Security Now! #765 - 05-05-20

An Authoritarian Internet?

This week on Security Now!

This week we add Bruce Schneier's thoughts about the theoretical feasibility of contact tracing apps; we touch on our government's feelings about DNS over HTTPS; we look at yet another whacky way of exfiltrating data from an air-gapped computer; we examine a new vulnerability that has already damaged some large high-profile enterprise infrastructures; we note Adobe's latest round of critical updates, another welcome service coming from Mozilla, a dispiriting bit of over-the-top political correctness from the UK, and Google's plans to clean up the mess which is the Chrome Web Store. We then share a bit of errata, miscellany and SpinRite news, then take a look at China's proposed changes to the fundamental operation of our global Internet.

Tech Talk | Computing | Software

10 Apr 2020 | 20:17 GMT

Cobol Programmers Answer Call to Shore Up Unemployment Benefits Systems

Retirees and newcomers want to help fix old software overloaded by new claims caused by the coronavirus pandemic

https://spectrum.ieee.org/tech-talk/computing/software/cobol-programmers-answer-call-unemployment-benefits-systems

Security News

Bruce Schneier on COVID-19 Contact Tracing Apps
https://www.schneier.com/blog/archives/2020/05/me_on_covad-19_.html

Bruce Schneier begins his blog posting titled: "Me on COVID-19 Contact Tracing Apps" by writing: "I was quoted in BuzzFeed:" and his quote reads:

> *"My problem with contact tracing apps is that they have absolutely no value,"* Bruce Schneier, a privacy expert and fellow at the Berkman Klein Center for Internet & Society at Harvard University, told BuzzFeed News. *"I'm not even talking about the privacy concerns, I mean the efficacy. Does anybody think this will do something useful? ... This is just something governments want to do for the hell of it. To me, it's just techies doing techie things because they don't know what else to do."*

Then, elaborating in his blog posting, Bruce continues:

I haven't blogged about this because I thought it was obvious. But from the tweets and emails I have received, it seems not.

This is a classic identification problem, and efficacy depends on two things: false positives and false negatives.

- False positives: Any app will have a precise definition of a contact: let's say it's less than six feet for more than ten minutes. The false positive rate is the percentage of contacts that don't result in transmissions. This will be because of several reasons. One, the app's location and proximity systems -- based on GPS and Bluetooth -- just aren't accurate enough to capture every contact. Two, the app won't be aware of any extenuating circumstances, like walls or partitions. And three, not every contact results in transmission; the disease has some transmission rate that's less than 100% (and I don't know what that is).

- False negatives: This is the rate the app fails to register a contact when an infection occurs. This also will be because of several reasons. One, errors in the app's location and proximity systems. Two, transmissions that occur from people who don't have the app (even Singapore didn't get above a 20% adoption rate for the app). And three, not every transmission is a result of that precisely defined contact -- the virus sometimes travels further.

Assume you take the app out grocery shopping with you and it subsequently alerts you of a contact. What should you do? It's not accurate enough for you to quarantine yourself for two weeks. And without ubiquitous, cheap, fast, and accurate testing, you can't confirm the app's diagnosis. So the alert is useless.

Similarly, assume you take the app out grocery shopping and it doesn't alert you of any contact. Are you in the clear? No, you're not. You actually have no idea if you've been infected.

The end result is an app that doesn't work. People will post their bad experiences on social media, and people will read those posts and realize that the app is not to be trusted. That loss of trust is even worse than having no app at all.

It has nothing to do with privacy concerns. The idea that contact tracing can be done with an app, and not human health professionals, is just plain dumb. ***[End of Blog]***

Okay. As we know, our April 14th podcast thoroughly analyzed the technology and found it to be essentially perfectly implemented. Then last week we looked at some tweaks that had been added. Encrypting the Bluetooth metadata was useful. The rest was just feel good nonsense designed to placate someone, somewhere. But the result is still a very robust system that would do what it's intended to do.

Perhaps it'll do some good. But Bruce makes a VERY GOOD point about adoption rate. Even without the misguided and outspoken pseudo-technical press completely misunderstanding how this works because, after all it IS complicated, a huge percentage of people would inherently mistrust it. And really, who would blame them? If we here participating in this podcast hadn't taken a close look at exactly what it does and how it works, we would be in the same boat. And, as I noted last week, we STILL need to take it on faith that those who use this underlying solidly privacy-protecting API as the foundation of their contact tracing apps do not deliberately, or perhaps inadvertently, bundle a bunch of deanonymizing technology on top of the API.

But Bruce ended that blog posting by saying: *"The idea that contact tracing can be done with an app, and not human health professionals, is just plain dumb."*

Okay, so what about human health professionals performing contact tracing? At the end of last week, New York State's governor, Andrew Cuomo explained that the previous day the state had identified 4,681 new cases of Coronavirus and that about the same would be happening again, that same day. And, pretty much every day. He explained that effective contact tracing required that every one of those 4,681 people be interviewed to determine the identities of everyone they had come into physical proximity to over the past FOURTEEN days. And then, that all of THOSE people would need to be interviewed and perhaps placed into isolation. 4 thousand, 6 hundred and 81... per day… that then expands to ALL the people any of them were in proximity to at any point during the previous 14 days.

Does anyone believe that's possible? I can't see any way that this is not utter nonsense.

It'll create some badly needed paying jobs. The most recent estimate is that 300,000 human contact tracers are needed. So that will be good. But it's not going to be effective in stopping the spread of this virus. We're learning that entirely asymptomatic people can be long-term carriers, and there was an instance of someone getting sick, recovering quickly and feeling fine. He wanted to be let out of the hospital but he needed to be cleared first by testing negative with a PCR test for the live virus. That took an additional 50 days! He was contagiously shedding the live virus for 50 days **after** he felt fine. And we're told to stay home when we get sick, not to go to the hospital unless needed. So how many of these viral shedders do we have out and about?

