

Security Now! #820 - 05-25-21

The Dark Escrow

This week on Security Now!

This week we examine Firefox's just-released and welcome re-architecture under code name "Fission." We look at a new and recently active ransomware player named "Conti" and at a recently paid, high profile megaransom. We then ask the question, "when they say IoT do they mean us?" and we examine the implications of a new industry term "Mean Time To Inventory." We'll then lighten things up a bit with a new form of CAPTCHA and, of all things, a screensaver I discovered that I cannot take my eyes off of. (Leo, it's not quite as bad as whatever that game is that you cannot stop playing, but still...) We'll then share an ample helping of Closing-The-Loop feedback from our terrific listeners, after which I want to conclude by predicting what I would bet we're probably going to next see emerge from the evolving ransomware business model — sad though it is to utter the phrase "ransomware business model."

On the theme of "being careful what you ask for"...



Another bit of brilliance from Scott Adams

Web Browser News

Firefox finally achieves sustained "Fission"

"Fission" is Mozilla's name for full site isolation. The original design of all web browsers was as a single application running in an operating system. Or, a bit more formally, a single process being hosted by the OS. The trouble was, that browsers have kept growing more and more complex to the point that the line is blurring between an OS app and what a browser can do with a web page. The other day I followed a link to a drag and drop mechanical electrical circuit design page where a capacitor was emulated by a balloon, an inductor was created from a flywheel, a resistor was a narrowing pipe, and so on. I initially took the page for granted, until it hit me that this amazing fully animated creation was not an app in the classic sense, but that this was being entirely done on a web page. So, yes, today's web browsers have become incredibly complex and capable.

And as we know, complexity is the enemy of security. It's not impossible to have both, but it turns out that it's impractically expensive to actually have both. So we get features that mostly work. And as for security, we hope for the best while we fix the flaws that are later uncovered on-the-fly after the fact.

But browsers are sort of a special case, because the presence of the near certainty of flaws raises the specter of containment. If flaws are inevitable, then what we want is to at least contain them. The last thing we want is for a flaw to be leveraged to reach across web browser pages, or really across domains, to infect an unrelated page to leverage its access to its site—for example, where we're logged into our banking site. A perfect analogy is to my insistence upon applying network segmentation for IoT devices which are inherently high risk. If we cannot guarantee that our light switches and blow dryers will not attack us, at least we can contain any potential damage by placing all of those less trustworthy devices onto their own low-trust network. It's a form of sandboxing. And this exactly mimics the designs that browsers have been working to adopt for several years.

It's taken years because switching from a mono-process architecture to a process-per-domain is far more easily said than done. Google initially released an experimental site isolation feature that they had already been hard at work on for quite a while, in Chrome 64 back at the end of 2017 and it became generally available by May of 2018. Mozilla realized that was where the future would be, so they began work on the same thing the month before its general availability in Chrome, in April of 2018. They formally announced their plans to do it nearly a year later in February of 2019 under the codename "Fission", and a little more than two years later, last week, it is finally here. But it's not yet turned on by default.

Anyone with a current release of Firefox can enable full site isolation by placing "about:config" into Firefox's URL. Then search for ""fission" and locate "fission.autostart", which should be set to "True". After making the changes, Firefox won't auto-prompt for a restart as Chrome does. But the "about:profiles" page contains a manual restart button that will do the trick.

Once that's done, Firefox will create a separate standalone OS process for each domain the browser pulls content from in order to prevent anything nasty from escaping. Even an off-site frame will now be run from within its own process and then merged onto its containing page.

So what's essentially been done by these two browsers is to redefine the division of labor between them and the underlying OS. All contemporary operating systems already have extremely mature enforcement of inter-process isolation which browsers historically took advantage of only to keep web foolishness contained within the browser and away from all other external applications and the larger OS. But browsers struggled to keep sites isolated and contained. By rewriting browsers under this process-per-domain model, they're able to leverage the OS's already mature process isolation to deliver very strong site isolation.


Ransomware

Conti. We have very thoroughly covered the Darkside ransomware as a service gang. So next we need to introduce "Conti"...

The Conti is believed to be used by a Russia-based if not backed cybercrime group known as "Wizard Spider." They use phishing attacks to install the TrickBot and BazarLoader trojans that subsequently provide remote access to the infected machines. Using that remote access they then move laterally through internal networks, stealing credentials and harvesting unencrypted data stored on workstations and servers. Once the attackers have stolen everything of value and gained access to Windows domain credentials, they wait for for a quite time, like very early Sunday morning local time, before triggering and deploying the encrypting Conti ransomware throughout the network.

