Security Now! #839 - 10-05-21

“Something Went Wrong”

This week on Security Now!

This week we, of course, look at the massive global outage that took down all Facebook services for 6 hours yesterday. But before we get there we look at this week's new pair of 0-day flaws which Google fixed in Chrome, we note the arrival of Windows 11 with a yawn and also caution about one known flaw that it's already known to have. We look at some potential for global action against ransomware, and some possible movement by the FCC to thwart SIM swapping and number transporting attacks. We also examine a widespread Android Trojan which is making its attackers far too much money, and speaking of money, there's a known flaw in Apple Pay when using a VISA card that neither company wants to fix. And finally, after a quick check-in on SpinRite, we're going to examine what exactly did “go wrong” at Facebook yesterday?

Earth’s Submarine Cable Network

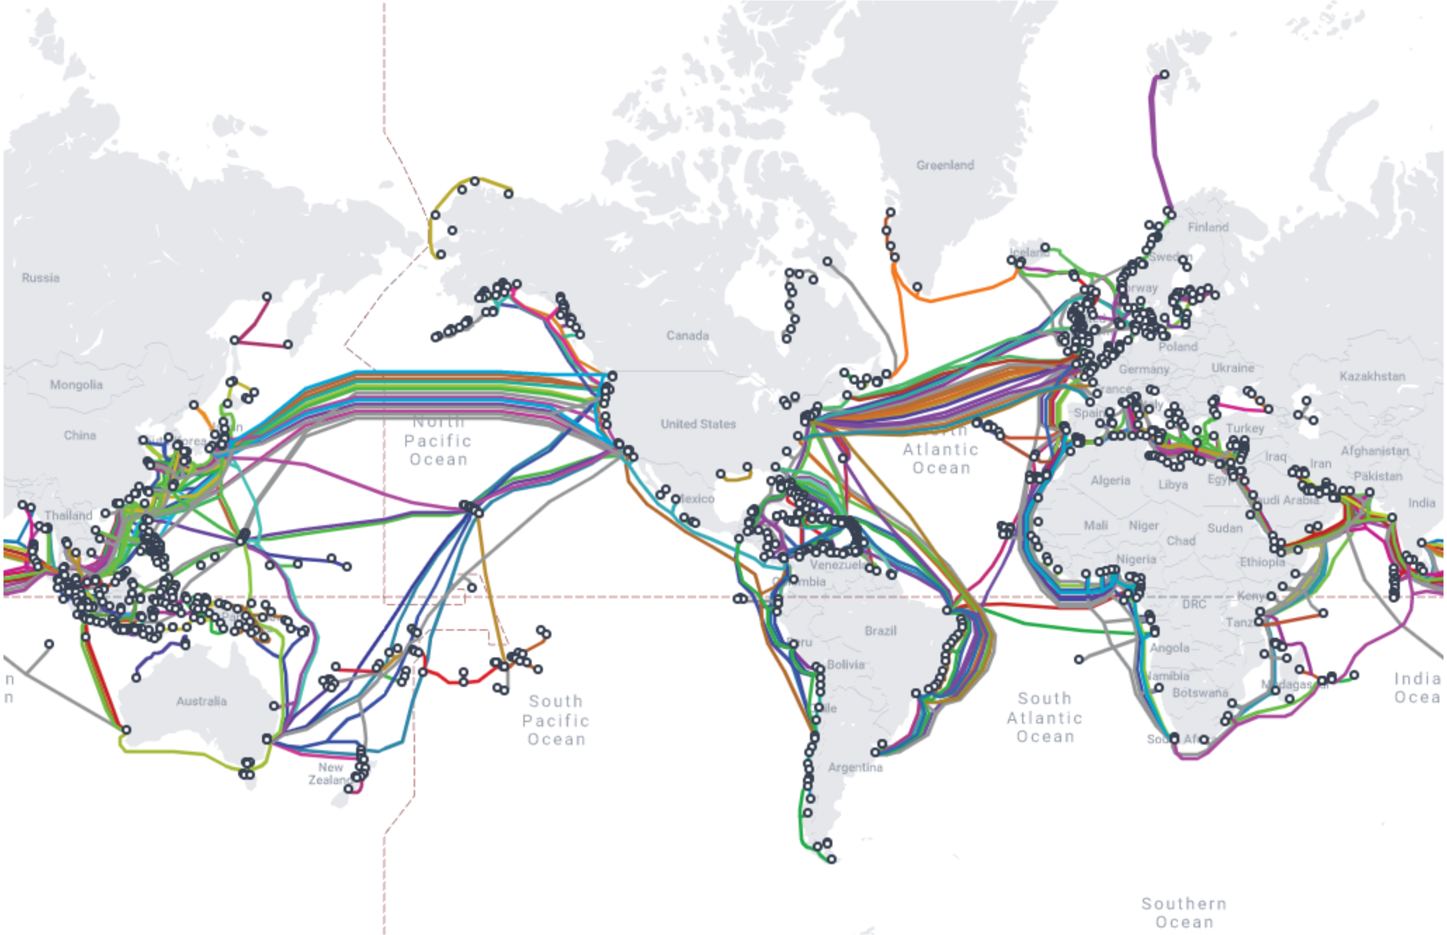

https://www.submarinecablemap.com/

0-Day Watch

Another two, in-the-wild, true 0-days found and fixed in Chrome

Chrome's most recent update late last week resolved yet another pair of authentic 0-day vulnerabilities which were found being exploited in the wild. For those keeping score at home this brings us to **14** so far this year as we finish with month 9 and start into month 10.

When I went to Chrome to ask it to update itself, I received the message: "An error occurred while checking for updates: Update check failed to start (error code 3: 0x80004002 -- system level)." The message contained a link for learning more, which I then clicked and was informed that Error 3 or 11 was: "Update failed" (that much I knew) but then also: "An error occurred while checking for updates: Update server not available" Now, as it happens, this was coincident with Monday afternoon's transient global Facebook outage. I didn't know whether one might have anything to do with the other. But I was stuck back on v94.0.4606.61 from the previous week. (Facebook later returned to the Internet, but Chrome's update server remains offline.)

