Security Now! #951 - 12-05-23
Revisiting Browser Trust

This week on Security Now!

How can masked domain owners be unmasked? What new and very useful feature has WhatsApp just added? How did Iranian hackers compromise multiple U.S. water facilities across multiple states? Did Montana successfully ban all use of TikTok statewide?, and is that even possible? How many Android devices are RCS-equipped? What's the EU's Cyber Resilience Act?, and is it good or bad? Is ransomware finally beginning to lose steam? What's the deal with all of these new top level DNS domains? Do they make any sense? Has CISA been listening to this podcast, or have they just been paying attention to the same things we have? What's up with France's ban on all "foreign" messaging apps?, and did the Prime Minister's nephew come up with an alternative? And I want to share two final insights from independent industry veterans regarding the EU's proposal to forcibly require our browsers and operating systems to trust any certificates signed by their member countries.

Enough said.

Security News

One of the things that was always chafing while I was with Network Solutions was the idea of my paying them additional money – annually, even – to redact the domain registration listings they themselves had created for ICANN's Internet's WHOIS database queries. The original idea behind domain registration was for it to be public. That was when we were all one big happy Internet. But it wasn't long before spammers and scammers were scraping the public domain registration WHOIS database for information and abusing it in every way imaginable. So it became prudent to have the data masked, and masking services appeared. Then the domain registrars themselves began offering this extra service with many seeing the provision of this masking as another revenue opportunity. One of the many reasons I'm so glad I left Network Solutions and moved to Hover is that domain registrations are masked by default at no extra cost.

Now, in the EU with the GDPR, things are somewhat different. As we know, the GDPR has had its pluses and minuses. One of the minuses we now all experience every day is the pervasive annoyance of every website being forced to wave its cookie policies in our faces and obtain our acknowledgement and consent. On the flip side, one of the pluses is that the GDPR includes a stringent data protection law that has forced domain registrars to redact information on domain owners from their publicly available WHOIS databases. This information is still present in the private databases of domain registrars and has historically been made available to some organizations, but usually only in a very limited fashion, such as through court orders, subpoenas, or following intelligence-sharing agreements.

I'm bringing this up because last Tuesday ICANN announced a new facility to improve the current situation for those, such as law enforcement, who have a legitimate need to obtain access to otherwise redacted domain ownership information. With a bit of editing, here's what ICANN said:

> *The Internet Corporation for Assigned Names and Numbers (ICANN) has launched the Registration Data Request Service (RDRS). The RDRS is a new service that introduces a more consistent and standardized format to handle requests for access to nonpublic registration data related to generic top-level domains (gTLDs).*
>
> *Personal data protection laws now require many ICANN-accredited registrars to redact the personal data from public records that was previously available in their "WHOIS" databases. With no one way to request or access such data, it can be difficult for interested parties to get the information they need. The RDRS helps by providing a simple and standardized process to make these types of requests.*
>
> *The RDRS can be an important resource for ICANN-accredited registrars and those who have a legitimate interest in nonpublic data, like law enforcement, intellectual property professionals, consumer protection advocates, cybersecurity professionals, and government officials.*
>
> *The RDRS is a free, global, one-stop shop ticketing system that handles nonpublic TLD registration data requests. The RDRS connects requestors of nonpublic data with the relevant ICANN-accredited registrars for TLD domain names that are participating in the service. The service will streamline and standardize the process for submitting and receiving requests through a single platform.*

> *The service does not guarantee access to requested registration data. All communication and data disclosure between the registrars and requestors takes place outside of the system.*
>
> *By utilizing a single platform and request form, the RDRS provides a consistent and standardized format for handling nonpublic TLD registration data requests. This simplifies the process for requestors by automatically identifying the correct registrar for a domain name and preventing the need to complete multiple forms with varying sets of required information managed by different registrars.*
>
> *The service also provides a centralized platform where requestors can conveniently access pending and past requests. They also have the ability to create new requests, develop request templates, and cancel requests when needed.*
>
> *Registrars can benefit from using the service as it provides a mechanism to manage and track all nonpublic data requests in a single location. Registrars can receive automated alerts anytime a request is submitted to them. The use of a standardized request form also makes it easier for the correct information and supporting documents to be provided to evaluate a request.*

To me, this makes so much sense. There are so many shenanigans going on with Internet domain names that abusers of the system need to know that their ability to hide is being reduced. And legitimate domain owners should have a reasonable expectation of privacy. So the idea of standardizing the process of obtaining this information seems like a long-missing piece that's finally being provided.

Meta introduces "Secret Code" for "Chat Lock"

Due to the strength of FaceBook, Meta's WhatsApp is the most popular messaging app in the world. And last Thursday WhatsApp announced a significant new feature which was missing when they announced "Chat Lock" last May. So first, here's what they announced on May 15th under the headline "Chat Lock: Making your most intimate conversations even more private":

> *Our passion is to find new ways to help keep your messages private and secure. Today, we're excited to bring to you a new feature we're calling Chat Lock, which lets you protect your most intimate conversations behind one more layer of security.*
>
> *Locking a chat takes that thread out of the inbox and puts it behind its own folder that can only be accessed with your device password or biometric, like a fingerprint. It also automatically hides the contents of that chat in notifications, too.*
>
> *We think this feature will be great for people who have reason to share their phones from time to time with a family member or those moments where someone else is holding your phone at the exact moment an extra special chat arrives. You can lock a chat by tapping the name of a one-to-one or group and selecting the lock option. To reveal these chats, slowly pull down on your inbox and enter your phone password or biometric.*
>
> *Over the next few months we're going to be adding more options for Chat Lock, including locking for companion devices and creating a custom password for your chats so that you can use a unique password different from the one you use for your phone.*

I would call that a good start. But I'm mentioning this today because last Thursday Meta announced the roll-out of that planned accompanying feature, which is the addition of a "something you know" password to separately and independently protect extra-sensitive chat content. Here's what Meta said under their headline "Introducing Secret Code for Chat Lock":

> *Earlier this year we rolled out* ***Chat Lock*** *to help people protect their more sensitive conversations. Today we're launching* ***Secret Code****, an additional way to protect those chats and make them harder to find if someone has access to your phone or you share a phone with someone else.*
>
> *With* ***Secret Code*** *you'll now be able to set a unique password different from what you use to unlock your phone to give your locked chats an extra layer of privacy. You'll have the option to hide the Locked Chats folder from your chatlist so that they can only be discovered by typing your secret code in the search bar. If that doesn't suit your needs, you can still choose to have them appear in your chatlist. Whenever there's a new chat which you want to lock, you can now long press to lock it rather than visiting the chat's settings.*
>
> *We're so happy our community has been loving Chat Lock, and hope that secret code makes it even more useful to them. Secret code starts rolling out today, and in the coming months will be available globally. We're excited to keep bringing more functions to Chat Lock to help people protect their privacy, let us know what you think.*

I think it makes total sense and I predict that it will become a heavily-used feature. From a privacy and security standpoint it makes sense for our devices to have multiple layers and levels of protection. We need to have more than just a device being locked or unlocked. That's no longer sufficient. And I don't think that using the same password or a biometric makes sense for an "inner" level of protection. Locking enhanced layers of privacy behind "something you know" makes the most sense.