The Wizard Spider / Conti gang then use the stolen data as leverage to force their victims into paying a ransom by threatening to release it on their ransom data leak site if they are not paid. Recent high-profile Conti ransomware attacks include the FreePBX developer Sangoma, IoT chip maker Advantech, Broward County Public Schools (BCPS), and the Scottish Environment Protection Agency (SEPA).

The big attack that recently captured everyone's attention was their infiltration, data exfiltration and subsequent encryption of Ireland's health service network, known as the HSE for Health Service Executive. It's Ireland's publicly funded healthcare system which was forced to shut down the Friday before last. The Irish national health service said: "We have taken the precaution of shutting down all our IT systems in order to protect them from this attack and to allow us fully assess the situation with our own security partners." That IT outage led to widespread disruption of the country's healthcare, resulting in limited access to diagnostics and medical records, transcription errors due to handwritten notes, and slow response times for healthcare visits.



As you already know, we infiltrated your network and stayed in it for more than 2 weeks(enough to study all your documentation), encrypted your file servers, sql servers, downloaded all important information with a total weight of more than 700 GB: personal data of patients(home addresses, phone numbers of the contract), employees (home addresses, employment contracts, scans of personal documents, phone numbers), contracts, customer bases, consolidated financial statements, payroll, settlements with partners, bank statements.
The good news is that we are businessmen. We want to receive ransom for everything that needs to be kept secret, and don't want to ruin your business
The amount at which we are ready to meet you and keep everything as collateral is \$ 19,999,000.

3 hours ago

Sensing that they had hooked a big fish, the attackers then demanded one thousand dollars shy of \$20 million ransom. The Prime Minister of Ireland, said that they would not be paying any ransom.

Then, as part of their ongoing negotiations, and perhaps in part due to the long shadow cast by the Darkside attack which shutdown the Colonial Pipeline, last Thursday the gang behind Conti posted a link to a "free" decryptor which will work for all of the HSE-encrypted data. Researchers have confirmed that the provided tool will decrypt the files, though it is still being inspected for any other malicious content.

However, the bad guys still insist upon being paid just shy of \$20 million ransom, threatening to go public with their trove of extremely confidential Irish citizen health records data. They said: *"We are providing the decryption tool for your network for free. But you should understand that we will sell or publish a lot of private data if you will not connect us and try to resolve the situation."*

And then, in an odd seeming bit of whimsy, the High Court of Ireland issued an injunction against the Conti Ransomware gang, demanding that the 700 GB of stolen HSE data be returned and not sold or published. What?!?! The High Court has issued an injunction against a hidden, probably Russian, cybercrime organization? To what end?

The injunction was received by the HSE against the Conti ransomware again from the High Court of Ireland. But without any formal method to service the Court's order, government representatives uploaded it to the Tor dark website associated with the gang. The order prohibits the attackers from publishing, selling, or sharing any of the stolen data with the public. And then, get this... it also demands that they return the stolen data and identify themselves by revealing their full and true names, email addresses, and physical addresses.

Which led me to wonder what sorts of mind-altering substances might be in use by the High Court? This feels like something more than a long afternoon in an Irish pub.

What's up with this apparent loss of sanity? No one believes for a moment that the Conti gang will acquiesce to the demands of the High Court of Ireland. But there's some outside hope that the country providing cover and domicile for the attackers might be willing to track them down and at least prevent them from leaking the stolen data. You have to imagine that Russian Intelligence knows exactly where these guys are holed up.

But the next day, last Friday, the Irish Times Security & Crime Editor, Conor Lally, explained in a pair of tweets that the injunction is never intended to prevent the Conti gang from leaking the data. Rather, it was issued to prevent the press, or anyone else, from publishing the contents of stolen data if it were to subsequently be leaked by the ransomware gang.

Conor first Tweeted:

The injunction is not a super injunction in the traditional sense. We can report about it and report if data is published. The Injunction is designed to stop/limit the distribution of the data/docs in Ireland after they are published by the attackers

— Conor Lally (@conormlally) May 21, 2021

And:

The injunction doesn't stop the media reporting if the attackers leak the data, which one assumes is likely. It just means media, and anyone else, cannot publish/share the actual documents when they are leaked.

— Conor Lally (@conormlally) May 21, 2021

CNA Financial pays up big ... (and insurance co's begin backing away)

I also think that it's worth taking note when a mega ransom is paid, since such events are certainly noticed by and affect the thinking of the dark denizens.