In any event, issues CVE-2021-37975 and 976 were found and immediately put to rest. They were fixed by a total of four patches surrounding a use-after-free flaw in the V8 JavaScript and WebAssembly engine as well as an information leak from Chrome's core. The discovery and report of the 975 flaw was credited to an anonymous researcher, whereas the 976 problem was found by the same researcher within Google's TAG team who also discovered the actively exploited use-after-free flaw in Chrome's Portals API that we talked about last week.

Eventually, Google's update server will be back online and everyone will be able to update from v94.0.4606.61 to 94.0.4606.71. I'd feel pretty confident that no one should treat this as an emergency. These incremental updates typically aren't. But Google has arranged to make them so transparent to its users that there's no reason for them to be delayed.

Google probably takes a bit more heat over these than they deserve, since whereas this does bring the total to 14 0-day vulnerabilities so far this year, Microsoft fixes scores of such problems every month, and even then, only after they decide to stop ignoring the many problems in the first place.

Security News

Windows 11 arrives

Yesterday, thanks to a tweet from Paul Thurrott, we have the official Windows 11 download page. So I tweeted:

> *"Anyone seeking additional pain can now obtain the official Windows 11, directly from Microsoft. (I would call it the "Final Windows 11" ... but who are we kidding?)"*

We never had a final Windows 10. And now we definitively know that they never intended Win10 to be the last Windows ever. So, every few years Microsoft will find something new to do to the UI that allows them to increment Windows' major version number, and to enforce some new random and arbitrary restriction on which machines it will run on. Because I'm a pragmatist, I'm pretty sure that Microsoft will never lose me as a user. Windows is far more practical for me than any alternative. But they have, and they continue to, needlessly discard bits and pieces of my respect for no good reason.

To make finding Windows 11 easy, I created the shortest practical and easy-to-remember grc.sc shortcut possible: https://grc.sc/11
https://www.microsoft.com/en-us/software-download/windows11

A known memory leak in Windows Explorer

https://www.pcgamer.com/windows-11-file-explorer-memory-leak-bug/
And speaking of this release certainly not being final for Windows 11, PC Gamer notes that "Windows 11 will soon be rolling out" and rhetorically asks: "but will File Explorer keep chomping through your RAM?"

It seems that among the many known new bugs which Windows 11 is introducing is a big new memory leak in Windows File Explorer. PC Gamer informs us that, <quote>:

> *The Windows 11 File Explorer memory leak bug, which surfaced a couple of months ago thanks to the keen eyes of one Windows 11 Insider preview user, is outlined in a post by user gyrohan269 on the Windows 11 subreddit. They note that with each instance of Windows File Explorer opened, the RAM usage stacks and doesn't disperse upon exiting. The post was met with thousands of upvotes from those with a similar issue, and plenty of comments from users who were able to replicate the issue.*
>
> *We're currently running version 22000.194 (which we believe is the release version) on our test bench, and I was able to reliably replicate the bug several times. And, no, the RAM usage still hadn't freed up even after about half an hour of waiting.*
>
> *If you want to check if this is a problem for your own Windows 11 version, open Task Manager now and sort your processes by highest memory usage, then spam Win+E. You'll notice Explorer rise up the list pretty fast, and once you've closed them all keep an eye out to see if the memory frees up.*
>
> *There doesn't seem to be an official acknowledgement of the issue anywhere, let alone news of a coming fix. Thankfully it has been logged in the Feedback Hub so, if you can replicate it, do pile in so Microsoft are aware it's a wide-ranging issue.*

It sounds as though this would mostly be a problem for someone who launches and shuts down Windows Explorer over and over without rebooting, thus compounding and accumulating the memory problem. I always run with a single instance of Explorer up and open on a monitor off to the right. I'll occasionally open another Explorer instance, but not often. And I assume that this will be fixed before long. So just something for early adopters of Windows 11 to be aware of.

Ransomware and cyber warfare.
On the ransomware front, the US announced last Friday that the administration will be conducting a series of online virtual meetings with representatives from 30 countries, including NATO allies and G7 partners, on the topic of cybercrime, with an explicit focus on ransomware and the concomitant abuse of cryptocurrency. In their press release on the topic, the White House said: "This month, the United States will bring together 30 countries to accelerate our cooperation in combating cybercrime, improving law enforcement collaboration, stemming the illicit use of cryptocurrency, and engaging on these issues diplomatically."

Accordingly to the release, additional topics to be discussed will include 5G technology, supply chain attacks, quantum computing, and AI.

You poke the bear enough and it rouses. As we know, last May was the attack against Colonial Pipeline which resulted in fuel shortages across the US East Coast. The next month, in June, the attack on JBS Foods disturbed the supply of meat across the US. Don't mess with the bear's meat! And then in July the massive series of attacks which leveraged flaws in Kaseya's IT management servers created disruptions at hundreds of companies across the US and more than 1,500 across the world.

We're told that President Biden first raised the issue of ransomware attacks carried out from within Russia's borders with President Putin during a face-to-face meeting last June. And that he again raised the issue in a phone call in early July, asking Mr. Putin to crack down on gangs operating within Russia. And we know now that if that was done — what actually transpired remains unknown — it apparently didn't last long, since multiple attacks resumed last month. And this led the FBI to conclude that they saw <quote> "no indication" that Russian officials had taken any effort to crack down on these groups.

There are a number of big, significant and very interesting macro trends taking shape. What governments ultimately decide to do about encryption they cannot crack is another one. So it's going to be very interesting to see how this plays out. We actually do appear to be entering a period of true inter-nation economic cyber warfare. That idea still startles and unsettles me. I don't want it to be true.

On the topic of thwarting SIM swapping attacks...
I titled our August 31st podcast "Life: Hanging by a PIN", and I know from the much greater than average level of feedback I received, that my painting a clear picture of just how vulnerable we would typically be if our smartphone number were to fall into the hands of an attacker, that a greater awareness of the trouble did hit home with our listeners. So I was interested to see, initially hopeful, and I wanted all of our listeners to know, that at least some of the US bureaucracy is awake to this threat and is beginning to move on it. The bad news is that the specific arm of the US bureaucracy... is the FCC.