Having looked at all the data, I believe that it's not a matter of WHEN we reopen, it's HOW we reopen. And frankly, the sooner the better… so long as it's done carefully. Because this is doing serious damage to our country and to the world. If sufficient care is exercised, the R-naught can be pushed and held below 1.0. That will keep the virus from exploding. But it does mean that it will slowly be winding its way through our world's population. All the numbers show that's now inevitable.

DHS's CISA says no to 3rd-party DoH.

First of all, in this abbreviation soup, DHS is, of course, the US Department of Homeland Security. And CISA is the Cybersecurity & Infrastructure Security Agency, which is an agency within the DHS. And we all know that DoH is DNS over HTTPS.

In a recent 4-page memorandum, the US Department of Homeland Security reminded all Federal CISOs that they must not use the DoH services which are coming to all of our web browsers.

https://www.cisa.gov/sites/default/files/publications/Addressing_DNS_Resolution_on_Federal_Networks_Memo.pdf

> *"The purpose of this memorandum, issued pursuant to authorities under section 3553(b) of Title 44, U.S. Code, and Title XXII of the Homeland Security Act of 2002, as amended, is to remind agencies of their legal requirement to use EINSTEIN 3 Accelerated's Domain Name System (DNS) sinkholing capability for DNS resolution [pause] and [to] provide awareness about recent security and privacy enhancements to DNS resolution protocols – in particular, DNS over HTTPS (DoH) and DNS over TLS (DoT)."*

This EINSTEIN 3 thing started off life as an intrusion detection system designed by the DHS's US-CERT. Version 1 allowed the Agency to monitor traffic across all government networks. Version 2 added the ability to spot suspicious traffic. And Version 3, where we are now, with Einstein 3 Accelerated, AKA Einstein 3A, went still further, preventing unwanted intrusions by known bad actors. It offers useful DHS-specific services like sink-holing that override public DNS records by blocking access to destinations that the DHS knows to be malicious. It also lets the DHS examine all DNS requests made by government users. And we know how useful THAT can be.

The memo notes that:

> *CISA encourages efforts to make network communications encrypted by default. Doing so increases user security, making it harder for attackers to monitor and modify communication. DoH and DoT add desirable security features to DNS resolution; however, federal agencies that use DNS resolvers other than E3A lose the protection that defensive DNS filtering provides, and E3A does not currently offer encrypted DNS resolution. CISA intends to offer a DNS resolution service that supports DoH and DoT in time. Until then, agencies must use E3A for DNS resolution.*
>
> *Required Action.*
> *In accordance with 6 U.S.C. §663 note, "Agency Responsibilities," ensure local DNS recursive resolvers use E3A as their primary (or ultimate) upstream DNS resolver.*

In other words, we looooooove DNS encryption, but not if we cannot monitor it for your own safety, and not if we are unable to apply our spiffy EINSTEIN DNS filtering bad-domain sinkhole system to keep all you government workers from going to naughty places. And unfortunately we don't currently offer DoH or DoT encryption of our value-added DNS, so be super-sure to turn off that on-by-default DoH resolution that Firefox and Google and soon Microsoft will all be using…

or you're going to get in trouble.

"POWER-SUPPLaY"

The somewhat nutty, but endlessly clever guys at the Cyber-Security Research Center of the Ben-Gurion University of the Negev, Israel, specifically doctor Mordechai Guri, have come up with yet another sneaky, if not entirely practical, means for exfiltrating an air-gapped computer's digital data, by way of its switching power supply: https://arxiv.org/pdf/2005.00395.pdf

Mordechai's research has been published with the title: "POWER-SUPPLaY: Leaking Data from Air-Gapped Systems by Turning the Power-Supplies Into Speakers"

His paper's Abstract reads:

> It is known that attackers can exfiltrate data from air-gapped computers through their speakers via sonic and ultrasonic waves. To eliminate the threat of such acoustic covert channels in sensitive systems, audio hardware can be disabled and the use of loudspeakers can be strictly forbidden. Such audio-less systems are considered to be audio-gapped, and hence immune to acoustic covert channels.
>
> In this paper, we introduce a technique that enables attackers to leak data acoustically from air-gapped and audio-gapped systems. Our developed malware can exploit the computer power supply unit (PSU) to play sounds and use it as an out-of-band, secondary speaker with limited capabilities.
>
> The malicious code manipulates the internal switching frequency of the power supply and hence controls the sound waveforms generated from its capacitors and transformers. Our technique enables producing audio tones in a frequency band of 0-24khz and playing audio streams (e.g.,WAV) from a computer power supply without the need for audio hardware or speakers. Binary data (files, keylogging, encryption keys, etc.) can be modulated over the acoustic signals and sent to a nearby receiver (e.g., smartphone).
>
> We show that our technique works with various types of systems: PC workstations and servers, as well as embedded systems and IoT devices that have no audio hardware at all. We provide technical background and discuss implementation details such as signal generation and data modulation. We show that the POWER-SUPPLaY code can operate from an ordinary user-mode process and doesn't need any hardware access or special privileges. Our evaluation shows that using POWER-SUPPLaY, sensitive data can be exfiltrated from air-gapped and audio-gapped systems from a distance of five meters away at a maximal bit rates of 50 bits/sec.

And we should remember that while no one is going to transfer anything massive from a computer at 50 bits per second, the elliptic curves which are coming into increasing use provide state-of-the-art security using only a 256-bit key... which could be sent in just 5 seconds. Since every bit of a key is critical and no bits can be inferred from others, I would encode the burst transmission with ample error correction added. That might add another 30% to its size but only a couple of seconds to its burst duration.