To that end, we have the news that the U.S. insurance giant CNA Financial reportedly paid \$40 million to a ransomware gang to recover access to its systems following an attack in March. This registers as one of the largest ransoms paid to date.

This was first reported by Bloomberg, citing "people with knowledge of the attack." The adversary that staged the intrusion is said to have allegedly demanded \$60 million a week after CNA, based in Chicago, began negotiations with the hackers, which culminated in the payment two weeks following the theft of company data.

In a May 12th statement, CNA Financial said it had "no evidence to indicate that external customers were potentially at risk of infection due to the incident."

The attack has been attributed to a new player on the scene, the "Phoenix CryptoLocker," according to a BleepingComputer's report at the time. The strain is believed to be an offshoot of WastedLocker and Hades, both which are known to have been used by the Russian cybercrime network, Evil Corp.

A year and a half ago, back in December 2019, U.S. authorities sanctioned Evil Corp. and filed charges against its alleged leaders Maksim Yakubets and Igor Turashev for developing and distributing the Dridex banking Trojan to plunder more than \$100 million over a period of 10 years. Law enforcement agencies also announced a reward of up to \$5 million for information leading to their arrest. Today, both individuals remain at large.

And remember that last October, 2020 the U.S. Treasury Department issued a guidance, warning of penalties against companies making ransom payments to any sanctioned person or group. This prompted ransomware negotiation firms to avoid dealing with blocked groups such as Evil Corp. The Treasury Department said: "Companies that facilitate ransomware payments to cyber actors on behalf of victims, including financial institutions, cyber insurance firms, and companies involved in digital forensics and incident response, not only encourage future ransomware payment demands but also may risk violating [Office of Foreign Assets Control] regulations."

The surge in ransomware attacks has also had an impact on the cyber insurance industry. One firm, AXA, announced earlier this month that it will stop reimbursing clients in France should they opt to make any extortion payments to ransomware cartels. This underscores the expense of ransomware indemnification where insurance firms grapple with successfully underwriting ransomware policies while confronting the rising payout costs that threaten their bottom lines.

And to that end, a report just released by the U.S. Government Accountability Office (GAO) Thursday revealed that the soaring demand for cyber insurance has driven insurers to raise premiums while limiting coverage. The amount of total direct premiums written jumped up 50% from 2016 to 2019, from \$2.1 billion to \$3.1 billion. The GAO noted that “The continually increasing frequency and severity of cyberattacks, especially ransomware attacks, have led insurers to reduce cyber coverage limits for certain riskier industry sectors, such as health care and education, and for public entities and to add specific limits on ransomware coverage.”

We've been watching this unfold on this podcast. It was inevitable... And now it's happened.

Security News

When they say IoT do they mean us?

A headline at ThreatPost from last Wednesday caught my eye. It read: “Keksec Cybergang Debuts Simps Botnet for Gaming DDoS.” with the subhead: “The newly discovered malware infects IoT devices in tandem with the prolific Gafgyt botnet, using known security vulnerabilities.”

So I thought... IoT devices? What IoT devices, exactly? Are they some obscure widget that no one we know uses? Turn out, no. They're referring to the incredible number of Linux-based NAT routers that virtually everyone is using.

The ThreatPost piece goes on to talk about a recently developed botnet named “Simps” which has emerged from the cyber-underground for the purpose of carrying out DDoS attacks aimed at gaming targets and others. It's hosted on other people's consumer routers and forms part of the toolset being used by the Keksec cybercrime group. The Simps botnet was first spotted in April being dropped on IoT devices by the Gafgyt botnet — I suppose one good botnet deserves another. Gafgyt is a Linux-based botnet that was first seen seven years ago in 2014. It targets vulnerable IoT devices such as routers made by Huawei, Realtek, ASUS and Dasa GPON home gateway devices. In other words, not some obscure unsupported lightbulb. And in the present infection campaign, Gafgyt is compromising Realtek and Linksys endpoints, and then fetches and installs the Simps bot using WGET. Simps itself then uses Mirai and Gafgyt modules for its DDoS functionality.

So our takeaway is this: These routers are only being discovered because they are, in some way, responding to incoming packets. Way back in 1999, on October 8th — my god, nearly 22 years ago — some guy named Paul Thurrott wrote the very first article for a site called WUGNET (Windows User Group Network) titled “Protect your Windows PC with Steve Gibson's 'Shield's UP!’”