Last Thursday the FCC — our Federal Communications Commission — announced its plans to introduce new rules for US mobile carriers to govern any changes made to their subscriber's telephone numbers in an attempt to address the growing problem of SIM swapping and port-out fraud attacks.

As we know, these attacks surround mobile carriers' failure to correctly verify the requesting party's identity when that party requests either that their service be transferred to a new SIM card, or to an account at another mobile operator. Podcast #834 painted a very clear and depressing picture of just how much devastation could result.

The US Justice Department has charged many individuals over the past few years with theft enabled by SIM swapping and port-out fraud. And some of the victims of these thefts have, understandably, brought lawsuits against their mobile carriers in an attempt to recover their monetary losses. Many of those lawsuits are still working their way through the US's delay-prone legal system.

In response to the problems, some US carriers have introduced additional verification measures to limit this form of fraud, but the SIM swapping gangs have also changed their tactics. Some groups have bribed carrier employees or used vulnerabilities in the carrier’s networked backend systems to perform their attacks, thus skipping the need to have any direct contact and “trick” the carrier’s support staff.

So, in its press release Thursday, the FCC announced that in response to having received “numerous complaints from consumers'' (I'll bet!) that it's initiating a “formal rulemaking process” by issuing a “notice of proposed rulemaking.” Oh my god. These are the same people who said they were going to outlaw and prevent telephone spam. So I think this means that it's still going to be up to us to protect ourselves.

An FCC spokesperson told "The Record", who reported on this, that <quote> “The FCC’s rulemaking process generally starts with a Notice of Proposed Rulemaking that asks questions and makes proposals. We then... have a period... during which we take public comments – generally made through our Electronic Comment Filing System.” (Just to remind everyone, that's the ineptly designed system that was brutally spammed during their Net Neutrality request for comments period.) Anyway, they conclude: “After that we review comments before taking any next steps.”

So, in conclusion, although I wrote above that some of the US bureaucracy is awake to this threat and is beginning to move on it, unfortunately, the bureaucracy in question is the US FCC. So, in other words, although I'll be keeping an eye on this for everyone, we shouldn't be holding our breath.

<<< SPONSOR BREAK HERE >>>

A widespread Android Trojan is making someone a bunch of money!
https://blog.zimperium.com/grifthorse-android-trojan-steals-millions-from-over-10-million-victims-globally/

Zimperium recently revealed their research into one of the best-named Android malware campaigns I've seen in a long time. It's a Trojan, right? So ya have a horse. And it signs its victims up to high-cost services to reap the rewards. So they named it “GriftHorse.”

What caught my eye was the tremendous amount of effort that **had** to have gone into the creation of GriftHorse's count’em **139 individually created** and infected Android Trojaned apps.

Zimperium explains:

The threat actors have exerted substantial effort to maximize their presence in the Android ecosystem through a large number of applications, developer accounts, and domains. The Zimperium zLab researchers have noticed the technique of abusing cross-platform development frameworks to stay undetected has been on the rise, making it more difficult for legacy mobile AV providers to detect and protect their customers.

The timeline of the threat group dates back to November 2020, suggesting that their patience and persistence will probably not come to an end with the closing down of this campaign. The threat to Android users will always be present, considering the innovative approaches used by malicious actors to infect the victims.

The numerical stats reveal that more than 10 million Android users fell victim to this campaign globally, suffering financial losses while the threat group grew wealthier and motivated with time. And while the victims struggle to get their money back, the cybercriminals made off with millions of Euros through this technically novel and effective Trojan campaign.

List of Applications

App Name	**Min**	**Max**
Handy Translator Pro	500,000	1,000,000
Heart Rate and Pulse Tracker	100,000	500,000
Geospot: GPS Location Tracker	100,000	500,000
iCare – Find Location	100,000	500,000
My Chat Translator	100,000	500,000
Bus – Metrolis 2021	100,000	500,000
Free Translator Photo	100,000	500,000
Locker Tool	100,000	500,000
Fingerprint Changer	100,000	500,000
Call Recoder Pro	100,000	500,000
Instant Speech Translation	100,000	500,000
Racers Car Driver	100,000	500,000
Slime Simulator	100,000	500,000
Keyboard Themes	100,000	500,000
What's Me Sticker	100,000	500,000
Amazing Video Editor	100,000	500,000
Heart Rhythm	100,000	500,000
Smart Spot Locator	100,000	500,000
CutCut Pro	100,000	500,000
OFFRoaders – Survive	100,000	500,000
Phone Finder by Clapping	100,000	500,000
Bus Driving Simulator	100,000	500,000
Fingerprint Defender	100,000	500,000
Lifeel – scan and test	100,000	500,000
Launcher iOS 15	100,000	500,000
Idle Gun	50,000	100,000

