minutes in, to just before Adrian begins to run his ancient copy of SpinRite on the emulator: https://grc.sc/837

https://www.youtube.com/watch?v=q__R8khTwo8&t=2483s

Cobalt Strike

Cobalt Strike dot Com proudly introduces newcomers to their quite pricey offering with the headline: “Software for Adversary Simulations and Red Team Operations”

https://www.cobaltstrike.com/

And then goes on to explain: *“Adversary Simulations and Red Team Operations are security assessments that replicate the tactics and techniques of an advanced adversary in a network. While penetration tests focus on unpatched vulnerabilities and misconfigurations, these assessments benefit security operations and incident response.*

Cobalt Strike gives you a post-exploitation agent and covert channels to emulate a quiet long-term embedded actor in your customer’s network. Malleable C2 lets you change your network indicators to look like different malware each time. These tools complement Cobalt Strike’s solid social engineering process, its robust collaboration capability, and unique reports designed to aid blue team training.

Raphael Mudge created Cobalt Strike in 2012 to enable threat-representative security tests. Cobalt Strike was one of the first public red team command and control frameworks. In 2020, HelpSystems acquired Cobalt Strike to add to its Core Security portfolio. Today, Cobalt Strike is the go-to red team platform for many U.S. government, large business, and consulting organizations.”

Unfortunately... it's so good that it has also rapidly become the go-to platform for non-simulated real world malware post-penetration network infiltration. We've seen many examples over the years of well designed and well-meaning utilities being commandeered and abused for malicious purposes.

For example, SysInternals' Mark Russinovich create PSEXEC. Its description at Microsoft (who, as we know, purchased SysInternals back in the summer of '06) says:

Utilities like Telnet, and remote control programs like Symantec's PC Anywhere, let you execute programs on remote systems, but they can be a pain to set up, and require that you install client software on the remote systems that you wish to access. PsExec is a light-weight telnet-replacement that lets you execute processes on other systems, complete with full interactivity for console applications, without having to manually install client software. PsExec's most powerful uses include launching interactive command-prompts on remote systems and remote-enabling tools like IpConfig that otherwise do not have the ability to show information about remote systems.
[It's at this point that we all chant in unison: “What could possibly go wrong?”]

Note: some anti-virus scanners report that one or more of the tools are infected with a "remote admin" virus. None of the PsTools contain viruses, but they have been used by viruses, which is why they trigger virus notifications.

Yeah. And as another example, I chose the lovely “Remote Utilities” solution for my own and Lorrie's remote system control needs for the same reasons that it has become the remote control solution of choice for nameless miscreants across the Internet. It's not its fault, of course. Just like a compiler compiles code. It can't help it if the code it's compiling is malware.

So today, **Cobalt Strike** deserves and receives our attention because it appears that we're just at the beginning of a wave of malicious exploitation which is being enabled by Cobalt Strike...

Back in July, the security firm “ProofPoint” published, and then recently updated, a report titled: “Cobalt Strike: Favorite Tool from APT to Crimeware.” — APT, of course being the abbreviation for Advanced Persistent Threat.

One of the terms we’ll encounter today and in the future is “Beacon” or “Cobalt Strike Beacon.” The Beacon is the bad bit that's infiltrated into an unwitting victim's machine. And unfortunately it was quite well designed. Listen to how Colbalt Strike themselves boast about Beacon's capabilities:

Beacon is Cobalt Strike's payload to model advanced attackers. Use Beacon to egress a network over HTTP, HTTPS, or DNS. You may also limit which hosts egress a network by controlling peer-to-peer Beacons over Windows named pipes.

Beacon is flexible and supports asynchronous and interactive communication. Asynchronous communication is low and slow. Beacon will phone home, download its tasks, and go to sleep. Interactive communication happens in real-time.

Beacon's network indicators are malleable. Redefine Beacon's communication with Cobalt Strike's malleable C2 language. This allows you to cloak Beacon activity to look like other malware or blend-in as legitimate traffic.

They then go into some additional detail which should chill the blood of any CISO:

“Right-click on a Beacon session and select interact to open that Beacon's console. The console is the main user interface for your Beacon session. The Beacon console allows you to see which tasks were issued to a Beacon and to see when it downloads them. The Beacon console is also where command output and other information will appear.