Stealing Data from Air-Gapped Devices

Using Power Supply as an Out-of-Band Speaker

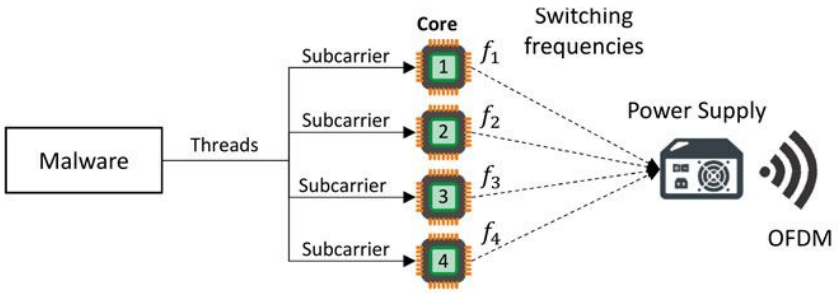

The system operates by causing a multicore processor to draw varying amounts of power from the system's power supply. As the power drawn changes, the switching power supply changes its switching frequency to supply more or less power as needed. It is these changes in switching frequency that can be used to send ultrasonic signals from an air-gapped computer system. Mordechai chose to use four distinct frequencies so that 2-bits could be sent at one time. Back in the era of analog telephone modems, similar multi-bit tone signalling was used to increase the number of bits that could be sent simultaneously.

The result is an effective system which can send 50 bits per second to a smartphone placed two and a half meters away. While doing this might seem of dubious value to us, I'd be surprised if our spooks at the NSA and CIA weren't thinking "Hmmm... I know just where to use this!"

An authorization bypass in SaltStack

An authorization bypass in SaltStack

We've never had the occasion to touch on SaltStack. But serious security vulnerability in this cloud resource management system changes that.

SaltStack livesover on GitHub: https://github.com/saltstack/salt where SaltStack explains itself:

> SaltStack makes software for complex systems management at scale. SaltStack is the company that created and maintains the Salt Open project and develops and sells SaltStack Enterprise software, services and support. Easy enough to get running in minutes, scalable enough to manage tens of thousands of servers, and fast enough to communicate with them in seconds.
>
> Salt is a new approach to infrastructure management built on a dynamic communication bus. Salt can be used for data-driven orchestration, remote execution for any infrastructure, configuration management for any app stack, and much more.
>
> Salt Open is tested and packaged to run on CentOS, Debian, RedHat Enterprise, Ubuntu, and Windows.

The good news is, we've never had the occasion to talk about it before now. The bad news is... Now we have.

F-Secure wrote that they discovered a number of vulnerabilities in the "Salt" management framework by the company SaltStack. They wrote:

> The open source Salt project is at the heart of SaltStack's (the company) product offerings but is also very popular as a configuration tool to manage servers in datacenters and cloud environments.
>
> Salt is used to monitor and update the state of servers. Each server runs an agent called a "minion" which connects to a "master", a Salt installation that collects state reports from minions and publishes update messages that minions can act on. Typically, such messages are updates to the configuration of a selection of servers, but they can also be used to run the same command in parallel over multiple, even all, managed systems asynchronously.
>
> The default communication protocol in salt is ZeroMQ. The master exposes two ZeroMQ instances, one called the "request server" where minions can connect to report their status (or the output of commands) and one called the "publish server" where the master publishes messages that the minions can connect and subscribe to.
>
> The vulnerabilities described in this advisory allow an attacker who can connect to the "request server" port to bypass all authentication and authorization controls and publish arbitrary control messages, read and write files anywhere on the "master" server filesystem and steal the secret key used to authenticate to the master as root. The impact is full remote command execution as root on both the master and all minions that connect to it.
>
> The vulnerabilities, allocated CVE ids CVE-2020-11651 CVE-2020-11652, are of two different classes. One being authentication bypass where functionality was unintentionally exposed to unauthenticated network clients, the other being directory traversal where untrusted input (i.e. parameters in network requests) was not sanitized correctly allowing unconstrained access to the entire filesystem of the master server.
>
> SaltStack engineers patched these vulnerabilities in release 3000.2 and users of Salt are encouraged to make sure that their installs are configured to automatically pull updates from SaltStacks repository server, see https://repo.saltstack.com/ for more information. A patch release for the previous major release version is also available, with version number 2019.2.4.
>
> Adding network security controls that restrict access to the salt master (ports 4505 and 4506 being the defaults) to known minions, or at least block the wider Internet, would also be prudent as the authentication and authorization controls provided by Salt are not currently robust enough to be exposed to hostile networks.
> Detection
>
> A scan revealed over 6,000 instances of this service exposed to the public Internet. Getting all of these installs updated may prove a challenge as we expect that not all have been configured to automatically update the salt software packages.

They then go into detail about the specific vulnerabilities.

It was interesting that F-Secure had difficulty getting the news of these critical vulnerabilities to SaltStack. In their disclosure timeline, F-Secure noted:

- 2020-03-12: The GPG key for the SaltStack security team published on saltstack.com had expired in 2018 and a request for an updated key was sent.

- 2020-03-16: Repeated request for an updated GPG key resulted in publication of a re-signed key to the security contact page. Full vulnerability report sent to SaltStack security team.

- 2020-03-20: Request for confirmation of receipt was sent to the SaltStack security team.

- 2020-03-24: The SaltStack security team confirm receipt of the vulnerability report and that they are reviewing it.

- 2020-04-07: SaltStack asks if F-Secure has requested CVE's for the reported vulnerabilities and F-Secure replies in the negative, recommending that SaltStack proceed to contact Mitre in order to reserve CVE IDs.