The case I made then, when I was the first to use the term “stealth” to refer to an Internet-connected device that did not respond in any way to incoming probes, was that it was worth deliberately violating a de facto rule of the Internet that all devices having TCP/IP stacks must respond to a “Ping”, and that closed ports should respond with either a TCP RST or an ICMP Port Unreachable. That, nearly 22 year old advice, has aged well. It has withstood the test of time.

So, when they refer to IoT devices they DO mean us.

At my other location I have an ASUS router. It's not on the front line. It's safely positioned behind a FreeBSD pfSense router since I need features such as pfSense's powerful static port translation and IP-based incoming packet filtering. But after a Tweet I received this morning I'm excited to get home to update my ASUS by hand... because it may be the last time I need to do so. I'll have more to say about that when we get to this week's Listener Feedback.

In the meantime, please please please make absolutely certain that the routers you're responsible for, and the routers of those you care about, do not have any connection-accepting ports statically exposed to the Internet. Any appearance of convenience is just not worth the risk.

"Mean Time to Inventory"

Everyone is familiar with the abbreviation MTBF: Mean Time Before Failure. Now the industry is coining a new abbreviation: MTTI — Mean Time To Inventory. This refers to the startling speed with which bad guys have begun to scan the Internet for vulnerabilities after that vulnerability's first public announcement. In this case the term "Inventory" refers to them adding penetrated devices to **their** "inventories."

In both the cases of MTBF and MTTI, we'd like them to be as long as possible. But at least in the case of this new MTTI, it turns out that a study recently released by Palo Alto Networks' Cortex Expanse research team reveals for the first time just how startlingly short today's MTTI actually is. They frame their research by explaining:

Malicious actors are opportunistic predators, constantly searching for vulnerable targets. Unfortunately, adversaries are much faster at finding vulnerable assets to attack than defenders are at finding those same assets to secure. It's not just an arms race in terms of conducting cyberattacks and protecting against them. There's also a sprint to detect systems vulnerable to cyberthreats.

To help enterprises gain ground, the Palo Alto Networks Cortex® Xpanse™ research team studied the public-facing internet attack surface of some of the world's largest businesses. From January to March 2021, we monitored scans of 50 million IP addresses associated with 50 global enterprises, including a subset of the Fortune 500, to understand how quickly adversaries can identify vulnerable systems for exploitation.

In this report, we share our key findings, information on the top threats in attack surface management, and insights on how to ensure your organization is secure.

We've been talking about this race-to-patch vs race-to-penetrate for a while now. But we've been lacking in metrics. What the Cortex Expanse team found was that potentially juicy 0-day vulnerabilities can prompt attackers to begin scanning within as little as 15 minutes following public disclosure.

Now, think about that. This is not just a scan for a port. This is a scan for a specific vulnerability located at a specific port. That means that attackers need to write and deploy custom code within 15 minutes in order to be scanning for not-yet-patched vulnerabilities. This really does change the landscape for serious remotely exploitable vulnerabilities.

And get this, the researchers noted that attackers worked even faster when it came to Microsoft Exchange, with first vulnerability scans detected within no more than five minutes.

Recall that I talked about how it seemed clear that the bad guys must already have mature databases indexed by port and probably also by what's known about what's answering queries to that port. So the moment a new vulnerability appears for instance in Exchange Server, they're able to immediately pull the list of all known Exchange servers currently accepting connections over the SMTP, POP and IMAP ports. Today, this literally happens in the blink of an eye. When this happened to Microsoft Exchange, the researchers at F-Secure commented that vulnerable servers were "being hacked faster than we can count."

It's also been noted that the general availability of inexpensive cloud services has helped not only well-established APT groups, but also smaller cybercriminal groups and individuals to take advantage of new vulnerabilities as they surface. The Cortex Expanse report notes: "Computing has become so inexpensive that a would-be attacker need only spend about \$10 to rent cloud computing power to do an imprecise scan of the entire internet for vulnerable systems. We know from the surge in successful attacks that adversaries are regularly winning races to [infect systems before they can be patched against] new vulnerabilities."

The research also highlights Remote Desktop Protocol (RDP) as the most common vector for security intrusions among enterprise networks. It alone accounts for 32% of all security problems. (Gee... who'd have thunk?) The report says: "This is troubling because RDP can provide direct admin access to servers, making it one of the most common gateways for ransomware attacks."

You know... if it weren't for the blessed pervasive ubiquity of NAT routers, behind which all of us with Windows systems are able to hide, the world as we know it today would have already ceased to exist. Can you imagine if the world's inventory of remotely accessible devices weren't just enterprises but were also every single last powered-on Windows machine and IoT device?