Really? You never changed the default password?

I said recently that one of the broad changes to the way we've always done things must somehow be the elimination of any initial default password from our devices. My first thought was to require the user to set a password themselves while preventing them from setting it to "password" or "Monkey123" by also embedding some minimum complexity requirements. But I don't think that's the right solution. I think the right answer is to have the device randomly assign a strong password when it's initially set up, and that's it. The user needs to write it down. Period. We've been talking for years about the need to be using strong passwords that we cannot recall. That needs to apply to equipment as well as websites.

Here's the news that brought me back to this train of thought. Get a load of this:

The US government has confirmed that an Iranian hacking group named Cyber Av3ngers has gained access to equipment at water facilities across multiple US states. CISA, the FBI, the NSA, and other agencies say the attacks began around November 22 and exploited PLCs (programmable logic controllers) manufactured by the Israeli company Unitronics. The group targeted Unitronics PLCs that were still using the default password "1111." Wow. So last week CISA asked US organizations to please change the default password, enable MFA, and remove the devices from the Internet. Gee, what a concept. US officials say the Cyber Av3ngers group is affiliated with the IRGC, an Iranian military and intelligence organization. According to the Shadowserver Foundation, between 500 to 800 Unitronics PLCs are currently exposed to the Internet, with 66 identified in Australia, 52 in Singapore, 42 in Switzerland, 37 known to be in the United States, and Estonia and Spain both with 31 and so on down the list across the globe.

Unlike web servers, PLC systems typically have no need to be exposed to the Internet. Doing so should require jumping through some real hoops. And under no circumstances should a device be produced where it both has a well-known default password and is also exposing any interface protected by that default password to the Internet. In today's world designing and selling such systems is nothing short of irresponsible.

We've talked in the past about countries becoming proactive in scanning their own Internet address space with an eye toward getting ahead of attackers and cleaning up some of these issues. This is the sort of thing that CISA in the US ought to be considering.

Montana's attempted ban on TikTok stalled

A while back we covered the news that a bunch of states were enacting legislation to block the use of TikTok on government devices within their jurisdictions. Doing that was likely within their power. But the state of Montana wanted to go further and outright ban **all** use of the TikTok service statewide. From a purely technical standpoint this would be somewhat tricky, since network boundaries and state borders are not currently aligned since there's never been any need to align them. But now it appears that it might not matter after a recent federal ruling on Thursday. NPR's coverage of this also provides some interesting background. They wrote:

> *A federal judge has blocked a law in Montana that sought to ban TikTok across the state, delivering a blow to an unprecedented attempt to completely restrict a single app within a state's borders. The ruling, which came on Thursday, means that Montana's TikTok ban, which was set to go into effect on Jan. 1st, has now been temporarily halted. U.S. District Judge Donald Molloy said Montana's TikTok ban "oversteps state power" and "likely violates the First Amendment."*
>
> *Molloy wrote that though officials in Montana have defended the law as an attempt to protect consumers in the state, there is "little doubt that Montana's legislature and Attorney General were more interested in targeting China's ostensible role in TikTok than with protecting Montana consumers." Montana, as a state, does not have authority over foreign affairs, Molloy said, but even still, he found the national security case presented against TikTok unconvincing, writing that if anything, the Montana law had a "pervasive undertone of anti-Chinese sentiment."*
>
> *The ruling is preliminary with a final determination to be made following a trial expected some time next year. TikTok, which has more than 150 million American users, has for years been under intense scrutiny over fears that its Beijing-based parent company, ByteDance, would hand over sensitive user data to Chinese authorities, or that Beijing would use the app as a propaganda tool — even though there is no public proof that either has ever happened.*
>
> *Although several states and the federal government have prohibited the app from being downloaded on government devices, Montana was the first state to pass an outright ban of the*

> *app. Some critics have accused it of government overreach. In May, TikTok sued the state over the law, arguing that it amounts to an illegal suppression of free speech. Lawyers for TikTok argued that the national security threat raised by officials in Montana was never supported by solid evidence.*
>
> *Molloy, the judge overseeing the case, was skeptical of the ban in an October hearing on the lawsuit. He pointed out that TikTok users voluntarily provide their personal data, despite state officials suggesting the app was stealing the data of users. He said state officials justified the Montana ban under a "paternalistic argument."*
>
> *As Washington continues to debate TikTok's future, states have been acting faster, and the law in Montana was considered an important test case of whether a state-level ban of the app would survive court challenges. Backing the Montana law were 18, mostly Republican-led states that were eyeing similar bans of TikTok. Aside from the legal hurdles to implementing such laws, cybersecurity experts have raised questions about how, from a technical standpoint, such a ban would even be possible.*

Right. Count me in that group. Those pesky technical details which keep tripping up the legislators who believe that they can simply have any magical technology they demand. NPR writes:

> *President Trump clamped down on TikTok and attempted to outlaw the app, but his efforts were twice struck down in the courts. National security experts say TikTok is caught in the middle of escalating geopolitical tensions between the U.S. and China, as Washington grows ever more concerned about the advancement of Chinese tech, like semiconductors, and the country's investments in artificial intelligence.*
>
> *Supporters of restricting or banning TikTok in the U.S. point to Chinese national security laws that compel private companies to turn information over to Beijing authorities. They also point to ByteDance, TikTok's corporate owner. It admitted in December that it had fired four employees, two of whom worked in China, who had improperly accessed data on two journalists in an attempt to identify a company employee who leaked a damaging internal report.*

I am by no means defending TikTok. But let's not forget that many domestic companies as well as many of our own US law enforcement agents have also been caught with their hands in the cookie jar. Access to personal and private data appears to be quite tempting. So it's not just Chinese misbehavior.