- 2020-04-15: F-Secure informs SaltStack that an Internet wide scan turned up over 6,000 publicly exposed salt masters and expresses concern that these will be at risk of compromise when the vulnerabilities are disclosed. Also, F-Secure requests information on SaltStacks plan for distributing fixes.

- 2020-04-16: SaltStack informs F-Secure that fixes are being tested and planned for release "early next week". F-Secure reiterates concerns over the number of salt masters exposed to the public Internet and requests information about how SaltStack plans to communicate this release to their customers.

- 2020-04-18: SaltStack informs F-Secure of their communication plans and requests the list of identified IP-addresses that expose a salt master to the public Internet as well as suggestions for alternative approaches to disclosure.

- 2020-04-20: F-Secure provides IP-address list and suggests a multiple-phase communication strategy as an alternative.

- 2020-04-23: SaltStack publishes advance notice to their users urging them not to expose salt masters to the Internet and to prepare to apply the patch once it is published on the 29th (https://github.com/saltstack/community/blob/master/doc/Community-Message.pdf).

- 2020-04-27: F-Secure requests information from SaltStack about the CVE ids allocated for the vulnerabilities.

- 2020-04-29: F-Secure reiterates request for CVE ids to SaltStack.
- 2020-04-29: SaltStack responds with the allocated CVE identifiers.
- 2020-04-29: SaltStack publishes version 3000.2 and 2019.2.4 addressing these issues.
- 2020-04-30: F-Secure publishes advisory.

Immediately upon publishing their security advisory, F-Secure was asked what they expected to see next. They said that they expect to see attacks in the wild very shortly:

> "We expect that any competent hacker will be able to create 100 percent reliable exploits for these issues in under 24 hours," the researchers said, citing the "reliability and simplicity" of exploitation."

And, sure enough, hackers wasted no time exploiting vulnerable Salt instances used in various infrastructures for server management and automation. Among the organizations that announced an intrusion are LineageOS, Vates (creators of open source Xen Orchestra), Ghost blogging platform, and even my very favorite certificate authority, DigiCert. By now hundreds of servers, both masters and clients (minions), if not thousands, have likely been compromised by now.

Because exploit code is trivial to create, F-Secure published nothing in order to protect companies that are slow to patch. However, several versions and proofs of concept have since emerged in public... So it's likely that more attacks are to come.

PoC's...
https://github.com/dozernz/cve-2020-11651
https://github.com/jasperla/CVE-2020-11651-poc
https://github.com/Imanfeng/SaltStack-Exp
https://github.com/0xc0d/CVE-2020-11651

In the few days since the attacks began more than 134 messages have been posted to Salt's bug page: https://github.com/saltstack/salt/issues/57057

On Sunday, Jeremy Rowley, DigiCert's Executive VP of Product posted the news to the Certificate Transparency group at Google:

> Hey all,
>
> I'm sad to report that we discovered today that CT Log 2's key used to sign SCTs was compromised last night at 7 pm via the Salt vulnerability (https://threatpost.com/salt-bugs-full-rce-root-cloud-servers/155383/). All other DigiCert CT logs are unaffected as they run on separate infrastructure. We are pulling the log into read-only mode right now. Although we don't think the key was used to sign SCTs (the attacker doesn't seem to realize that they gained access to the keys and were running other services on the infrastructure), any SCTs provided from that log after 7pm MST yesterday are suspect. The log should be pulled from the trusted log list.
>
> Happy to answer any questions about what happened, the infrastructure running the other logs, or what remediation we are taking..
>
> Jeremy

Adobe's Big Last Tuesday, Non-Patch Tuesday, Update
Last Tuesday Adobe released emergency updates for three of its widely used products that patch dozens of newly discovered critical vulnerabilities. The affected software was Adobe Illustrator, Adobe Bridge, and the Magento e-commerce platform, containing a total of 35 vulnerabilities where each one of them is affected with multiple critical arbitrary code execution flaws. Yikes.

It's unclear to me how Adobe Illustrator 2020, a quite capable drawing tool used by millions of artists around the world, could contain 5 critical remote code execution vulnerabilities. But these days, but no one is ever content to leave anything alone. It's a race to add features. So I wouldn't be surprised if it uses UPnP to open up an incoming port from the Internet. Who knows? What I do know is that anyone using Illustrator, Bridge or Magneto would be well advised to update their products immediately.

Another welcome service coming from Mozilla...
Remember how we have "Send" from Mozilla. It's a painless, limitless, locally encrypted TNO large file transfer facility that unregistered users can use for file transfers up to 1 Gigabyte and registered users can use for file transfers up to 2.5 Gig. It's available for free at: https://send.firefox.com/

Right. So, Mozilla will be addressing another constant source of annoyance with Firefox Private Relay, a one-click eMail alias creation service.

When Apple announced a similar forthcoming service as part of their "Sign in with Apple" service at the 2019 WWDC I thought that it was a cool idea, and I was a bit envious since it wasn't clear how that would help someone using mail over on Windows. But Mozilla's forthcoming Firefox extension may be just the ticket.

The extension will generate unique aliases on-the-fly whenever you need an address but really don't want to give out your actual primary address. I know that we all maintain throwaway accounts for that, but this is better since it offers a nice built-in e-mail alias UI to manage aliases.

The forthcoming service entered testing last month and is currently in a closed beta, with a public beta currently scheduled for later this year. Mozilla explains: "We will forward emails from the alias to your real inbox. If any alias starts to receive emails you don't want, you can disable it or delete it completely."