The nutty IP purists, with their heads well positioned far up their own you-know-whats where the sun don't shine, have always decried the use of NAT. They say that the Internet was designed for every device to be directly addressable and accessible to every other. Thank god that never happened. Just because every IPv6 user will be receiving their own personal 64 thousand IPv6 address space, don't ever consider directly mapping those external IPs through to your devices inside. It's already bad enough that Microsoft gave us UPnP so that Xboxes could autonomously solicit incoming traffic. The last thing we need is to step out from behind the protection of those billions of little hardware firewalls that everyone is using today. With a "Mean Time To Inventory" numbered in the low minutes, none of us would stand a chance.

Miscellany

The "Doom" CAPTCHA

This week's GRC shortcut is the "DOOM CAPTCHA" —> <https://grc.sc/820>
(<https://vivirenremoto.github.io/doomcaptcha/>)

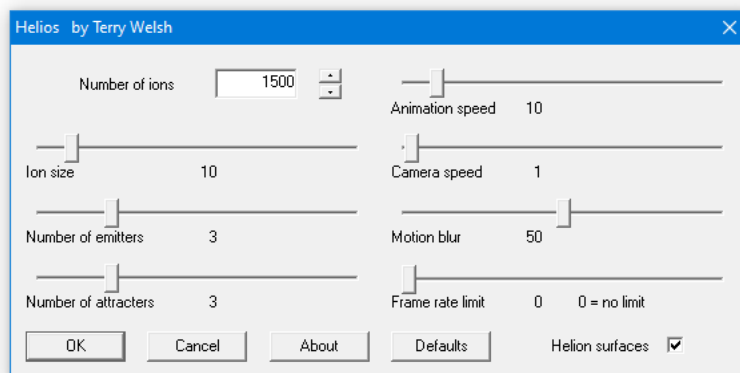
It's just a joke, but it made #1 Product of the Day over on Product Hunt:
<https://www.producthunt.com/posts/doom-captcha>

The “Helios” screensaver

This is completely random, I know. But a few months ago I looked at Windows 10’s built-in screen savers and realized that as far as Microsoft appears to be concerned, screensavers have fallen into disfavor. Or perhaps they’ve migrated to the Windows Store. I didn’t go there. Since I often leave my Win10 machine on and unattended, during dinner, taking a walk after dinner or whatever, and since my workstation is in our family room so that I’m able to work while still being nominally present (clanky mechanical keyboard notwithstanding), I decided that I wanted something fun on the screen. So I went looking for a satisfying screensaver.

Several months later, I am enjoying what I found so much that I decided I needed to share it with my listeners who often tell me that my taste matches theirs. A developer named Terry Welsh has written a collection of open source OpenGL screensavers for Windows, many of which have also been ported by others to OS X and Linux, including **HELIOS** which is the one that for some reason I find to be transfixing: <http://www.reallyslick.com/screensavers/>

However, I was not a fan of Helios’ default settings. Out of the box it was way too busy for me. So it needed some tweaking to match my taste and I’ve captured the settings panel I use with it so that others can see it the way I see it:



Steve’s recommended Helios settings

Terry has written a total of 12 screensavers. So even if you don’t fancy Helios, you might find one or more among those others that you do. And, in fact, I was playing with his “Microcosm” screensaver last night. If really cool looking liquid 3D objects is more your thing, check out “Microcosm” — though it, too, benefits from some tweaking.

Closing the Loop

JP: *On today's Security Now you discussed TOR and HTTPS. By definition, Tor hidden sites are not going to use HTTPS - getting a cert would make mockery of being hidden. And, a self signed cert these days just pops up flags.*

JP’s assertion was interesting to me since I had never looked into the issue of Tor .onion sites themselves using HTTPS. JP was a bit confused about our previous discussion, since in that discussion of the problems with Tor exit node security we were talking about non-HTTPS connections being made to external websites on the Internet, not to internal .ONION sites. But that left the interesting question about obtaining TLS certs for .onion Tor hidden sites.

It seemed to me that any ACME-based TLS certificate issuer such as Let's Encrypt would be what one wanted. And it turns out that until recently only EV certs could be issued for .ONION domains. And it was truly by coincidence that DigiCert, my chosen certificate authority appears to be the choice for EV .ONION certs:

<https://www.digicert.com/dc/blog/ordering-a-onion-certificate-from-digicert/>

It was initially unclear to me why EV was required, and it would also seem that needing to authenticate oneself to the level required to obtain an Extended Validation certificate could hamper some of the value provided by .onion domains. But as we'll learn in a moment, there was a rational security reason for requiring EV.