> *TikTok says China-based employees no longer have access to U.S. user data under a new firewall it has put in place with the help of Oracle. With this change, dubbed Project Texas after Oracle moved its corporate headquarters to Austin, all Americans' data will be stored on servers owned and maintained by Oracle, with additional oversight from independent auditors.*

TikTok is obviously an extremely successful and valuable service. It seems to me that they're making every effort to legitimately assuage concerns of secret Chinese influence. Today's social media is all about influence. But such influence is as pervasive with FaceBook and 'X' as it is anywhere else.

RCS is now enabled on more than one billion Android devices

We recently noted Apple's announcement that they would be upgrading their non-iMessage messaging use of SMS and MMS to RCS. So it was noteworthy that last Thursday Google announced that its RCS messaging system is now enabled on more than one billion Android devices. It appears that Android users will be ready once Apple joins them with RCS next year.

The EU does some good things under the "EU Cyber Resilience Act"

Much as I'm becoming increasingly annoyed with the EU over their move to commandeer our web browser's well established system of trust – which spend some time examining more deeply at the end of this podcast – it appears that the EU's European Council and Parliament have reached a useful agreement known as the Cyber Resilience Act. This is a piece of legislation designed to improve the security of smart devices sold within the European Union. The new regulation applies to products ranging from baby monitors and smartwatches to firewalls and routers. Under the new rules, vendors **must** establish processes to receive reports about vulnerabilities, and **must** support products for at least five years. Moreover, products will be required to come with free and automatic security updates as the default option, must ensure data confidentiality using encryption, and vendors must inform authorities of any attacks. This won't be happening immediately, however. The requirement set by the new rules will come into effect three years after the Cyber Resilience Act is formally voted on the EU Parliament floor. Given the requirements which will likely require some redesign and new infrastructure, that seems reasonable. At least in this regard the EU is usefully leading in the direction we need to head.

There's (still) big money in ransomware

The forensics industry is getting better at tracking cryptocurrency flows, and cyber insurance firms are being more forthcoming about what they're seeing. So we know more now than we have previously.

For example, one of the newer upper echelon ransomware groups is known as Black Basta. This gang is believed to have netted more than $107 million in ransom payments since it first appeared and began operations early last year. Since we're closing out 2023, that's $107 million dollars in less than two years' time. That $107 million dollar number represents payments made by more than 90 victims of the **329** organizations known to have been hit by the gang. Think about that for a minute. There are 365 days in most years. Yet in less than two years, 329 individual organizations were breached by this gang. On average about one every other day. The largest payment was $9 million, while the average ransom payment was $1.2 million. This is according to joint research published by the blockchain tracking company Elliptic, and the cyber insurance provider Corvus Insurance.

Unfortunately, what this shows is that there's a great deal of money to be made through cyber extortion, and the hostile governments (or government, since it's known to be Russia) harboring these criminals are more than happy to turn a blind eye. This means that a great deal of pressure will continue to be placed on the security of our networks and systems. And, unfortunately, as the last few months of many very serious large weaknesses and compromises continue to show, our networks and systems are not up to the challenge. Years of laxity in the design, operation, configuration and administration of these systems is catching up with us. We know that thanks to the inherent inertia which works against change, we're not going to fix these endemic problems all at once. But they're never going to get fixed at all if we don't apply constant effort in that direction.

Google's .meme TLD

Google is offering a new ".meme" top-level domain for anyone who wants to play with meme-related Internet properties. It's difficult to keep up with all of the new TLD's appearing, and it feels as though this aspect of the Internet's original design, which is to say the concept of a hierarchy of DNS domains, anchored by just a few major classifications, is not evolving well. There are companies that attempt to snatch up their existing "dot COM" second level domain name in each of the other TLD's, presumably to preserve their brand and trademark. But that's not in keeping with the spirit of creating additional DNS hierarchies for future growth. I have no interest in "grc.meme", and "grc.zip" would have caused all kinds of confusion – what is that, GRC's entire website in a ZIP archive? No thanks.

Secure by Design Alert

Last Wednesday, CISA introduced a new series of publications called *"Secure by Design"* with its first alert titled *"How Software Manufacturers Can Shield Web Management Interfaces From Malicious Cyber Activity."* And if I didn't know – as I do – that anyone who's focused on security would naturally come up with the same thoughts, I would think that they had been listening to this podcast. Get a load of what's in this document. For example, CISA writes:

> *Malicious cyber actors continue to find and exploit vulnerabilities in **web management interfaces**. In response, software manufacturers continue to ask why **customers** did not harden their products to avoid such incidents.*
>
> *"Secure by design" means that **software manufacturers** build their products in a way that reasonably protects against malicious cyber actors successfully exploiting vulnerabilities in their products. Baking in this risk mitigation, in turn, reduces the burden of cybersecurity on customers. Exploitation of vulnerabilities in web management interfaces continues to cause significant harm to organizations around the world—but can be avoided at scale. CISA urges software manufacturers to learn from ongoing malicious cyber activity against web management interfaces by reviewing the principles below.*
>
> ***Principle 1: Take Ownership of Customer Security Outcomes***
>
> *This principle focuses on key areas where software manufacturers should invest in security: application hardening, application features, and default settings. When designing these areas, software manufacturers should examine the default settings of their products. For instance, if it is a known best practice to shield a system from the public internet, **do not rely on customers to do so**. Rather, have the product itself enforce security best practices. Examples include:*
>
> - *Disabling the product's web interface by default and including a "loosening guide" that lists the risks—in both technical and non-technical language—that come with making changes to the default configurations.*

> *• Configuring the product so that it does not operate while in a vulnerable state, such as when the product is directly exposed to the internet.*
>
> *• Warning the administrator that changing the default behavior may introduce significant risk to the organization.*
>
> *Additionally, software manufacturers should conduct field tests to understand how their customers deploy products in their unique environments and whether customers are deploying products in unsafe ways. This practice will help bridge the gap between developer expectations and actual customer usage of the product. Field tests will help identify ways to build the product so customers will securely use it.*
>
> *Furthermore, software manufacturers should consistently enforce authentication throughout their product, especially on critical interfaces such as administrator portals.*

Amen to all that. Of course, not one of those concepts will come as news to the listeners of this podcast, but it would be great if those manufacturers to whom CISA is addressing this Alert would immediately take heed. We know that it's going to take time for any such changes to work their way through the entire supply chain, from drawing board into final deployment. It would have been nice if we could have started that "Secure by Design" process ten years ago, but we haven't even fully started it today. The fact that this alert has been published, with what it says, is a very good sign. I suspect that this is the first step toward beginning to hold the designers of these systems accountable for their default security.