Scanner App Scan Docs & Notes	50,000	100,000
Chat Translator All Messengers	50,000	100,000
Hunt Contact	50,000	100,000
Icony	50,000	100,000
Horoscope : Fortune	50,000	100,000
Fitness Point	50,000	100,000
Qibla AR Pro	50,000	100,000
Heart Rate and Meal Tracker	50,000	100,000
Mine Easy Translator	50,000	100,000
PhoneControl Block Spam Calls	50,000	100,000
Parallax paper 3D	50,000	100,000
SnapLens – Photo Translator	50,000	100,000
Qibla Pass Direction	50,000	100,000
Caller-x	50,000	100,000
Clap	50,000	100,000
Photo Effect Pro	10,000	50,000
iConnected Tracker	10,000	50,000
Smart Call Recorder	10,000	50,000
Daily Horoscope & Life Palmestry	10,000	50,000
Qibla Compass (Kaaba Locator)	10,000	50,000
Prookie-Cartoon Photo Editor	10,000	50,000
Qibla Ultimate	10,000	50,000
Truck – RoudDrive Offroad	10,000	50,000
GPS Phone Tracker – Family Locator	10,000	50,000
Call Recorder iCall	10,000	50,000
PikCho Editor app	10,000	50,000
Street Cars: pro Racing	10,000	50,000
Cinema Hall: Free HD Movies	10,000	50,000
Live Wallpaper & Background	10,000	50,000
Intelligent Translator Pro	10,000	50,000
Face Analyzer	10,000	50,000
*TrueCaller & TrueRecoder	10,000	50,000
iTranslator_ Text & Voice & Photo	10,000	50,000
Pulse App – Heart Rate Monitor	10,000	50,000
Video & Photo Recovery Manager 2	10,000	50,000
Fitness Trainer	10,000	50,000
ClipBuddy	10,000	50,000
Vector arts	10,000	50,000
Ludo Speak v2.0	10,000	50,000
Battery Live Wallpaper 4K	10,000	50,000
Heart Rate Pro Health Monitor	10,000	50,000
Locatoria – Find Location	10,000	50,000
GetContacter	10,000	50,000
Photo Lab	10,000	50,000
AR Phone Booster – Battery Saver	10,000	50,000
English Arabic Translator direct	10,000	50,000
VPN Zone – Fast & Easy Proxy	10,000	50,000
100% Projector for Mobile Phone	10,000	50,000

Forza H Mobile 4 Ultimate Edition	10,000	50,000
Amazing Sticky Slime Simulator	10,000	50,000
Clap To Find My Phone	10,000	50,000
Screen Mirroring TV Cast	10,000	50,000
Free Calls WorldWide	10,000	50,000
My Locator Plus	10,000	50,000
iSalam Qibla Compass	5,000	10,000
Language Translator-Easy&Fast	5,000	10,000
WiFi Unlock Password Pro X	5,000	10,000
Pony Video Chat-Live Stream	5,000	10,000
Zodiac : Hand	5,000	10,000
Ludo Game Classic	5,000	10,000
Loca – Find Location	5,000	10,000
Easy TV Show	5,000	10,000
Qibla correct Quran Coran Koran	5,000	10,000
Dating App – Sweet Meet	5,000	10,000
R Circle – Location Finder	5,000	10,000
TagsContact	5,000	10,000
Muslim Prayer Times	1,000	5,000
Qibla Compass	1,000	5,000
Soul Scanner – Check Your	1,000	5,000
CIAO – Live Video Chat	1,000	5,000
Plant Camera Identifier	1,000	5,000
Color Call Changer	1,000	5,000
Squishy and Pop it	1,000	5,000
Keyboard: Virtual Projector App	1,000	5,000
Scanner Pro App: PDF Document	1,000	5,000
QR Reader Pro	1,000	5,000
FX Keyboard	1,000	5,000
You Frame	1,000	5,000
Call Record Pro	1,000	5,000
Free Islamic Stickers 2021	1,000	5,000
QR Code Reader – Barcode Scanner	1,000	5,000
Bag X-Ray 100% Scanner	1,000	5,000
Phone Caller Screen 2021	1,000	5,000
Translate It – Online App	1,000	5,000
Mobile Things Finder	1,000	5,000
Proof-Caller	1,000	5,000
Phone Search by Clap	1,000	5,000
Second Translate PRO	1,000	5,000
CallerID	1,000	5,000
3D Camera To Plan	500	1,000
Qibla Finder – Qibla Direction	500	1,000
Stickers Maker for WhatsApp	500	1,000
Qibla direction watch (compass)	500	1,000
Piano Bot Easy Lessons	500	1,000
CallHelp: Second Phone Number	500	1,000
FastPulse – Heart Rate Monitor	500	1,000

Caller ID & Spam Blocker	500	1,000
Free Coupons 2021	100	500
KFC Saudi	100	500
Skycoach	100	500
HOO Live – Meet and Chat	100	500
Easy Bass Booster	10	50
Coupons & Gifts: InstaShop	10	50
FindContact	10	50
Launcher iOS for Android	10	50
Call Blocker-Spam Call Blocker	10	50
Call Blocker-Spam Call Blocker	10	50
Live Mobile Number Tracker	10	50
Total:	**4,287,470**	**17,345,450**

Zimperium summed up this threat by writing:

> *These mobile applications pose a threat to all Android devices by functioning as a Trojan that subscribes unsuspecting users to paid services, charging a premium amounting to around 36 Euros per month which is about $42.*
>
> *The campaign has targeted millions of users from over 70 countries by serving selective malicious pages to users based on the geo-location of their IP address with the local language. This social engineering trick is exceptionally successful, considering users might feel more comfortable sharing information to a website in their local language.*
>
> *Upon infection, the victim is bombarded with alerts on the screen letting them know they had won a prize and needed to claim it immediately. These pop ups reappear no less than five times per hour until the application user successfully accepts the offer. Upon accepting the invitation for the prize, the malware redirects the victim to a geo-specific webpage where they are asked to submit their phone numbers for verification. But in reality, they are submitting their phone number to a premium SMS service that would start charging their phone bill over €30 per month. The victim does not immediately notice the impact of the theft, and the likelihood of it continuing for months before detection is high, with little to no recourse to get one's money back.*
>
> *These cybercriminals took great care not to get caught by malware researchers by avoiding hardcoding URLs or reusing the same domains and filtering / serving the malicious payload based on the originating IP address's geolocation. This method allowed the attackers to target different countries in different ways. This check on the server-side evades dynamic analysis checking for network communication and behaviors.*
>
> *Overall, GriftHorse Android Trojan takes advantage of small screens, local trust, and misinformation to trick users into downloading and installing these Android Trojans, as well frustration or curiosity when accepting the fake free prize spammed into their notification screens.*

I wanted to put this particular campaign on our listener’s radar because, just as when we first talked about ransomware many years ago and I commented that it felt to me as though it was really going to become a problem in the future, this feels the same way to me.