But first, I discovered that .ONION domains have a long history of HTTPS:// access. Seven years ago, when Tor users would attempt to visit their Facebook accounts, Facebook's geofencing security would trigger to lock that user's account because the traversal through Tor would cause the user to appear to be connecting from some foreign land. To fix this problem while also allowing Tor users to have a better experience when connecting to Facebook, way back in 2014 Facebook launched the dedicated Tor address <https://facebookcorewwi.onion/>. Using this SSL/TLS authenticated and encrypted onion site, Tor users could access the site directly without fear that doing so might freak out Facebook's geo-aware security. So, yes... traditional security certificates have long been available for Tor's .ONION sites.

However, I also found a CAB Forum ballot where certificate issuers (15 different certificate authorities) and four consumers (Apple, Microsoft, Google, Mozilla) voted unanimously (with a few abstentions, but no "nays"), in favor of opening up DV and OV certs to .ONION domains, and "Let's Encrypt" was among those voting in favor of this happening:

<https://cabforum.org/2020/02/20/ballot-sc27v3-version-3-onion-certificates/>

The explanation of the purpose of the ballot was also quite informative. It says:

Purpose of Ballot:

This ballot will permit CAs to issue DV and OV certificates containing Tor onion addresses using the newer version 3 naming format.

In ballot 144, later clarified by ballots 198/201, the Forum created rules for issuing EV certificates containing onion addresses. A primary reason for requiring EV level validation was that onion addresses were cryptographically weak, relying on RSA-1024 and SHA-1. More recently a newer "version 3" addressing scheme has removed these weaknesses. For much the same reason that EV certificates are not always a viable option for website operators (e.g. sites operated by individuals), many onion sites would benefit from the availability of DV and OV certificates for version 3 onion addresses.

The Tor Service Descriptor Hash extension required in the EV Guidelines to contain the full hash of the keys related to the .onion address is no longer needed as this hash is part of the version 3 address.

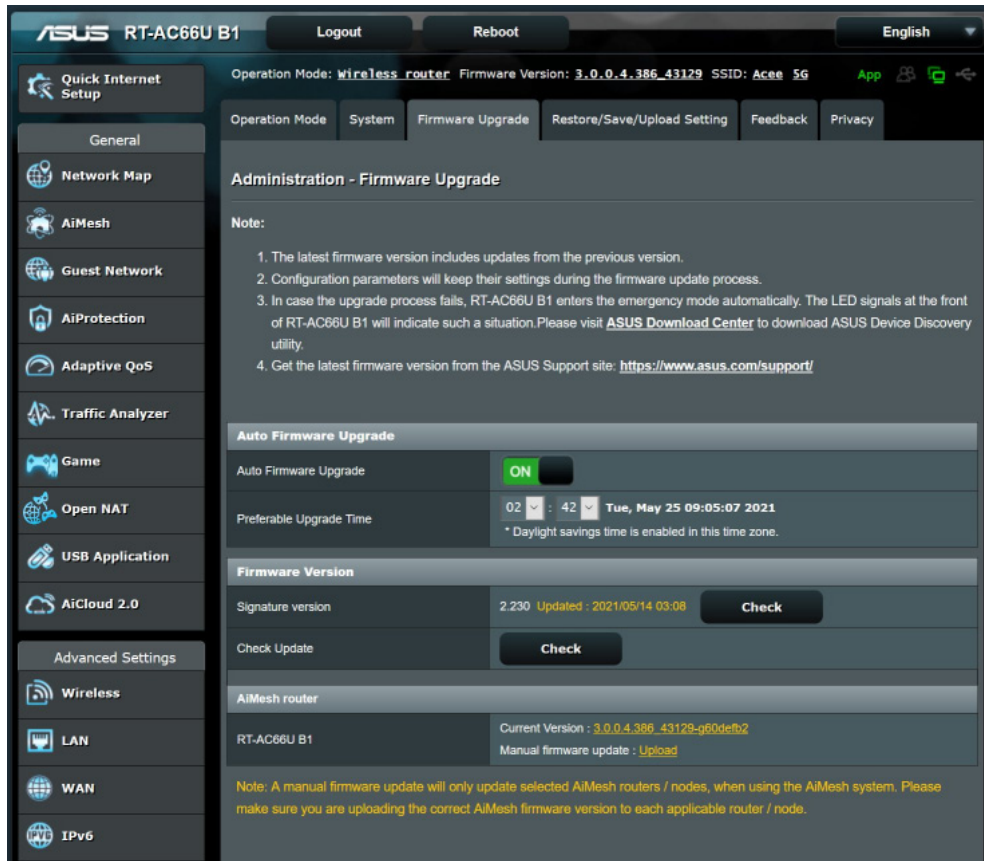
Older version 2 onion addresses are still in use, so this ballot does not remove the existing EV Guidelines requirements for onion names.