Unfortunately, due to the "hold harmless" nature of software and equipment licensing agreements, accountability is difficult to create. I intensely dislike the idea of having government criminalize insecure design. That's a slippery slope that's not that far from what the EU is planning to do with their eIDAS 2.0 web certificate overreach. Legislation and technology rarely make great bedfellows. But one of the ways we've seen government influence things for the better is by using its own purchasing power to create voluntary incentives.

With CISA, the U.S. government finally has a highly effective and worthwhile cybersecurity agency. Based upon what CISA just published last Wednesday, it would not be a stretch to imagine adding exactly those default network behavioral requirements to any future software and equipment purchasing made by state and federal government agencies. That would affect voluntary change overnight. Vendors would be required to legally attest that their equipment abides by this new set of requirements and if it was later found not to be true, then they could be held liable for damages resulting from the functional out-of-spec behavior of their equipment. And just to be clear, not for bugs in their systems, but for the deliberate design of those systems. As I've repeatedly observed, anyone can make a mistake, but vendors can and should be held responsible for their policies.

France to ban "foreign" messaging apps

And while we're on the subject of things governments do, also last Wednesday, France's government announced a near immediate ban – as in 10 days from then – on the use of what they called "foreign end-to-end encrypted messaging apps".

France has banned government officials from using “foreign” encrypted messaging services including specifically Telegram, Signal, and WhatsApp. The government is notifying its ministers and their cabinet staff that they **must** uninstall any such applications from their devices by this coming Friday, December 8th. French officials have been told to use the French developed alternative messenger known as Olvid. Officials cited privacy risks and a need to *“advance towards greater French technological sovereignty.”*

Okay. So what the heck is Olvid? Even though we’ve never talked about it here, I have to say that it looks pretty good. It’s open source for both Android, iOS, macOS and Windows, and it’s living over on GitHub. Here’s how it describes itself:

> *Olvid is a private and secure end-to-end encrypted messenger.*
>
> *Contrary to most other messaging applications, Olvid does not rely on a central directory to connect users. As there is no user directory, Olvid does not require access to your contacts and can function without any personal information. The absence of directory also prevents unsolicited messages and spam.*
>
> *Because of this, from a security standpoint, Olvid is not "yet another secure messenger". Olvid guarantees the total and definitive confidentiality of exchanges, relying solely on the mutual trust of interlocutors. This implies that your privacy does not depend on the integrity of some server. This makes Olvid very different from other messengers that typically rely on some "Trusted Third Party", like a centralized database of users or a public blockchain.*
>
> *Note that this doesn't mean that Olvid uses no servers (it does). It means that you do not have to trust them: your privacy is ensured by cryptographic protocols running on the client-side (i.e., on your device), and these protocols assume that the servers were compromised from day one. Even then, your privacy is ensured.*

So this is less looney than it might seem at first, though it does have some feeling of nationalism and protectionism with the French government labeling everything else “foreign” and talking about the need to increase France’s technological sovereignty. But that said, Olvid is not some random homegrown messaging app designed by the Prime Minister’s nephew.

I’ve not had time to look at it closely, but it looks like the real deal and the more I look at it, the more I like it. Over on Olvid’s website, which is https://olvid.io, they proudly note that: *“Olvid does not require any personal data: no phone number, no email, no name, no surname, no address, no date of birth. No nothing. Unlike your previous messenger, Olvid will never request access to your address book.”* Those are some compelling features. And under the headline “Compatible with what you already have” they say: *“Olvid is available for your macOS and Windows computers, as well as for your iPhones, iPads, Android smartphones and tablets. No SIM? No problem. No SIM card required. Wifi is all you need. Since Olvid needs no phone number to work, you can use any of your devices. And they'll stay in sync. Olvid even works in an emulator. Geeks will love it.”*

Olvid uses something known as SAS-based authentication where SAS stands for “Short Authenticated Strings.” The concept of SAS was proposed and formalized in a 311-page PhD thesis by a French cryptographer, Sylvain Pasini, back in 2009. Here’s what Pasini explained in the first two paragraphs of his thesis' Abstract:

> *Our main motivation is to design more user-friendly security protocols. Indeed, if the use of the protocol is tedious, most users will not behave correctly and, consequently, security issues occur. An example is the actual behavior of a user in front of an SSH certificate validation: while this task is of utmost importance, about 99% of SSH users accept the received certificate without checking it. Designing more user-friendly protocols may be difficult since the security should not decrease at the same time. Interestingly, insecure channels coexist with channels ensuring authentication. In practice, these latters may be used for a string comparison or a string copy, for example, by voice over IP spelling. The shorter the authenticated string is, the less human interaction the protocol requires, and the more user-friendly the protocol is. This leads to the notion of SAS-based cryptography, where SAS stands for Short Authenticated String.*
>
> *In the first part of this thesis, we analyze and propose optimal SAS-based message authentication protocols. By using these protocols, we show how to construct optimal SAS-based authenticated key agreements. Such a protocol enables any group of users to agree on a shared secret key. SAS-based cryptography requires no pre-shared key, no trusted third party, and no public-key infrastructure. However, it requires the user to exchange a short SAS, for example, just five decimal digits. By using the just agreed secret key, the group can now achieve a secure communication based on symmetric cryptography.*

Since 2009 this SAS protocol has received a great deal of further scrutiny and it has held up. So this works by having the users at each end initially discover each other by sharing the short tokens being displayed on each other's devices. For that some form of already-authenticated out of band channel is used, like an audio or video call, to exchange the information that each user's device presents. And this simple process has been proven to be cryptographically sound.

I really like its integration with the desktop. That's something I've been missing as a cross-platform iOS and Windows user. And Signal is annoying with its required tie to a phone number. Signal claims that's needed to prevent spam, but with Olvid there's no possibility of being spammed. A real world out-of-band interaction is required to establish a channel between two participants or among the participants in a group. After that, the devices remain linked for further communication.

So what pays for this? The system runs on a "freemium" model. All bidirectional text messaging and incoming audio is free. You get unlimited messages, unlimited attachments, secure group discussions, unsend and edit messages, remote deletion, ephemeral messages, multiple profiles, user mention, markdown, etc., Olvid Web (whatever that is) and inbound secure audio calls. The system is financially supported at a 5€ Euros per month by those who want to be able to initiate secure voice calls as well as use Olvid on multiple devices that receive all messages and synchronize. There are also more powerful enterprise plans with more features. It's interesting that the French government is telling their ministers and cabinet staff that they must switch to Olvid. Since only text messaging is completely free, perhaps that's all that will be presumed.