The trouble is that the attackers behind this were known to be netting several million dollars per month. That creates a lot of motivation. And many of the new facets of this approach have no been proven to be surprisingly successful. And unlike a single ransomware attack which inherently creates a single point of failure for the attacker if their victim chooses not to comply, the individual damages in this attack are small, incremental and widely dispersed across 70 countries and many languages. They slip under the radar by siphoning a small cash flow from many millions of individual end users and in the process create a stable cash flow of illicit funds. The GriftHorse campaign not only managed to fly under the radar and avoid A/V detection, it likely surpassed hundreds of millions of dollars in the total amount plundered from its victims until Zimperium responsibly notified Google of their discovery and Google purged those identified apps from the Play Store. But even so, those apps continue to be available on untrusted third-party app repositories, which again underscores the risks associated with sideloading arbitrary applications.

I’m afraid that we can expect to see more of this style of attack in the future, probably on the Android platform since that’s where it makes the most sense.

There's a problem with Apple Pay and Visa

Both companies have been informed, but neither have fixed it, because they’re arguing about whose fault it is, even though either one of them could mitigate the trouble on their end.

A handful of researchers at the University of Birmingham and the University of Surrey, both in the UK, have written up their research in a paper which will participate in the 2022 IEEE Symposium on Security and Privacy. In their paper's Abstract they explain rather densely:

> *Relay attackers can forward messages between a contactless EMV bank card and a shop reader, making it possible to wirelessly pickpocket money. To protect against this, Apple Pay requires a user’s fingerprint or Face ID to authorise payments, while Mastercard and Visa have proposed protocols to stop such relay attacks. We investigate transport payment modes and find that we can build on relaying to bypass the ApplePay lock screen, and illicitly pay from a locked iPhone to any EMV reader, for any amount, without user authorisation. We show that Visa’s proposed relay-countermeasure can be bypassed using rooted smart phones. We analyse Mastercard’s relay protection, and show that its timing bounds could be more reliably imposed at the lower protocol level, rather than at the EMV protocol level. With these insights, we propose a new relay-resistance protocol (L1RP) for EMV. We use the Tamarin prover to model mobile-phone payments with and without user authentication, and in different payment modes. We formally verify solutions to our attack suggested by Apple and Visa, and used by Samsung, and we verify that our proposed protocol provides protection from relay attacks.*

So what does this mean for us practically?

Elsewhere they make that more clear, explaining:

> *Contactless Europay, Mastercard, and Visa (EMV) payments are a fast and easy way to make payments and are increasingly becoming a standard way to pay. However, if payments can be made with no user input, this increases the attack surface for adversaries and especially for relay attackers, who can ferry messages between cards and readers without the owner's knowledge, enabling fraudulent payments. Payments via smart-phone apps generally have to be confirmed by a user via a fingerprint, PIN code, or Face ID. This makes relay attacks less of a threat.*
>
> *However, Apple Pay introduced the "Express Transit/Travel" feature (May 2019) that allows Apple Pay to be used at a transport-ticketing barrier station without unlocking the phone, for usability purposes. We show that this feature can be leveraged to bypass the Apple Pay lock screen, and illicitly pay from a locked iPhone, using a Visa card, to any EMV reader, for any amount, without user authorisation.*
>
> *Furthermore, Visa has proposed a protocol to stop such relay attacks for cards. We show that Visa's proposed relay-countermeasure can be bypassed using a pair of NFC-enabled Android smart phones, of which one is rooted.*

So, the bottom line is:

- The Apple Pay lock screen can be bypassed for any iPhone with a Visa card set up in transit mode. The contactless limit can also be bypassed allowing unlimited EMV contactless transactions from a locked iPhone.

- An attacker only needs a stolen, powered on iPhone — it does NOT need to be unlocked. The transactions could also be relayed from an iPhone inside someones handbag or pocket, without their knowledge and the attacker needs no assistance from the merchant.

- Backend fraud detection checks have not stopped any of their test payments.

- Just to be clear, this attack is made possible by a combination of flaws in both Apple Pay and Visa's system. It does not, for instance, affect Mastercard on Apple Pay or Visa on Samsung Pay.

- Their research provides formal modelling that shows that either Apple or Visa could mitigate this attack on their own. As I noted above, both companies have been informed months ago, yet neither have fixed their system, so the vulnerability remains live and workable today.

- The researchers recommend that all iPhone users check that they do not have a Visa card set up in **transit mode**, and if they do... it would need to be disabled to prevent this attack from succeeding against that phone and card.

It's not that it's going to happen, but it could. Hopefully, Apple will get some flack now that this research is public and fix this. And we have another instance of Apple only fixing something when someone makes them. That's a bit disappointing.

Sci-Fi

Foundation
No improvement after the 3rd episode. It’s just not fun. It’s slow and heavy and boring. So I’m hoping that when hostile aliens invade the Earth this Friday in Apple’s new series “Invasion” that will provide some much-needed Sci-Fi excitement.

Matt Kopit / @mkopit
Regarding the sound quality on Foundation: I noticed the same thing with muddled dialogue and found that disabling Dolby Atmos in my AppleTV's audio settings improved things dramatically. YMMV, but feel free to give it a try.

SpinRite

Over the weekend I posted an update to the patient gang who's been waiting to test and pound on a next test release of what will become SpinRite v6.1. I explained that I had reached the point where, as far as I can tell, everything is working. I tracked down a problem with the non-DMA IDE driver and made the BIOS associator a bit more bulletproof. I had previously invested a bunch of time updating SpinRite's original benchmarking code for 64-bits worth of sectors. But I hate the way it's working and I've decided to scrap and rework it now... since I cannot live with it.

The trouble is that the ONLY timing reference the original SpinRite had was the PC's counter/timer which runs at 1.193 megahertz, or one count every 838 nanoseconds. The counter/timer is still there, and that's plenty of resolution. But the counter/timer chip was a ripple-counter which could sometimes be sampled and caught mid-count (or mid-ripple), thus producing a false count result. The chip had a latching function, but I encountered some that still produced bogus results. So I developed a sanity-filter that took three successive readings from the counter and only believed the final one if each of the previous two readings showed the same count or one that was increasing. (The filter also had to be smart about the 16-bit counter's wrap-around, since that happened 18.2 times per second.) If the filter thought that it had obtained a mid-count reading, it would try again.