This balloting occurred back in February, 2020, so more than a year ago. I haven't been able to

locate any evidence of anyone other than DigiCert issuing .ONION certs. But it appears that Lets Encrypt is all onboard for this, so I'd imagine that it's just a matter of time before it begins happening.

Mikael Falkvidd / @mfalkvidd

Hi. I just wanted to let you know that our Asus routers now support auto update. This feature was not included in the release notes but fix for FrAg attacks were included.



Jared Stein / @jaredstein

Steve, Since you always talk about science fiction books, I was wondering if you could suggest a starting book, one that would either lead me into continuing to read it or determine it's not for me. I do really appreciate security now and all you bring to the community.

*Regards,
Jared*

There are as many science fiction authors as there are musicians. And taste for the work of this or that author is probably as personal as taste in music. But one of my favorite authors is Peter F. Hamilton and his "Fallen Dragon" is a great example of his imagination. If you like that, then his pair of novels "Pandora's Star" and "Judas Unchained" would be next in line.

Kerry Blue Life / @KerryBlueLife

Steve, Thanks for the Martian recommendation years ago. Andy Weir has done it again with Hail Mary. I was surprised how emotional it could be. Cheers Geoffrey Burkholder

Craig Manske / @albion0

*@SGgrc Thanks for Hail Mary. Work on Monday is so much closer now. *sigh* BTW, great narrator, same as the Bobaverse books. I kept waiting for Bob or piggies to show up.*

Okay. Now I have no idea who Bob and the Piggies are... but Leo... How did the Andy Weir interview go Friday? I've stayed away from it since I definitely plan to read "Hail Mary" just as soon as the book series I'm rereading hits a slow spot.

Rob Mitchell / @TheBTCGame

Whoa, just heard your talk about Eufy at the end of Security Now. I recently bought some Eufy cameras and Eufy has my email address because they asked me to pay for their hosting service and responded to a support question. But they did not notify me about this security issue. I had so far been impressed with them...

I did not take time to study the issue in detail, which is why we only mentioned it in passing last week. But if we wanted to assume that they were acting responsibly, it may have been that only a subset of their users were affected and that they were able to determine who they were.

Mike Lawrence / @MikeLwrnc

@SGgrc Have you given any thought to Vaccine e-passport verification/privacy? One concern is verification can induce a record for tracking activities, but I feel like there's a public key solution to that.

Yes, there is indeed a public key solution to that. In fact, there's a public key solution for most things!

I played around with this a bit Monday. The largest possible standard QR code for 8-bit binary data can encode 3K bytes. I took a head shot of myself and used JPEG compression to reduce its size to 3K and pasted it into today's show notes . It's entirely recognizable as me:

This is a 3K byte JPG Image:



And a QR code's digital data could be signed with a government's private key and subsequently **unspoofably** verified with the government's matching public key. So, for example, a credit card size vaccination ID could be created containing its holder's photo, its signed QR code and its signature side by side. Then, when needing to prove vaccination status, say, at an arena, for indoor dining, or perhaps even dating, people could present their card to the ticket agent, the receptionist or their date. Upon scanning the card, the validity of the QR code would be instantly — and locally — verified without any need to phone home, and the QR code would also present the user's face on the verification screen for comparison to what's shown on the card and with the face of the person presenting the card.

To be clear, I'm NOT proposing that this is what we or anyone should do. I'm just noting that today's crypto technologies provide us with such a flexible toolkit that they can be used to provide solutions to nearly any problem.

Mike also asked about privacy. The system I've described is able to validate that a visually recognizable and digitally scannable image of an individual can be signed using a private key and that signature can therefore be authenticated without any communications. But to be more useful moving forward, if a query were allowed to be trusted, non-tracking, privacy advocacy group, then the individual's relevant vaccination history could be retrieved, without any logging, to make this single card also useful for any possible boosters or next pandemic vaccines.

The Dark Escrow

I wanted to finish up this week by making an observation about something that's going on in the dark underworld.