Since Olvid is authentically interesting, open source, mature and apparently secure messaging technology – the company also participates in a bug bounty program – I wanted to make all of our listeners aware that it existed. It might suit many other peoples' needs. https://olvid.io

Revisiting Browser Trust

We've been covering the news of the now-impending EU eIDAS 2.0 legislation mostly from the standpoint of the two open letters that those in the industry and academia have authored and co-signed. And by *"those in the industry and academia"* I mean a number now totaling more than 500 individuals who are truly concerned about what the EU is about to unilaterally place into law. Four weeks ago this podcast was titled Article 45. So I understand that we've already talked about this. But I just encountered two new pieces of commentary from two well-placed technologists. So I decided to share their appraisals to create some *"what it would really mean to the world"* perspective.

The first person's name is Ivan Ristić. I was immediately curious when I saw that Ivan had chosen to weigh-in and address this issue. If Ivan's name doesn't immediately jump out and mean anything to you, the well-known website and service he created "**SSL Labs**" probably does. For as long as I can remember, Ivan's SSL Labs at https://ssllabs.com has been the go-to site for checking the security at both the server and browser ends of secured connections. Ivan is also the author of two books: "Bulletproof TLS and PKI, Understanding and deploying SSL/TLS and PKI to secure servers and web applications". Its first edition was published nearly ten years ago, in 2014 and the book is now in its second print edition with added coverage of TLS v1.3. It's also available as an eBook. Ivan's second book is the "OpenSSL Cookbook, the definitive guide to the most useful command-line features." That one is in its 3rd edition, and is available for free. Ivan and his wife, Jelena, are based in London. His piece, which he wrote last Thursday, is titled: *"European Union Presses Ahead with Article 45"*. He wrote:

> *The European Union continues on its path to eIDAS 2.0, which includes the controversial Article 45 that basically tells browsers which certification authorities (CAs) to trust. eIDAS, which stands for electronic identification and trust services, is a framework aimed at regulating electronic transactions. As part of this proposal, the EU wants to support embedding identities in website certificates. In essence, the goal is to bring back Extended Validation (EV) certificates.*
>
> *Browsers—of course—don't want that, but the real problem is the fact that, with the legal text as it is at the moment, in its near-final form, the EU gets the final say in which CAs are trusted.*
>
> *The global security community has been fighting against Article 45 for more than two years now; we wrote about it on a couple of occasions. As of November 2023, the European Council and Parliament have reached a provisional agreement. The next step is for the law to be put to the vote, which is usually a formality.*
>
> *In November, ahead of the crucial vote, the campaigns intensified, with browser providers (Google, Mozilla), civil society groups (EFF), other companies, and more than 500 security experts voicing their concerns. In the end, it didn't help: the bureaucrats drafted the text and voted behind closed doors with little acknowledgement of the protests.*
>
> *And therein lies the main problem: the EU doesn't understand the global technical community. Internet standards are developed collaboratively and organically, with careful deliberation of the details. The EU, on the other hand, prefers a top-down approach that ignores the*

> *details—and apparently involves no debate. They expect everyone to trust that the details will turn out all right. The text voted on was published only after the fact.*
>
> *The EU might have the right to govern its territory, but when it comes to these global matters, it also has a duty to respect and compromise with the rest of the world. Above all, care must be taken to separate technology and politics as much as possible.*
>
> *After all, it took the world a very long time to achieve reasonable security of global website authentication. A decade ago, we were witnessing hackers breaking into CAs and government agencies issuing certificates for Google's properties. Today, we have much stricter issuance and security standards, and we also have Certificate Transparency, which provides visibility and auditing. No one knows what's going to happen with that, and the EU doesn't engage.*
>
> *Where are we now? The EU wants browsers to display legal identities embedded in the qualified certificates, but it also wants to control who issues them. It so happens that the same certificates are used to store the identities and authenticate websites. It's not at all clear if the EU cares about the latter part. In fact, the following statement appears in the recitals in the provisional agreement:*
>
> *"The obligation of recognition, interoperability and support of QWACs is not to affect the freedom of web-browser providers to ensure web security, domain authentication and the encryption of web traffic in the manner and with the technology they consider most appropriate."*
>
> *So can browsers recognise and show legal identities from the EU-approved CAs, but continue to require full compliance with current technical standards in order to fully trust qualified certificates? Or can browsers require two certificates, one for TLS and the other for identities, like Mozilla proposed last year? We'll need to wait and see.*

Okay. So, from Ivan Ristic we have a terrific summary.

The second piece, which Ryan Hurst wrote a little over two weeks ago is titled: *"eIDAS 2.0 Provisional Agreement: Implications for Web Browsers and Digital Certificate Trust"*. It goes further than anything I've seen so far to provide an assortment of interesting facts to clarify the way things are today, and to examine what the EU's proposed changes would mean to the industry and to the world. Get ready to be surprised. Leading with a summary, Ryan writes:

> *This document contains my notes on the problematic elements of the provisional agreement on the EU eIDAS 2.0 legislation reached by EU legislatures on November 8th* (https://www.europarl.europa.eu/cmsdata/278103/eIDAS-4th-column-extract.pdf).
>
> Ryan extracts six main points:
>
> - ***Mandatory Trust in EU-Approved CAs:*** *Browsers will be required to trust certificate authorities approved by each European member state. This could lead to scenarios where the government forces the trust of CAs that put global users at risk.*
>
> - ***Lower standards for EU-Approved CAs:*** *Establishes a lower standard for European CAs. Limiting the browser's ability to protect users from underperforming EU certificate authorities.*

> - ***EU to Override Browser CA Trust Decisions:*** *In cases where an EU investigation does not lead to the withdrawal of a certificate's qualified status, the EU can request browsers to end precautionary measures, forcing them to trust the associated CA.*
>
> - ***Establishes global precedent for further undermining encryption on the web:*** *When a liberal democracy establishes this kind of control over technology on the web, despite its consequences, it lays the groundwork for more authoritarian governments to follow suit with impunity.*
>
> - ***Browsers are forced to promote legal identity for authentication of websites:*** *Browsers will be required to have a user interface to support the display of legal identity associated with a website, potentially reversing previous design choices made based on user behavior and research.*
>
> - ***The Inconsistencies of Recitals with the Substantive Legal Text:*** *The recitals in the legal text have ambiguities and contradictions which will cause long-term negative consequences for the web.*

Ryan explains:

> *The text says browsers must either directly or indirectly take a dependency on the EU Trust List (https://eidas.ec.europa.eu/efda/tl-browser) to determine if a CA is trusted for "website authentication". This is a list of CAs as determined by each member state to be in conformance with the legal obligations under eIDAS.*
>
> *To put this into context, based on the currently authorized organizations on this list, we can expect to see* ***43 new organizations*** *added to both the Mozilla and Chrome Root Stores. This is just a number though, let's give that a little color.* ***Today there are 7 organizations in the WebPKI that are responsible for 99% of all certificate issuance.***