The point is back then, that's all I had. Today's RDTSC (read timestamp counter) instruction did not exist as it does now. My new code, the code I wrote for the ReadSpeed benchmark, uses the RDTSC instruction for various delays needed when waiting for hardware to settle and for all of its performance measurements. But SpinRite is not using it for the benchmark timing since I was trying to minimize the rewriting, and I figured that what SpinRite was already doing was fine. But I had to go back into that old code to make sure it would work, and it's horrendous. So I need to rip it out and replace it with the newer code that I have already developed for the ReadSpeed benchmark. Although it doesn't need it yet, SpinRite will be needing the new system's insane picosecond resolution in the future, and v6.1's code will be surviving for years. So now is the time.

I've already worked out all of the details for ReadSpeed, and I have all of the formatting and everything ready to go. So it's just a matter of removing the old bad code and moving the newly written good code over. I'm sure that once that's done I'll be
glad that I did.

“Something Went Wrong”

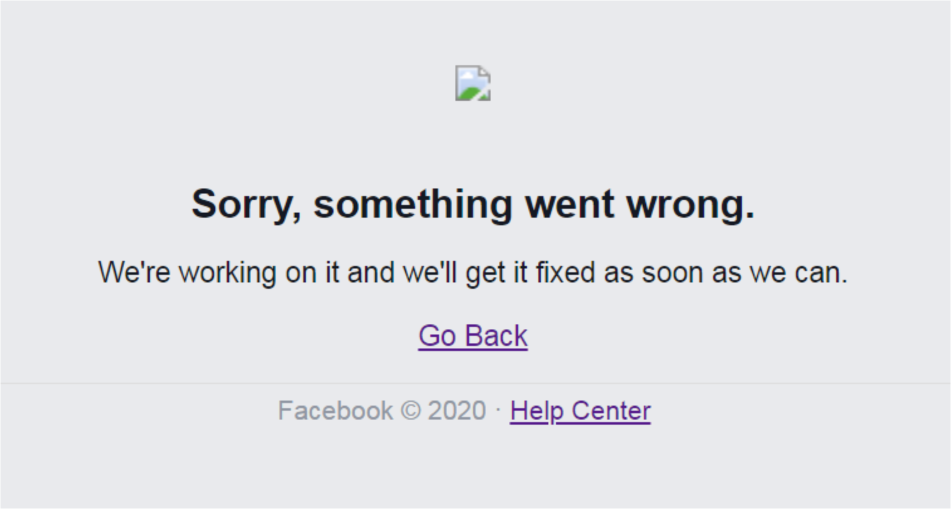

Sorry, something went wrong.

We're working on it and we'll get it fixed as soon as we can.

Go Back

Facebook © 2020 · Help Center

The Internet’s big iron routers are connected to each other by their peering interface links. If you had a piece of paper covered with a bunch of circles representing routers, then drew straight lines between each of the circles close to other circles, those straight lines would be peering links. Those peering links carry the low-level packet traffic which routers route, and the routers also use those same links to maintain persistent TCP connections over which the BGP protocol flows.

BGP, the Border Gateway Protocol is a TCP-hosted protocol like HTTP, FTP or SMTP. But in this case BGP is a peering protocol exclusively used by routers to talk to their neighbors. The conversation they have has the purpose of synchronizing their respective routing tables with each other—with the routers to which they peer—which is to say, to which they are directly connected. When a router receives a change to its own routing table, either introduced by a local admin or received from another router peer, it updates its own table appropriately to reflect the changes, then sends news of the changes it has made to its other peers so that they, too, might make any needed adjustments.

A large enterprise which has its own Autonomous System (AS) number (Facebook’s is AS32934), will use their large backbone routers to connect their public enterprise network to the rest of the public Internet. Their routers will know which IP ranges belong to its enterprise owner and it will contain entries to route traffic bound for those IP ranges to the enterprises’ router interfaces. And, since every such router shares its routing table with its peers, all of its peering routers will also know that any traffic they receive for those IPs should be forwarded to that Autonomous System router. And since **those** routers also share **their** routing tables with **their** peers, all of the routers they connect to will also know... and so it goes, over and over, peer-by-peer, until every big router on the Internet knows where to send traffic that’s bound for any of that enterprise’s IP ranges.

In the weird parlance of BGP routing, we say that that original router is “advertising” the routes which it alone is able to handle by forwarding any incoming traffic to its Internet-connected enterprise.

So yesterday, at 11:39 a.m. Eastern Time, shortly before noon (15:39 UTC), someone at Facebook updated the routing information for Facebook's networks to something that no longer worked. Even now, following Facebook's official post-recovery blog posting, they're not telling the world exactly what happened. Initial reports suggested that it might only be the routing to Facebook's four authoritative DNS servers **[a,b,c,d].ns.facebook.com** that may have been messed up to render those crucial servers unreachable. But some additional evidence suggests that the trouble was much bigger than that.

The reports of a DNS-based failure came from those whose own DNS servers suddenly started returning SERVFAIL errors, indicating that none of those four authoritative Facebook DNS servers responsible for "facebook.com" were replying to their queries. Presumably, those four authoritative Facebook DNS servers were still up, but they had just become unreachable due to a routing error. The reason I believed that DNS was only a part of a much bigger, probably network wide all-of-Facebook problem, was that, as shown in the diagram below, traffic being served by Facebook dropped like a rock off a cliff:

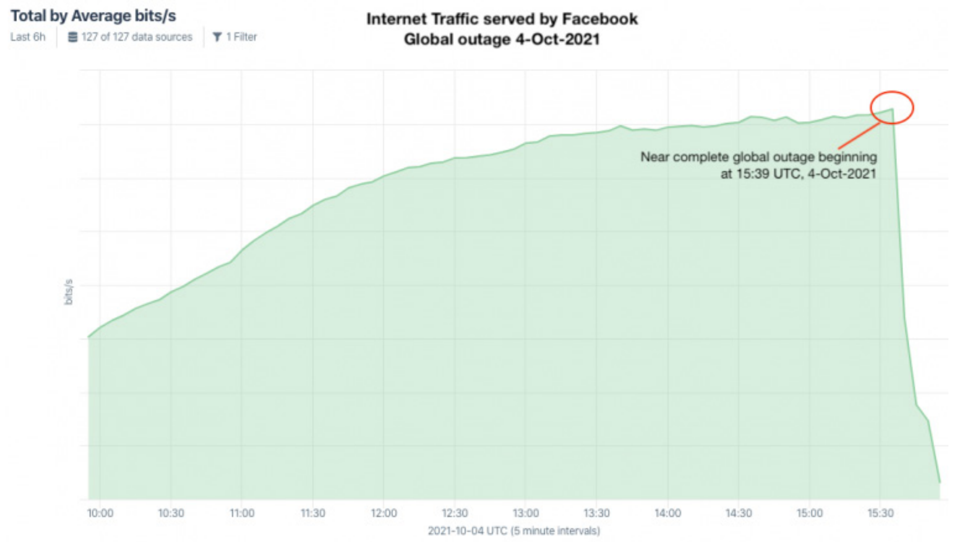

As we know, DNS caches. And caching is a core key feature of DNS. If this were "just" a major DNS outage, we would expect to see a far more gradual reduction in traffic to FaceBook over a span of hours rather than minutes as the Internet's massively dispersed and distributed DNS caches, which exist at every level, even within individual end-user smartphones and desktop PCs independently expired their own local DNS caching to only then go in search of an IP address update. And only when an expired entry could not be updated would the user's local machine report that Facebook had become unavailable. Or, as their copyrighted page wonderful stated: "Something went wrong."

Yesterday, during the outage, just like the rest of the world I was unable to query Facebook's authoritative DNS servers. So I was unable to obtain the details of how Facebook had setup their DNS. But this morning, with Facebook back on the air, we can more closely examine the way they have their DNS configured. I used NSLOOKUP to pull the SOA, the Start Of Authority, DNS record directly from one of Facebook's authoritative DNS servers: **a.ns.facebook.com**.

Here's what I received:

```
> set querytype=SOA
> facebook.com a.ns.facebook.com
Server:  a.ns.facebook.com
Addresses:  2a03:2880:f0fc:c:face:b00c:0:35
          129.134.30.12

facebook.com
        primary name server = a.ns.facebook.com
        responsible mail addr = dns.facebook.com
        serial  = 3954289570
        refresh = 14400 (4 hours)
        retry   = 1800 (30 mins)
        expire  = 604800 (7 days)
        default TTL = 300 (5 mins)
```

The last two items tell the story: The record expiration time and the default TTL.

Facebook's default TTL for their DNS records is only 5 minutes. This means that, unless it's been overridden by a per-record TTL (and Facebook's A and AAAA records for its IPv4 and IPv6 IPs are not overridden), once every 5 minutes, every DNS client on the planet will see that its local cached copy of Facebook's IP address has reached its end of life and should be re-fetched from that client's configured DNS server. But if that update query fails, what happens? That's where the SOA's specified record expiration time keeps you from being SOL. Although Facebook uses a short TTL, they also use an expiration of 7 days. And that will be seven days from the most recent previous successful update, which would have been within the past five minutes.

This means that even if Facebook's DNS had dropped off the face of the Earth, so long as the rest of Facebook was still present, client on the planet would have happily continued using Facebook for the next week while their DNS patiently and periodically attempted to check-in for a DNS update.

In other words, all of the evidence points to this being a massive whole-of-Facebook Internet routing error. Add to that the fact that Facebook's Instagram also disappeared from the Internet, despite the fact that Instagram's DNS is hosted by Amazon, not Facebook.

DNS caches. But what doesn't cache is BGP. If someone at Facebook made the colossal mistake of deleting all of Facebook's routes from the Internet, that change would have propagated at the speed of the Internet and within a few minutes — just as we saw from that sheer Internet traffic cliff — no Internet backbone routers anywhere in the world would have had any idea what to do with a Facebook IP address... even though all of those users would have still had them cached for the next week.

Since BGP is just a protocol like any other, anyone with access to Autonomous System level routers which peer with others, or is on the inside of Inter-Network Operations, can monitor BPG traffic and activity directly. And Monday afternoon Cloudflare was doing just that:

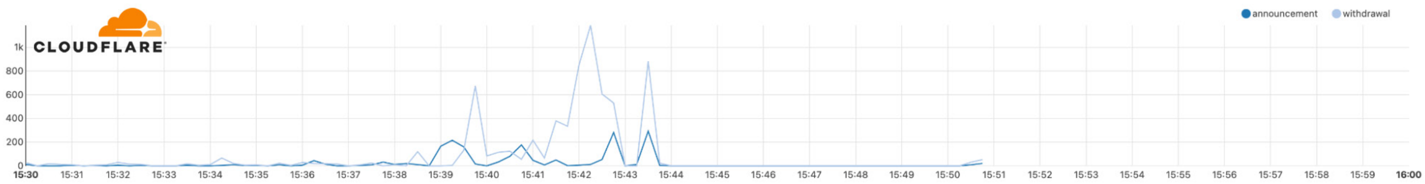

The chart above is a 30-minute graph, taken from 15:30 to 16:00 UTC, of Facebook-related BPG traffic. Normally, there's nothing happening since routing tends to be boring — except when it's not! And there was nothing boring about routing yesterday. The chart shows that starting at exactly 15:39, in the darker blue, some new routes were announced and a bit later there was a similar volume of routes withdrawn. It may have been that that was deliberate, because it looks like things were being moved around. But then, about a minute and a half later, we see a massive flurry of routes being withdrawn in the lighter blue, where the area under the curve dwarfs that of any new replacement announcements. Maybe it was supposed to be that way. The goal may have been to consolidate routes, which is much better for routing efficiency. In that case, you **would** expect to see many more withdrawals than new announcements. Except that we also know that the result of this sudden flurry of activity was a catastrophe.