Political Correctness hits cybersecurity
Leo, though for the most part I feel young and wonderful, I think that I must actually be getting old. Since things are seeming increasingly weird.
https://www.ncsc.gov.uk/blog-post/terminology-its-not-black-and-white

Get this: The UK government's cyber-security agency wrote last week that it would stop using the terms “whitelist” and “blacklist” due to stigma and racial stereotyping surrounding the two terms. Instead, the UK National Cyber Security Centre said that going forward, it would use the terms “allow list” and “deny list” instead of the other two.

Emma W., who heads up Advice and Guidance at the NCSC said:

> *"It's fairly common to say whitelisting and blacklisting to describe desirable and undesirable things in cyber security. However, there's an issue with the terminology. It only makes sense if you equate white with 'good, permitted, safe' and black with 'bad, dangerous, forbidden'. There are some obvious problems with this. So, in the name of helping to stamp out racism in cyber security, we will avoid this casually pejorative wording on our website in the future."*

So, I'm confused. Does that mean we can no longer have whitehat and blackhat hackers? Do they need to become "AllowList" hackers and "DenyList" hackers?

When we say we'll add those IPs to the blacklist, no one in the world imagines that this has anything to do with the African American race -- or anything to do with race. Maybe I'm being insensitive, but this makes me wonder if it's not possible to be too sensitive?

Google has announced its impending clean-up of the Chrome Web Store

It's way past time to do some serious house cleaning. Last Wednesday, Google announced new rules for the Chrome Web Store which should cut down the number of shady Chrome extensions submitted and listed there. Google explained that due to Chrome's success as today's top browser platform -- as we know, by a wide margin -- the Chrome Web Store has seen an influx of spammers and fraudsters.

Google says that these malicious entities have been behind a rising number of duplicate, spammy, and purely malicious extensions that are now poisoning and drowning the Chrome Web Store in low-quality content. So... starting on Thursday, August 27 (plenty of time from now), Google will begin enforcing a welcome new set of rules, which will result in a large number of extensions being immediately delisted. "Buh Bye!" These rules are meant to crack down on a series of practices extension developers have been recently employing to flood the Web Store with shady extensions or boost install counts for low-quality content. They include:

- Developers cannot submit duplicate extensions, for example wallpaper extensions that have different names but provide the user with the same wallpapers when installed.
- Extensions are not allowed to use "keyword spam" techniques to flood metadata fields with multiple terms and have the extension listed across multiple categories to improve the extension's visibility in search results.
- Developers are not allowed to use misleading, improperly formatted, non-descriptive, irrelevant, excessive, or inappropriate metadata. Extension metadata needs to be accurate, and Google intends to be strict about it.
- Developers are now forbidden from inflating product ratings, reviews, or install counts by illegitimate means, such as fraudulent or paid downloads, reviews, and ratings.
- Extensions that have only one purpose, such as launching a web page or an app will no longer be allowed.
- Extensions that abuse browser notifications to spam users with ads or other messages will also be banned.

On that Thursday (August 27), Google says it intends to take down every extension that violates these new rules. And once that happens, many thousands of junk Chrome extensions will

disappear from the Web Store. This will make searching for useful content on the site easier and safer than it has recently become. The Chrome Web Store currently lists more than 200,000 extensions.

Warning about RDP is not crying wolf

Just as the Shodan Internet-wide application search engine saw a 41% upward jump in RDP endpoints appearing at the beginning of last month, Kaspersky has had their eye on the Internet and has seen a somewhat startling increase in RDP scanning and brute force (what we're now calling "credential stuffing") attacks:

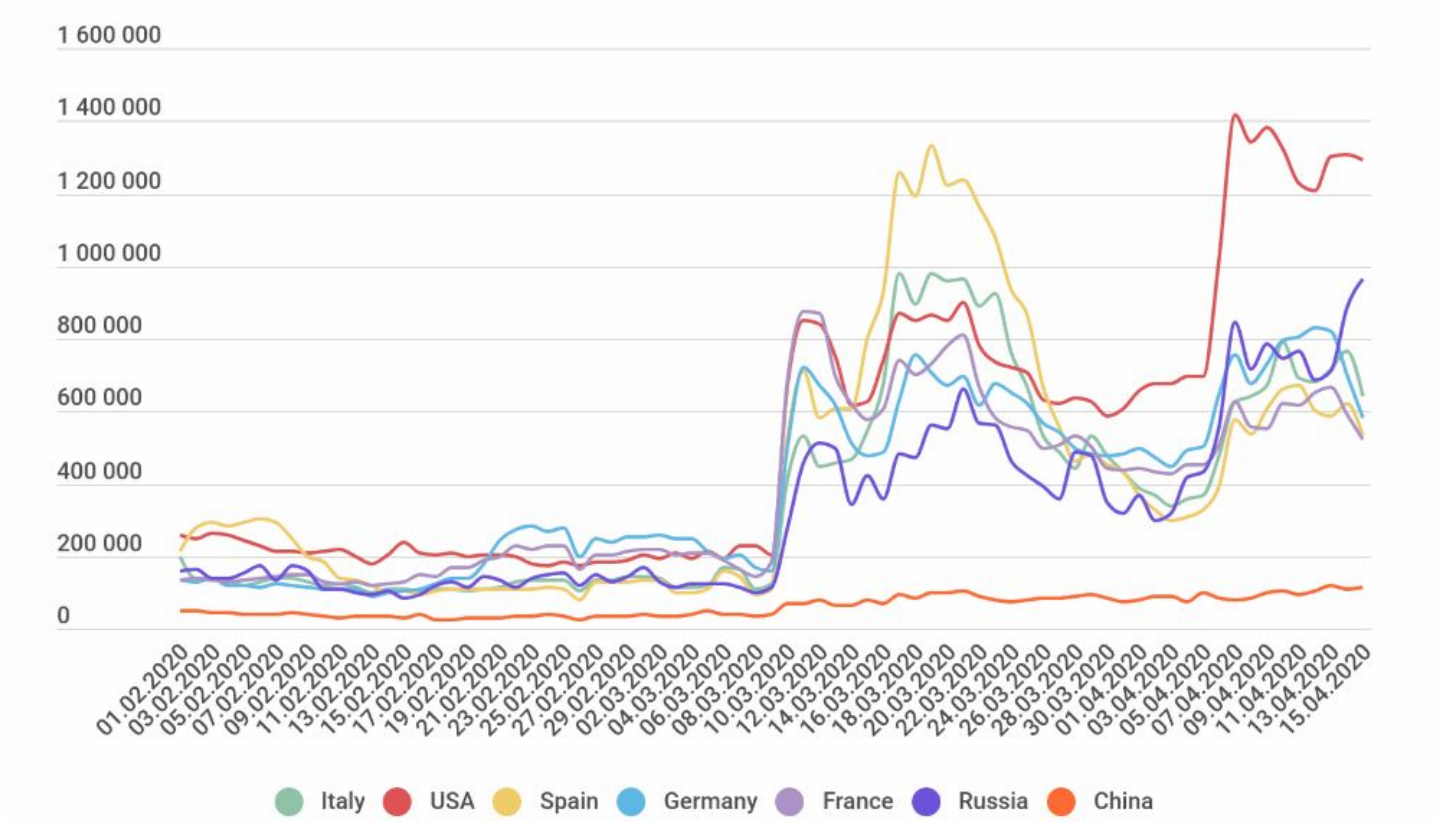

Whereas from the beginning of the year scans for open RDP ports were purring along between 150,000 and 200,000 per day. That all changed around the beginning of March. Many countries saw a quadrupling in attack rate and at one point toward the end of March, Spain hit about 1.2 MILLION attacks per day. Later, the USA even exceeded that by crossing the 1.4 Million attacks per day threshold. Although, overall, there has been some ebb and flow, securing RDP from random Internet access is crucial.

Errata

Recall that last week we talked about RPKI and BGP and I noted, with apologies, my utter inability to properly pronounce the name of one of the major players behind the work to secure BGP. His name is spelled: “Job Snijders” and I still have no idea how HE pronounces it. But I do know that he was aware of my failed attempts. He tweeted: “Hey @SGgrc - thanks for the shout out! You did a great butcher job of my name haha ;-) let me know if you want to talk more about RPKI & BGP in your show!”

So I don’t know if he’s a regular listener or if, more likely, someone who is, shot him a note saying “Oh, boy, did Gibson ever mess up your name!”

In any event, the entire industry VERY MUCH appreciates the vital work he is doing, however his name is pronounced!

Miscellany

- **“Devs”**
  - Produced by FX, streaming on Hulu / Purchase from Amazon for $13.
  - 6hrs:49mins, 8-episode miniseries / Amazingly GOOD. / Sound track!
- **Tales from the Loop**
  - Amazon Prime: Not awful, but not great.

SpinRite

I told everyone last week that I hoped to be able to report some first actual results from the hardware level AHCI driver development... and I can. We now have drive enumeration and communication with every system our testers have thrown at the code. We’re now working with Intel, AMD, ASMedia, Samsung and Marvell AHCI chipsets -- which is every chipset we have encountered. The newer Intel chipsets initially caused a glitch because they handled hardware interrupts a bit differently from everyone else. And the Marvell chips were just brought online yesterday morning before I began working to assemble this podcast. The trick with the Marvell chips was that they fail to implement two status bits that are clearly part of the AHCI specification. So I've allowed support for those two status bits to be optional.

At this point, the new AHCI driver code is running on everything that all of our testers have thrown at it. So, while I'm waiting for someone to dig up something that it doesn't work with, I'm going to re-enable the support for all of the older interface standards: IDE, ATA and ACHI when it’s in Legacy mode. We should then have enumeration and visibility of every drive on every system... all without ANY BIOS involvement.

People love Benchmarks. GRC's DNS Bench benchmark is our most often downloaded freeware. It's been downloaded 5,620,883 times and is currently being downloaded at the rate of a little over 4,000 times per day... every day. People love Benchmarks.

Since I want to verify our ability to transfer data, and since it'll be fun to know how long the next SpinRite will take to perform a problem-finding intolerant read scan of our big drives, I'm going to then turn it into a nice DriveBench benchmark which will allow everyone to test and compare their drive's true performance. I

And, since I would like this used and tested far and wide to verify the compatibility of this new code with everyone's hardware, I'll be producing an easy-to-boot solution so that everyone listening to this podcast can take their drives out for a test drive if they wish.

So, anyway... lots of good progress on SpinRite is underway.

An Authoritarian Internet?

The Internet gets more explicitly political. China proposes a wholesale revamp which other authoritarian countries including Iran, Russia, and Saudi Arabia support.

There's quite a lot of heated rhetoric surrounding this topic, and as with contact tracing where "it's complicated" the rhetoric and the technology might be quite different. So in working to get some feel for what was actually going on I started with a powerpoint slide presentation which is mostly about the technological bullet points. No one could implement anything from this, but it's at exactly the right level to gain some overall understanding.

https://www.itu.int/en/ITU-T/Workshops-and-Seminars/2019101416/Documents/Sheng_Jiang_Presentation.pdf

The presentation was made to a workshop/seminar of the ITU, the International Telecommunications Union.