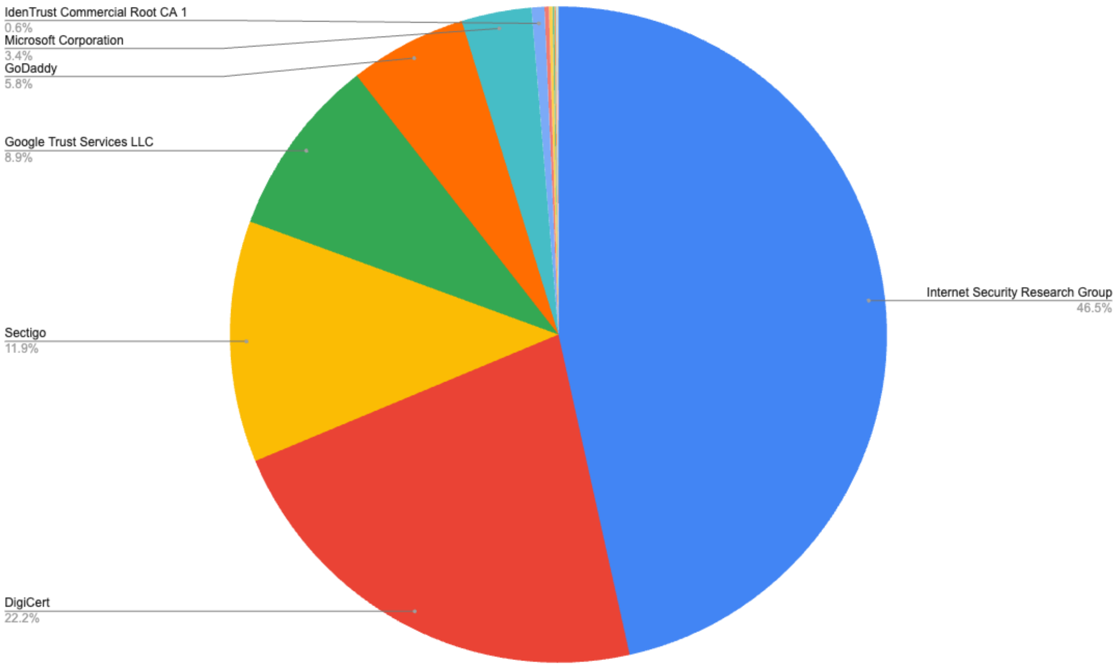

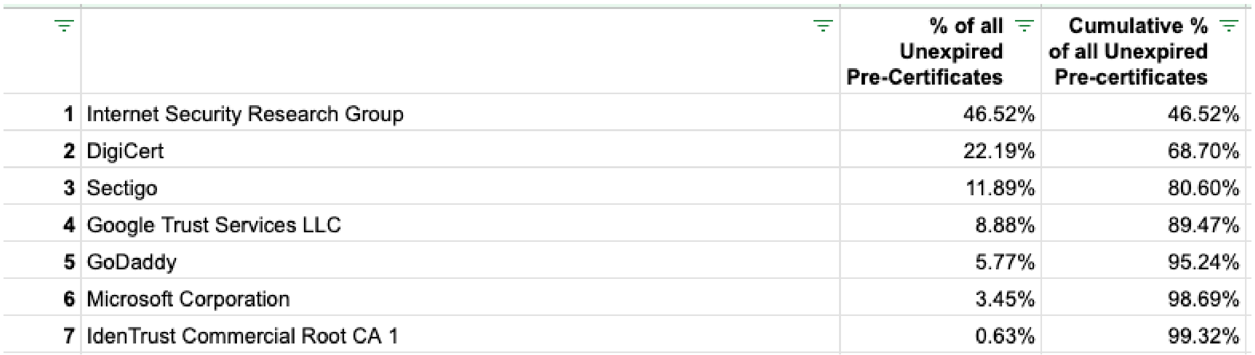

		% of all Unexpired Pre-Certificates	Cumulative % of all Unexpired Pre-certificates
1	Internet Security Research Group	46.52%	46.52%
2	DigiCert	22.19%	68.70%
3	Sectigo	11.89%	80.60%
4	Google Trust Services LLC	8.88%	89.47%
5	GoDaddy	5.77%	95.24%
6	Microsoft Corporation	3.45%	98.69%
7	IdenTrust Commercial Root CA 1	0.63%	99.32%

And we have a VERY COOL chart. The most significant fact is that just 7 certificate authorities in total account for 99.32% of all currently unexpired web certificates. In other words, if our browser root stores were to contain the trusted root certificates of only 7 certificate authorities we would be trusting the web certificates of all but 0.68% of websites... and something tells me that we could probably safely live without those few, which are likely far from the mainstream. Interestingly, this also says that everyone is obtaining their certificates from the same very small group of certificate authorities.

So, who are they? I suppose that it won't surprise anyone to know that nearly half of the pie chart is blue. But "big blue" in this case does not refer to IBM. It's ISRG, the Internet Security Research Group, better known by the name of its wildly popular service: Let's Encrypt. That's right. The free DV - Domain Validation - certificates being produced by the automated ACME protocol today account for 46.52% of all non-expired certs. So very nearly half. In second place holding onto nearly half of that is my CA, DigiCert, with nearly a quarter share of the global market at 22.19% percent. And reducing by half again, is Sectigo with 11.89%. Google Trust Services has 8.88%, GoDaddy is represented with 5.77%, Microsoft at 3.45%, and InenTrust with a bare 0.63%.

This count of 7 should bring everyone up short, since this means that the industry, in the guise of the CA/Browser forum, has been incredibly permissive about extending our global browser trust to organizations whom we really have very little **need** to trust. Yet today we're inherently trusting the signatures of certificates that most of us are never going to see. And, of course, with that unnecessary trust comes vulnerability... and it's about to become significantly worse.

Here's what Ryan has to say about this. He writes:

There are between 75 and 85 organizations in the various root programs constituting the [entire] WebPKI that can issue certificates for the entire web. If we use the higher estimate [of 85, the addition of the EU's 43 member countries]~~, this~~ represents an increase of over 50% in the number of organizations that are trusted for this.