Be aware that Beacon is an asynchronous payload. Commands do not execute right away. Each command goes into a queue. When the Beacon checks in (connects to you), it will download these commands and execute them one by one. At this time, Beacon will also report any output it has for you. If you make a mistake, use the clear command to clear the command queue for the current Beacon.

By default, Beacons check in every sixty seconds. You may change this with Beacon’s sleep command. Use sleep followed by a time in seconds to specify how often Beacon should check in. You may also specify a second number between 0 and 99. This number is a jitter factor. Beacon will vary each of its check in times by the random percentage you specify as a jitter factor. For example, sleep 300 20, will force Beacon to sleep for 300 seconds with a 20% jitter percentage. This means, Beacon will sleep for a random value between 240s to 300s after each check-in.

To make a Beacon check in multiple times each second, try sleep 0. This is interactive mode. In this mode commands will execute right away. You must make your Beacon interactive before you tunnel traffic through it. A few Beacon commands (e.g., browserpivot, desktop, etc.) will automatically put Beacon into interactive mode at the next check in.

Running Commands:

Beacon's shell command will task a Beacon to execute a command via cmd.exe on the compromised host. When the command completes, Beacon will present the output to you.

Use the run command to execute a command without cmd.exe. The run command will post output to you. The execute command runs a program in the background and does not capture output.

Use the powershell command to execute a command with PowerShell on the compromised host. Use the powerpick command to execute PowerShell cmdlets without powershell.exe. This command relies on the Unmanaged PowerShell technique developed by Lee Christensen. The powershell and powerpick commands will use your current token.

The psinject command will inject Unmanaged PowerShell into a specific process and run your cmdlet from that location.

The powershell-import command will import a PowerShell script into Beacon. Future uses of the powershell, powerpick, and psinject commands will have cmdlets from the imported script available to them. Beacon will only hold one PowerShell script at a time. Import an empty file to clear the imported script from Beacon.

The execute-assembly command will run a local .NET executable as a Beacon post-exploitation job. You may pass arguments to this assembly as if it were run from a Windows command-line interface. This command will also inherit your current token.

If you want Beacon to execute commands from a specific directory, use the cd command in the Beacon console to switch the working directory of the Beacon’s process. The pwd command will tell you which directory you’re currently working from.

Last, but not least, Beacon can execute Beacon Object Files without creating a new process. Beacon Object Files are compiled C programs, written to a specific convention, that run within a Beacon session. Use inline-execute [args] to execute a Beacon Object File with the specified arguments.

Use the spawn command to spawn a session for a listener. The spawn command accepts an architecture (e.g., x86, x64) and a listener as its arguments.

By default, the spawn command will spawn a session in rundll32.exe. An alert administrator may find it strange that rundll32.exe is periodically making connections to the internet. Find a better program (e.g., Internet Explorer) and use the spawnto command to state which program Beacon should spawn sessions into.

The spawnto command requires you to specify an architecture (x86 or x64) and a full path to a program to spawn, as needed. Type spawnto by itself and press enter to instruct Beacon to go back to its default behavior.

Type inject followed by a process id and a listener name to inject a session into a specific process. Use ps to get a list of processes on the current system. Use inject [pid] x64 to inject a 64-bit Beacon into an x64 process.

The spawn and inject commands both inject a payload stage into memory. If the payload stage is an HTTP, HTTPS, or DNS Beacon and it can't reach you—you will not see a session. If the payload stage is a bind TCP or SMB Beacon, these commands will automatically try to link to and assume control of these payloads.

Use "dllinject" [pid] to inject a Reflective DLL into a specific process.

Use the shinject [pid] [architecture] [/path/to/file.bin] command to inject shellcode, from a local file, into a process on target.

Use shspawn [archicture] [/path/to/file.bin] to spawn the "spawn to" process and inject the specified shellcode file into that process.

Use dllload [pid] [c:\path\to\file.dll] to load an on-disk DLL in another process.

Use ppid [pid] to assign an alternate parent process for programs run by your Beacon session. This is a means to make your activity blend in with normal actions on the target. The current Beacon session must have rights to the alternate parent and it's best if the alternate parent process exists in the same desktop session as your Beacon. Type ppid, with no arguments, to have Beacon launch processes with no spoofed parent.