There are now around 20, individually identifiable ransomware groups with most now—though not all—operating under the ransomware as a service (RaaS) model because it is proving to create additional value through specialization of function. And thus arises the problem of the responsibility to pay affiliates their share of a ransom after payment has been received from the victim. In the case of Darkside's decision to abruptly shutter their operations and/or being forced to do so—we don't have full visibility into exactly what went on there—affiliate postings have begun appearing on dark web forums complaining about non payment of affiliate commissions which have been earned and which are due. And yeah, I agree... "Oh, Boo Hoo." But the existence of well known dark web forums means that this is happening in plain sight for all other RaaS service operators and current and would-be affiliates to witness. Service operators want to attract affiliates and affiliates want to be reliably paid. In this environment, the presence of dark web forums means that we now have a marketplace with low friction information flow and communications. This, in turn, means that three things are likely to happen:

First, the notion that not all RaaS operators are the same will evolve. RaaS operators will begin to acquire a reputation for reliable and timely payment. And since, to a large degree, ransomware is ransomware and cash is king, reliability of payment will largely dictate future affiliate choices. And affiliates can be expected to be extremely fickle. ANY mistake or payment dispute will instantly doom any RaaS service. It's obviously over for Darkside. No one would ever trust them again, especially as the underworld is witnessing the many complaints about non-payment of earned commissions.

Now there's competition among RaaS service providers. So I expect that the second thing we're going to see is some jockeying over commission rates. The only thing the affiliates care about is money. So as ransomware matures we can expect to see commission rates settle and mature too. Those services which have built the best reputations will be able to charge a higher price for the use of their ransomware by taking a larger piece of the pie. And newer upstart services which are not yet established and trusted may need to lure new affiliates to use their product by offering the use of their ransomware at a greater discount than the more well established operators.

"Hooked a big fish? Want to keep more of what you're about to earn? Consider encrypting your target's network with Newbieware... we only take 5%! We also have the fastest encryption around so your target will never know what hit them. And we have big pipes, hosting your ill gotten goods on AWS, the industry's most reliable cloud storage."

And finally, the third thing I expect we're going to be learning about before long, will be the emergence of another player in the evolving increasingly specialized multi-component RaaS ecosystem: Ransom escrow services.

The publicity surrounding the Darkside collapse is likely to bring about the emergence of a neutral intermediary to manage the money. Escrows are a time honored system for holding and transferring valuable assets among untrusted parties. The Cryptocurrency system makes it easy to have the ransom paid to a wallet that's not under the control of either the RaaS service or their affiliates. This helps to insure that the affiliate will be paid no matter what might happen to the RaaS service. When tens of millions of dollars are at stake — and when payment might be all or nothing — shaving off a quarter point for escrow commision will probably seem like a wise investment. It will give affiliates the payment assurance they will now be clamoring for, and it will allow RaaS services to bootstrap themselves by offering to escrow ransom payments as an option. Given the dynamics of this dark ecosystem, I'll be surprised if we don't learn about ransomware escrow services appearing before long.

For completeness, I should note that there have been some informal first stabs at this. We noted a year or two ago when the REvil ransomware gang deposited \$1 million worth of bitcoin into a different hacking forum as a means of attracting affiliates. And to show that they, too, meant business, Darkside placed 22 bitcoin on deposit with the admin of the XSS forum. But this was more in the form of a guarantee than an escrow. And 22 bitcoins won't begin to cover the sorts of ransoms we're seeing recently. So this also wasn't an escrow. As we detailed last week, Darkside victim funds flowed directly into Darkside's bitcoin wallet, which had been identified by Eclipse. And, moreover, claims now being made against the Darkside bitcoin guarantee stash, which the XSS forum is administering, are meeting with trouble being paid. This also suggests why that XSS admin may have said that ransomware is no longer welcome there. This admin is likely quickly getting fed up with the unwanted responsibility of being Darkside's unpaid defacto guarantor. And one wonders what happens to any unclaimed guaranteed funds? Whose are they?

It does appear that DarkSide did the best they could after they stumbled into the Colonial Pipeline nightmare. Their affiliates who successfully encrypted victim networks have all received the corresponding decryption keys which allow them to pursue negotiations with their victim companies independently. Though that's not what they signed up for, either.

After writing the above, it occurred to me to Google the phrase "cryptocurrency escrow" which returned more than two and a half million hits. So escrowing cryptocurrency is obviously not a new concept. However, no reputable commercial cryptocurrency escrow service wants to receive a letter from Ireland's High Court demanding that they freeze funds in escrow. So the dark web will need to establish its own dark escrow services. I expect we'll be hearing about such before long.