*Why is all this significant? While it is true that there are numerous CAs in the WebPKI beyond the seven mission-critical ones, each additional CA represents an increased surface area for **all** users of the web. The "long-tail" CAs, those lesser-relied upon entities, are part of the WebPKI because they ostensibly meet the same objective technical and procedural standards as their more prominent counterparts.*

In other words, we trust all of those essentially unneeded CAs because the way the system has evolved, it would be considered rude not to give someone the benefit of the doubt and trust their work unless and until they give the world reasonable cause not to. However, it's also likely that with the ISRG's Let's Encrypt having changed the rules, no one in their right mind today would attempt to establish a new commercial CA. Given today's startling distribution of signed web certificates my feeling is that we ought to be running in the exact opposite direction than what the EU proposes. If IdenTrust's 0.63% was also eliminated – presumably also with minimal impact – we could reduce CA trust to just 6 well proven certificate authorities. That sure seems more like the future as opposed to adding 43 new and highly political trust roots. Ryan continues, writing:

> *Web browsers set these standards to participate in their programs, striving for objectivity, openness, and consistency. This approach not only keeps the web open and fosters the development of sovereign digital capabilities in various countries but also involves a balancing act: mitigating the risks associated with the expanded attack surface that each new CA introduces.*
>
> *[What the EU proposes] ~~This simply~~ tips the scales of this system by lowering the bar for European CAs allowing them to meet a lower standard while at the same time putting these governments in charge of which CAs meet that bar.*
>
> *To put this in context, consider this case in 2013 where a French agency that was allowed into the web PKI was caught minting SSL certificates that impersonated major sites like Google (https://arstechnica.com/information-technology/2013/12/french-agency-caught-minting-ssl-certificates-impersonating-google/). Putting 27 governments in a position to add more CAs that are trusted by the world means they can do this at scale if they decide to do so, or an attacker uses these governments' ability to do so for their own benefit.*
>
> *It is worth noting that the browsers distrusted this CA when it did this. In this new world **that won't be possible** - more on that later.*
>
> *Now consider for a moment that there are 195 sovereign nations in this world - for now. The 27 member states of the EU will be the only countries in the world with the ability to force browsers to trust arbitrary CAs like this or to add their pet features. How long do we think that will last if browsers become compliant with this new legislation?*

Ryan linked to an article which appeared in ArsTechnica ten years ago in 2013. I was pleased to see that this podcast had covered every one of those incidents cited, but that was ten years ago and we have many more recent listeners. So I'm going to digress for a moment to edit down and cite a few paragraphs from that ArsTechnica piece to fill-in that past for our more recent followers. Ars wrote:

> *Rekindling concerns about the system millions of websites use to encrypt and authenticate sensitive data, Google caught a French governmental agency spoofing digital certificates for several Google domains.*
>
> *The SSL credentials were digitally signed by a valid certificate authority. In fact, the certificates were unauthorized duplicates that were issued in violation of rules established by browser manufacturers and certificate authority services.*

> *The certificates were issued by the French cyberdefense agency better known as ANSSI. After Google brought the certificates to the attention of agency officials, the officials said the certificate was used in a commercial device on a private network to inspect encrypted traffic with the knowledge of end users. Google updated its Chrome browser to reject all certificates signed by the intermediate authority and asked other browser makers to do the same. Firefox developer Mozilla and Microsoft followed suit. ANSSI later blamed the mistake on human error. It said it had no security consequences for the French administration or the general public, but the agency has revoked the certificate anyway.*
>
> *An intermediate certificate authority is a crucial link in the "chain of trust" that's key in connections protected by SSL and its successor TLS. Because intermediate certificates are signed by a root certificate embedded in the browser, they have the ability to mint an unlimited number of digital certificates for virtually any site. The individual certificates will be accepted by default by most browsers. The issuance of an intermediate certificate that spoofs certificates for domains of Google or other third-party websites is a significant breach of protocol. It could represent a major threat if one of the individual certificates—or, worse, the intermediate certificate—were ever to fall into the wrong hands.*
>
> *The incident is only the latest to underscore gaping vulnerabilities in the SSL system. In early 2012, critics called for the ouster of **Trustwave** as a trusted issuer of SSL certificates after the security firm admitted to minting a credential a customer used to impersonate websites it didn't own. In both cases, the unauthorized certificates were used to help network operators inspect the encrypted traffic flowing over their systems. While the holders of these unauthorized certificates didn't intend to use them to attack Internet users, critics have slammed the practice because it has the potential to harm third-party bystanders.*
>
> *Even more worrisome are actual security breaches on certificate authorities that on at least one occasion have allowed attackers to create counterfeit credentials used to compromise third-party Web services. That's precisely what happened in 2011, when security researchers spotted a bogus certificate for Google.com that gave attackers the ability to impersonate the website's mail service and other offerings. The counterfeit certificate was minted after attackers pierced the security of Netherlands-based **DigiNotar** and gained control of its certificate-issuing systems. Within a few days of the discovery, most of the major Web browsers issued updates to block the certificate, but not before some 300,000 people, many located in Iran, had been exposed to the certificate as they accessed Gmail servers.*
>
> *Taken together, the incidents highlight one of the key shortcomings of the SSL system. Despite its importance to millions of online banks, e-commerce services, and other sites, the entire system can be undermined by a single point of failure. With hundreds of issuers trusted by the typical browser, all it takes is for one of them to be compromised, somehow.*

And now we learn that we probably only ***really*** need to trust six or seven CAs to obtain trust coverage of 99.32% of the entire web. It's clear that we're heading in the wrong direction.

So, next, Ryan makes that point that the EU's legislation *"... requires browsers to have a user interface to support the display of legal identity associated with a website."* Ryan writes:

> *Extended Validation certificates, once used by about 8-10% of websites, now represent only about 3.8% of all certificates on the web. (https://www.ssllabs.com/ssl-pulse/).* [Ryan cites the statistic from SSL Labs' ssl-pulse page.] *WebPKI CAs originally marketed these as tools for*

> *increasing conversion rates, among other supposed benefits, but there was never any data supporting these claims. Over time, it became apparent that they provided little to no value and, in some cases, even harmed users. A notable example of the confusion arising from this paradigm is a case where a security researcher demonstrated the ability to quickly and inexpensively create a legitimate company with a name conflicting with a well-known organization, without needing to reveal their identity.*

For those interested, the researcher was a guy named Ian Carroll. Ian filed the necessary paperwork to incorporate a business called Stripe Inc. He then used the legal entity to apply for and receive an EV certificate to authenticate the Web page https://stripe.ian.sh/. Of course, Ian was unable to get “stripe.com” because the read Stripe owned that domain. But creating a “stripe.com” sub-domain under his own “ian.sh” domain was sufficient. This was because at the height of EV certificate usage, the domain’s EV certificate details would be shown instead of that messy http URL. So what visitors to Ian’s demo site saw was simply: “Stripe, inc.” And I’ll note that this followed three months after a different researcher, James Burton, established a valid business entity named "Identity Verified" to demonstrate how the resulting EV certificate might be used to add the air of authenticity to a scam site. The bottom line was that since typical users don’t actually have any idea what’s going on, all of this extra special ’ness was abandoned. Or, as Ryan puts it:

> *This incident, along with several others and research based on large-scale analysis of user reliance on browser trust indicators, led to the de-emphasis of these affordances in the browser UI. The previous UI, which highlighted this information, was redesigned and demoted in the visual hierarchy, setting it on a path for a likely eventual removal as a result.*

Which is were we are today. But the legislation that is poised to become law in the EU, after several years of the industry warning against all of this in the strongest possible terms, **requires** browsers to bring this back and for the EU to be able to add their own identity assertions to the browser’s location bar display.