The runu command will execute a command with another process as the parent. This command will run with the rights and desktop session of its alternate parent process. The current Beacon session must have full rights to the alternate parent.

The spawnu command will spawn a temporary process, as a child of a specified process, and inject a Beacon payload stage into it. The spawnto value controls which program is used as a temporary process.

This thing is a dream come true for the underworld. It was deliberately and consciously designed by advanced cyber security experts specifically to operate as covertly as possible, using every advanced trick in the book. And it's now loose, being proactively used not for red vs blue team training, but in the wild by threat actors. And its use is spreading, not just in depth, but dramatically in breadth. Here's what Proof Point describes and found:

In December 2020, the world learned about an expansive and effective espionage campaign that successfully backdoored the popular network monitoring software SolarWinds. Investigators revealed tools used by the threat actors included Cobalt Strike Beacon. This campaign was attributed to threat actors working for Russia's Foreign Intelligence Service – a group with Cobalt Strike in their toolbox since at least 2018. This high-profile activity was part of a clever attack *chain enabling advanced threat actors to surreptitiously compromise a relatively small number of victims. The tool used, and customized to fit their needs, is almost a decade old but increasingly popular.*

Cobalt Strike debuted in 2012 in response to perceived gaps in an existing red team tool, the Metasploit Framework. In 2015, Cobalt Strike 3.0 launched as a standalone adversary emulation platform. By 2016, Proofpoint researchers began observing threat actors using Cobalt Strike.

Historically, Cobalt Strike use in malicious operations was largely associated with well-resourced threat actors, including large cybercrime operators like TA3546 (also known as FIN7), and advanced persistent threat (APT) groups such as TA423 (also known as Leviathan or APT40). Proofpoint researchers have attributed two-thirds of identified Cobalt Strike campaigns from 2016 through 2018 to well-resourced cybercrime organizations or APT groups. That ratio decreased dramatically the following years – between 2019 and present, just 15 percent of Cobalt Strike campaigns were attributable to known threat actors.

In other words, the word got out that "Cobalt Strike" was the tool to use, and everyone began picking it up and using it. ProofPoint notes that:

"*Threat actors can obtain Cobalt Strike in a variety of ways: purchasing it directly from the vendor's website, which requires verification; buying a version on the dark web via various hacking forums; or using cracked, illegitimate versions of the software. In March 2020, a cracked version of Cobalt Strike 4.0 was released and made available to threat actors."*

And in making even more clear the appeal that Cobalt Strike offers, ProofPoint explains:

Cobalt Strike is used by a diverse array of threat actors, and while it is not unusual for cybercriminal and APT actors to leverage similar tooling in their campaigns, Cobalt Strike is unique in that its built-in capabilities enable it to be quickly deployed and operationalized regardless of actor sophistication or access to human or financial resources.

Cobalt Strike is also session-based — that is, if threat actors can access a host and complete an operation without needing to establish ongoing persistence, there will not be remaining artifacts on the host after it is no longer running in-memory. In essence: they can hit it and forget it.

Threat actors can also use the malleability of Cobalt Strike to create customized builds that add or remove features to achieve objectives or evade detection. For example, APT29 frequently uses custom Cobalt Strike Beacon loaders to blend in with legitimate traffic or evade analysis.

For defenders, customized Cobalt Strike modules often require unique signatures, so threat detection engineers may be required to play catch-up to Cobalt Strike use in the wild. Cobalt Strike is also appealing to threat actors for its inherent obfuscation. Attribution gets more difficult if everyone is using the same tool. If an organization has a red team actively making use of it, it is possible malicious traffic could be mistaken as legitimate. The software's ease of use can improve the capabilities of less sophisticated actors. For sophisticated actors, why spend development cycles on something new when you already have a great tool for the job?

Proofpoint data shows Cobalt Strike is a popular tool for everything from strategic compromises to noisy, widespread campaigns.

So, a top-of-the-line tool, which was designed to enable good guys to steathfully penetrate and inhabit their own networks for the purpose of testing and training anti-penetration and intrusion detection teams, and which licenses for $3500 per seat... got loose and out of control from its publisher. Now, this high-end professional tool, which no script kiddie would have ever been able to create for themselves, is being used in the wild on a more than daily basis.

There's little doubt that we'll be encountering the term "Cobalt Strike" in the future. Now we'll all know exactly what that means and entails.