Key network technology requirements:

- Need more flexible variable-length IP addresses to adapt to diverse scenarios and backwards compatible with IPv4, IPv6
- Supports addressing and routing optimized based on communication entity semantics and digital twin relationships. (Some time is spent talking about these "Digital Twins" but what they are is never made very clear through the slides. Perhaps we had to be there.)
- Support real-time large-flux communication in virtual-physical fusion scenarios combining ubiquitous AI theory. (There is a lot of talk of the future need for truly massive bandwidth and guaranteed low-latency in support of holographic transmission and applications such as tele-surgery.)
- Support secure, reliable and resilient connection among massive heterogeneous networks
- Future network architectures supporting multi-ID space and digital twin relationships

"Identity-based Routing" is one of the most interesting aspects of this and I think it's what frightens people. The slides explain: "Instead of mapping all information into network addresses, diverse IDs are used to indicate the destination, which improves routing capabilities." In other words, this system proposes that it would be possible to address a connection to an individual based upon some sort of identity designator and the network, necessarily knowing where this individual is, would route the connection to them.

The slides explain: "Packets are addressed by semantic content metadata, of which old style IP addresses become a subclass." Examples of semantic content are:

- Content ID
- Device ID
- People ID
- Service ID

They suggest that this would enable network layer deterministic forwarding rather than our current "best effort" forwarding.

And the overall global network would be redesigned with a sort of super-QOS (quality of service) structure which, they assert, will be needed to satisfy future scenarios. So the network's available transit bandwidth would be divided into "slices" and they suggest examples:

- AR/VR slice has low latency <20ms
- Self driving slice (I guess autos) has latency <5ms
- Teleprotection slice offers jitter of <50us.

A lot of this is then pulled together on a slide titled "End-to-End Communication Requirements for Intrinsic Security"

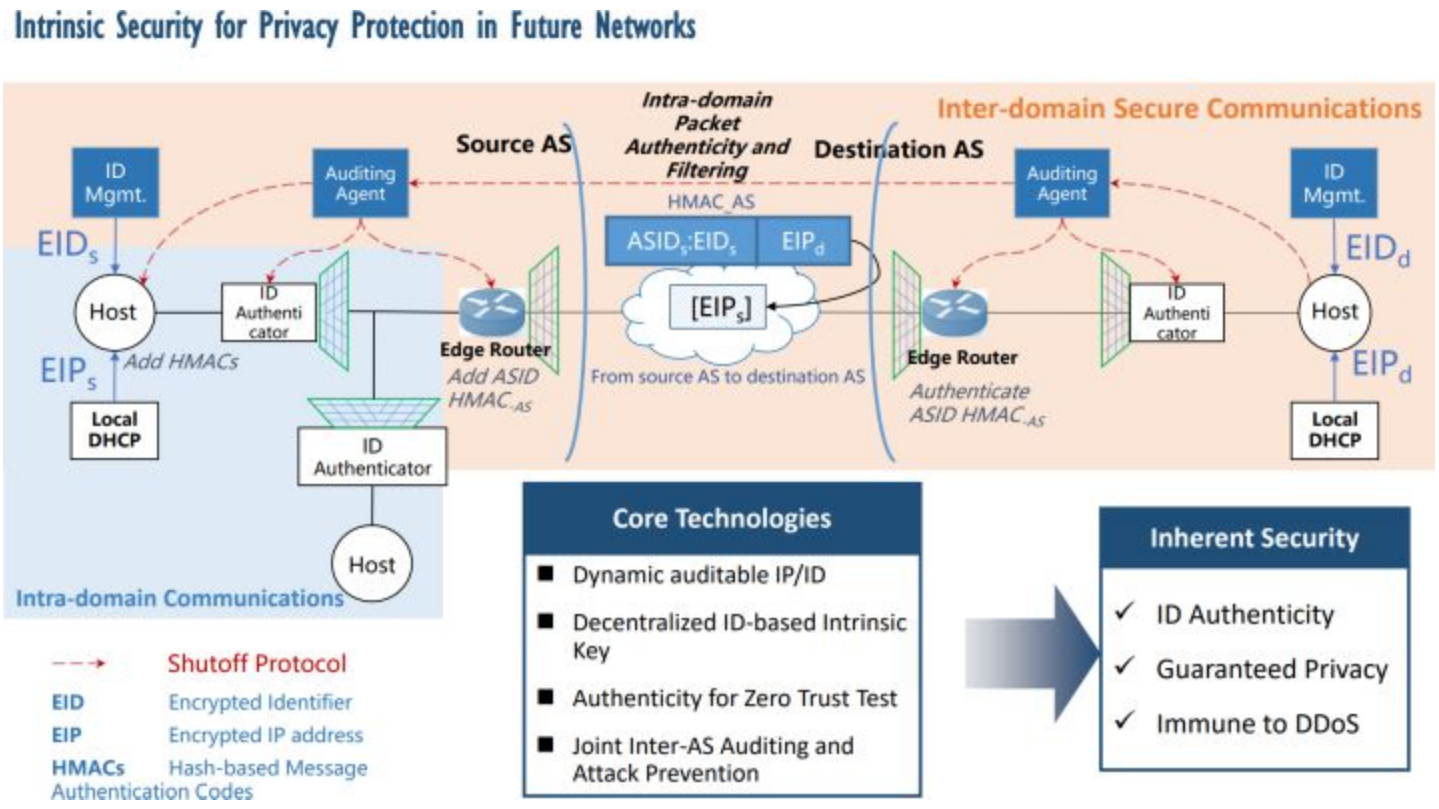

Authenticity: More than 1.3rd of autonomous domains do not have prevention mechanisms for IP address spoofing.

Accountability vs Privacy: Tradeoff → Exposing IP, port and TTL on the wire decreases privacy but anonymizing them decreases accountability. IP headers design should take into account a tradeoff between accountability and privacy.

Confidentiality & Integrity: The current key exchange mechanism has many vulnerabilities. Intrinsic Identity keys and not relying upon 3rd parties.