Another point of serious concern is that the EU’s forthcoming legislation explicitly and deliberately limits the ability of browsers to protect users from poor-performing EU certificate authorities. He were just reminded of several instances where it was crucial that browser have the freedom to do what they feel is right for their users. It’s no exaggeration to say that we depend upon our browsers to have our backs in countless ways. Ryan writes:

> *Today CAs are removed as trusted for a vast range of reasons, for example last year, a Turkish CA, E-Tugra, demonstrated they lacked the most basic security practices and could not effectively respond to a security incident and were distrusted as a result. Not due to having made any mistake, but because their service was clearly shown to be unconscionably insecure.*
>
> *Under this new legislation, browsers will no longer have the ability to distrust European CAs that are trusted for QWACs except for “breaches” and “loss of integrity of an identified certificate” whatever that means??*
>
> *Each of the CAs trusted within the WebPKI represents a risk to users, this is why it is so important that browsers, acting as the agents of their users, are empowered to establish*

> *uniform criteria to ensure all the CAs meet minimum best practices and have the ability to remove them when those minimum best practices are not met. The text reduces the cases substantially in which they may do that. Unfortunately, it gets worse.*

Ryan then cites some of the legislation that will take effect:

> *"Web browsers may take precautionary measures related to a certificate or set of certificates in case of substantiated concerns related to breaches of security or loss of integrity of an identified certificate. When such measures are taken, the browsers must notify their concerns in writing without undue delay, along with a description of the measures taken to mitigate those concerns. This notification should be made to the Commission, the competent supervisory authority, the entity to whom the certificate was issued, and the qualified trust service provider that issued the certificate or set of certificates. Upon receipt of such notification, the competent supervisory authority is expected to issue an acknowledgement of receipt to the web-browser in question."*

Ryan replies to this heavy handed nonsense, writing:

> *Today, no notification is required to the "the entity to whom the certificate was issued" because every CA is required to use Certificate Transparency which means that that entity and everyone else can keep an eye out for mis-issued certificates. However, there is no such requirement for the EU trust-list approved CAs, and the language in the legislation prevents the browsers from adding such a requirement when they are forced to trust them.*

"Certificate Transparency" is a fabulous system initiated by Google about ten years ago. It's use has become a requirement for CAs to be trusted by today's web browsers. This requirement means that certificate issuance is no longer private between a CA and its client. Certificates cannot be issued secretly. CAs must and do publish notification into global certificate logs which are rapidly searchable, cryptographic, append-only ledgers of certificates. Because these logs are distributed and independent, anyone can query them to see what certificates have been included in the log, and when. Unfortunately, the EU's language is being interpreted to mean that they will be operating independent of the world's now-established certificate transparency system. The legislation contains language with phrases such as:

"... shall not be subject to any mandatory requirements other than the requirements laid down earlier." (Which say nothing about the adoption of certificate transparency) And *"... shall not take any measures contrary to their obligations set out in Article 45."* (Again, the EU is flatly asserting absolute authority over the trust that browsers will place in any certificates issued by their member states. And the text also says that even in the cases where there were *"breaches of security or loss of integrity of an identified certificate"* the EU can override the browsers and force them to trust the associated CA anyway: *"When the outcome of an investigation does not result in the withdrawal of the qualified status of the certificate or certificates the supervisory authority shall inform the web-browser accordingly and request it to put an end to the precautionary measures referred to."*

The more time I’ve spent looking into this, the worse it seems. The world has spent a great deal of time slowly and carefully evolving an equitable system of trust. And now, for essentially commercial reasons, to force the display of website digital identity through the equivalent of their own system of EV certs, this legislation would force all web browsers to accept root certificates from every EU member state, which would then use them to assert the identity of anything they chose ... and there’s nothing any browser could do about it.

What I’m most wondering now is what gives the EU the right to dictate the operation of our web browsers? To me this seems like uncharted waters. Users currently have some say over the certificates which populate their root stores. If they wish to remove trust from some certificate authority, nothing prevents them from doing so. But the EU is stating that browsers will be required to honor these new, unproven and untested certificate authorities and thus any certificates they issue, without exception and without recourse. Does that mean that my instance of Firefox will be legally bound to refuse my attempt to remove those certificates?

If the EU wants to create their own “EU Browser” based upon a fork of Chromium, embellish it with their own certificates and a user interface display of whatever those certificates wish to assert... then require that their own citizens use it, the only people who would have any problem with that would be their own citizens – who could then decide whether they want to keep those legislators in office. To me, that appears to be the only feasible course of action.

What’s completely unclear, and what I haven’t encountered anywhere, is an explanation of the authority by which the EU imagines it’s able to dictate the design of other organization’s software. Because that’s what this comes down to. The UK tried to do this with end-to-end encryption. Every last publisher of that technology said no and the UK blinked.

Edge and Chrome on Windows obtain their root stores from Windows. So the EU is telling Microsoft that they must add and unilaterally trust 43 new root certificates to their operating system’s root? And what about Linux? Who’s going to make Linux do this? Good luck sneaking this past Linus! That’s never going to happen.

And then there’s the very real specter of what other doors this opens: If the EU shows the rest of the world that it can successfully dictate the terms of trust for the independent web browsers used by its citizens, what other countries will follow with similar laws? Now everyone gets to simply require that their own country’s certificates get added. This takes us in exactly the wrong direction.

None of this is good. And it’s not even as if there’s some actual problem that needs to be solved here. The more I think about it, the more I like the idea of disabling Firefox’s newly added “trust the certs in the underlying operating system’s root store” option – which was just added in Firefox 120 – then pruning all but six of Firefox’s current root certificates. That trusts 99% of the Internet's certificates, and likely 100% of any certificates that I would choose to trust.