By now, Facebook certainly knows precisely what happened. But based upon their public statement once this had all been resolved it doesn't appear as though we're going to get much clarity from them, unless it eventually leaks out. But what they did say confirms our hypothesis from the evidence. They posted:

> *To all the people and businesses around the world who depend on us, we are sorry for the inconvenience caused by today's outage across our platforms. We've been working as hard as we can to restore access, and our systems are now back up and running. The underlying cause of this outage also impacted many of the internal tools and systems we use in our day-to-day operations, complicating our attempts to quickly diagnose and resolve the problem.*
>
> *Our engineering teams have learned that configuration changes on the backbone routers that coordinate network traffic between our data centers caused issues that interrupted this communication. This disruption to network traffic had a cascading effect on the way our data centers communicate, bringing our services to a halt.*
>
> *Our services are now back online and we're actively working to fully return them to regular operations. We want to make clear at this time we believe the root cause of this outage was a faulty configuration change. We also have no evidence that user data was compromised as a result of this downtime.*
>
> *People and businesses around the world rely on us everyday to stay connected. We understand the impact outages like these have on people's lives, and our responsibility to keep people informed about disruptions to our services. We apologize to all those affected, and we're working to understand more about what happened today so we can continue to make our infrastructure more resilient.*

What was most curious—and really sort of unbelievable while this was all happening—was that Facebook didn't quickly pop back up on the air. If someone entered a bad routing update, how difficult could it be to at least roll back the change? As it happens, while this was going on, I was already at work on today's podcast, so I was able to participate in the drama with my followers.

My first Tweet read:

> *Facebook may have "deplatformed" itself, along with Instagram and WhatsApp. Hope no one depends upon "Login with Facebook!" Whoopsie! Somehow, the BGP entries for Facebook's DNS resolvers have been withdrawn from the Internet's routing tables. Insider? Attack? Who knows. Wow.*

After a bit more digging around I added:

> *Someone on the Facebook recovery effort has explained that a routine BGP update went wrong, which in turn locked out those with remote access who could reverse the mistake. Those who do have physical access do not have authorization on the servers. Catch-22.*

So, we were beginning to get some clarification of the trouble. Apparently, the unexpected loss of Internet routing had cut off their own engineers from the routers they needed to access in order to rollback the changes.

Whoopsie!

I know this is actually a very real concern. I manage GRC's network at Level3 remotely, and I sometimes need to alter the operation of the equipment that manages GRC's border with the Internet. My own management traffic is crossing the same boundary as the one I'm managing, and I WANT my own management traffic to be crossing that same boundary, since it's a security boundary and my management traffic needs to be highly secure. But that also means that I need to be careful not to edit some rule that locks me out. And it has happened, which necessitated a drive over to the physical plant to logon to the machine directly and correct a mistake.

A New York Times reporter Tweeted:

> *Was just on the phone with someone who works for Facebook who described employees unable to enter buildings this morning to begin to evaluate extent of outage because their badges weren't working to access doors.*

So, apparently everything at Facebook is IoT, on the Internet, and dependent upon some super-secure access control server somewhere... which cannot be reached during the routing outage.

After some additional digging, my next Tweet was:

> *Reports are that Facebook employees cannot enter their headquarters because their badges don't work, and those inside are unable to enter various rooms because access is dependent upon obtaining authorization from remote Facebook servers.*
>
> *Those who live by technology...*

And in light of all the recent news about Facebook's internal awareness that an obsession with Instagram postings may not be good for many of the youngsters who have flocked to it, I added a tongue-in-cheek Tweet:

> *Meanwhile... there's been a noted global decrease in reports of teenage depression and poor*

self image. Mental health is on the rise. But fear not, BGP is sure to be restored soon. :-/

And finally, to top it all off, five and a half years ago, back in April of 2016, Facebook acquired the domain registrar "RegistrarSec.com". Ever since then, they've been their own domain registrar... because... why not?

But what happens then, when Facebook's own domain registrar also goes offline due to a routing outage? Thanks to the automated systems being used by some less clued-in registrars, which are continually searching for previously registered domains that appear to be expired, abandoned or recently vacated... yesterday, and I kid you not, several different registrars were listing the domain **Facebook.com** as being available for sale!

I grabbed a screenshot of one of those several for the show notes. It announced with an exclamation point that "facebook.com" is for sale!

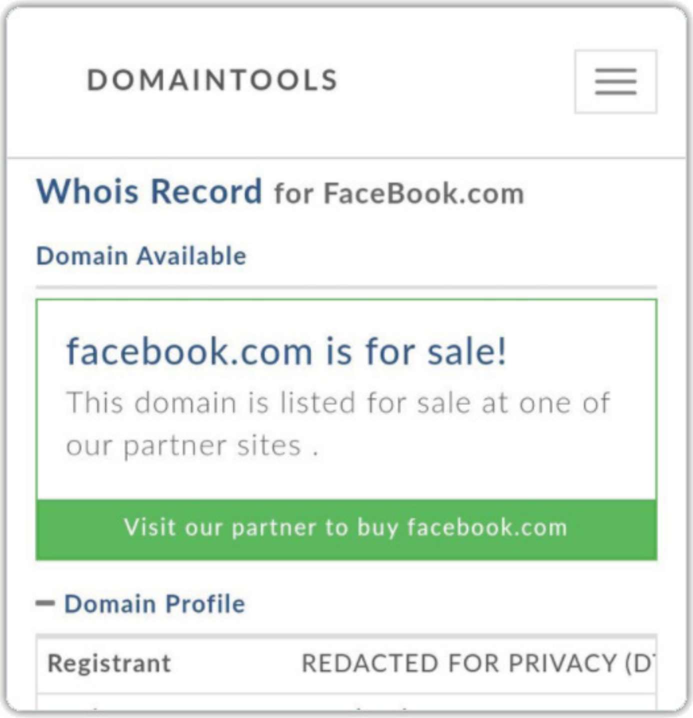

Whois Record for FaceBook.com

Domain Available

facebook.com is for sale!

This domain is listed for sale at one of our partner sites .

— Domain Profile

Registrant	REDACTED FOR PRIVACY (D

"Something Went Wrong" . . . Indeed.