Availability: Avoid the unavailability of target network resources, computing resources, storage resources, etc. caused by DDoS attacks. It is necessary to combine authenticity and accountability to build a multi-level verification filtering system for inter-domain and intra-domain traffic.

So, the system encrypts the sender's ID and the sender's Local IP. They are authenticated along with the hosting network. The traffic moves to the Autonomous System's edge router which adds an ASID (autonomous system ID) HMAC. The traffic traveling over the Internet contains the ASID verifier, an encrypted ID and an encrypted IP.

As the packet attempts to enter the destination autonomous system, the sender's ASID HMAC is authenticated to permit it to enter, then its content is made available and the destination address pieces are decrypted.

In the flowchart diagram, an overriding "Auditing Agent" appears across the top with dotted red lines flowing down to all of the key elements, and in a box with a legend describing some of the abbreviations the dotted red line is labeled "Shutoff Protocol." Apparently, if the auditing agent is unhappy, that traffic does not flow.

There's also ample attention paid to the needs imposed by "holographic" communication which appears to literally mean the display of holograms. Under that discussion we have:

• Ultra-high Throughput
Along with the evolution of media technologies, the future applications, especially the holographic communication, potentially require ultra-high throughput to the network

• Customizable Priority and Strategy
The priority and requirement of application data is different. Besides choosing the transport layer protocol, application should own the capability to indicate transport strategy.

• Reduced Complexity and Indeterminacy
Lossy transmission affects the quality of content however retransmission (lossless) potentially decreases the throughput. The new transport should consider combining with new technologies, such as network coding, to deal with the packet loss and provide better end-to-end capability.

• Inherent Network-awareness
Besides packet loss, more parameters, such as bandwidth, queue, delay and jitter, will influence the transport strategy. New transport should be network-aware.

So now we have some broad overview sense for what the Chinese government is proposing to the world. And as I noted at the start, most of the world - certainly the democratically organized portions of the world - have not accepted this proposal with open arms.

RIPE, which is the EU's Internet governance body, blogged about this proposal under the questioning title: "Do We Need a New IP?" That full blog post is rather long and winding, but the official response is short and quite to the point. It's titled: Response to "New IP, Shaping Future Network" proposal, and it reads:

https://www.ripe.net/participate/internet-governance/multi-stakeholder-engagement/ripe-ncc_tsag_new-ip.pdf

During this group's last meeting in September 2019, a number of proponents introduced a proposal titled "New IP, Shaping Future Network", proposing what was described as an opportunity for a "strategic transformation [of the Internet]". The RIPE NCC appreciates the opportunity to respond to the proposal by means of this contribution.

Since first being described in 1974, Internet protocol architecture has totally transformed our societies and economies, as the design philosophy of an open and flexible Internet that has allowed for an unprecedented number of life-changing innovations.

It is true that, boosted by privatisation and increased competition, these technological breakthroughs have also dramatically changed the telecommunications landscape. While it was once common to carry TCP/IP data streams across traditional PSTN infrastructure, those original telephone services have now also evolved and are carried using the Internet protocol.

Throughout its lifetime, the Internet protocol has also adapted to accommodate changing requirements and new technological insights. The most noticeable and impactful change was the redesign towards IPv6, with its 128 bits address space to overcome scalability issues and support growth beyond expectations.

We believe it is exactly that open and adaptable nature, not only of the technical architecture but also the surrounding governance models, that is fundamental to the Internet's unprecedented success. Growing alongside its technical infrastructure, the Internet's governance model has also evolved to include new stakeholders and accommodate innovations, unforeseen use cases and unimagined growth as the Internet became the fundamental technology powering our societies and economies.

The open, inclusive, multistakeholder approach throughout its development – both in the technical forums that created its standards, as well as in the governance of its resources – has made the Internet what it is today.

While we recognise the need for both the technical standards as well as these governance models to continue evolving, we strongly believe such evolution should take place from within the organisations and structures that invented the Internet and have supported its evolution throughout its history. We also strongly believe that any rationale for change must be carefully evaluated by all stakeholders in an open and transparent process to achieve consensus.

The RIPE NCC is deeply concerned by what has been proposed here. We are especially concerned by the notion that this proposal represents an opportunity to steer away from the traditional "bottom-up" decision-making model.

We also believe the technical rationale presented is flawed and find the suggested alternative designs to be both unrealistic and unproven. Furthermore, if any of the proposed solutions could be developed to a mature and production-ready standard, market adoption is very uncertain and will take decades to accomplish.

The RIPE NCC is of the opinion that the proposal is premature and that following through with any of the suggested work would create significant overlap with the ongoing work of other stakeholders – in particular, that of the Internet Engineering Task Force.

Although some of the issues mentioned may warrant further study, we insist that any work on the evolution of the IP protocol layers and the associated technical standards be left to the Internet Engineering Task Force and be conducted under its governance.

The RIPE NCC recommends that TSAG, and the ITU Telecommunication Standardization Sector in general, not make any of the suggested changes to its structure or pursue any work items related to this proposal that would evolve the Internet protocol stack under the ITU's remit.

So this is interesting from a sort of theoretical "what if" model, and it sure did ruffle a bunch of feathers. But it's clear that those of us who have some deep affection for the operation of today's Internet, warts and all, have little to fear. Let's remember that we can't even get people to upgrade their version of TLS let alone scrap their ENTIRE investment in ALL of their existing Internet networking equipment.

If such a thing ever should happen it's a long long ways off. And the only way I can see anything like that EVER happening would be as an overlay of an entirely new next-generation network which would run in parallel alongside today's Internet and would gradually, over time, eventually replace the older Internet… which is exactly the way the original Internet was born.