Deep Dive Into Digital Privacy and Suggestions For a Maximally Private System

Mikołaj Z. Wysocki, Kuba B. Herka, Wojciech Kowal
08.2025

1. Introduction

Privacy is one of the most cherished wants of human beings in the current age. It's difficult to provide a consistent definition; Thomas Cooley defines it as "right to enjoy life and be left alone". Unfortunately, in the current day and age, we are gradually losing our privacy due to increased usage of the internet and regulations requiring more data from customers. For example, the recent UK Online Safety Act requires users to verify themselves, using their IDs, before accessing certain adult-deemed content, which could include political themes and discussions. In some sense, we traded privacy for money and our data has become a commodity. That's partially the reason we decided to create this document. We wish to educate people and provide them with ways of increasing their privacy on the internet. To achieve this goal, we created a privacy system that combined four different domains under which we could increase our privacy. Those domains are: Operating System, Network, Payments, and Services. To make this system usable, we introduced two key assumptions, which are internet access and ease of use. We also decided to introduce some variety by providing two distinct versions. The first one will try to maximise privacy even at a loss of "ease of use" and the second one will try to find a compromise between these two ideas.

2. Table of contents

3. Operating systems

Windows and Linux in terms of privacy

Windows, as the dominant operating system, has long been criticized for its invasive data collection practices. Since Windows 10, Microsoft has implemented extensive telemetry systems that collect data about your usage patterns, installed applications, and even your

activity. While some of this can be disabled, much of it cannot, leaving users with little control over what data is being sent to Microsoft servers. This is compounded by the fact that Windows is proprietary software, meaning its source code is closed and cannot be audited by the public. Users are forced to trust Microsoft's claims about what the system does with their data, with no way to verify them.

Adding to this, Windows 10 is approaching its end of life in October 2025, meaning Microsoft will stop providing updates and security patches. This will render millions of older PCs vulnerable to security threats, effectively forcing users to either upgrade to Windows 11 (which has stricter hardware requirements) or abandon their devices. This planned obsolescence is not only wasteful but also coercive, pushing users toward newer systems that continue the same invasive practices.

Privacy benefits of Linux-based operating systems

The most significant characteristic of Linux-based systems is their open-source nature. Users can check for themselves how the system operates and aren't forced to trust third-party corporations. An additional benefit is the vast configuration possibilities that allow for personalized privacy options. They also do not collect telemetry data by default, which differentiates them from some OSes like Windows.

Another important aspect of the Linux ecosystem is its distributions, often called "distros". These are unique operating systems that use a shared Linux kernel. They vary in terms of functionality, programs installed, and overall structure. There exists a distribution for almost every use, from day-to-day usage like Ubuntu or MintOS to privacy-focused ones like Tails (The Amnesic Incognito Live System) OS.

The Tails OS is a system consisting of several tools aimed at anonymity of its users. It is created to boot from a USB stick or previously a DVD, and all files are deleted after a system reboot. You can still save files through, for example, another removable media like a USB stick. The Tails OS also has a tool called "Nautilus Wipe" which ensures a thorough deletion of the files, because a lot of Operating systems only remove the file name and the pointer to the place in the system's memory where the file resides. Not to leave any digital footprint, all the traffic is conducted through the TOR Network. Tails is equipped with an array of cryptographic tools to encrypt files, emails, and other sources of data. It also comes with software to create a virtual keyboard to protect its users from hardware keyloggers. It was created entirely using free software, because the proprietary software isn't trustworthy. As a side note, I'll add that it also used to come with a camouflage that would make the OS look like Windows 8 to avoid attracting unwanted attention when working in public places.

From the ease of use perspective, it's also worth mentioning that the Linux-based systems offer the possibility to work on older devices with limited resources due to their light-weight nature. Because of the open-source nature of Linux, all the distros are free. A big issue, a lot of users switching to Linux face it lack of support for many commonly used applications like Adobe Creative Suite. There are some applications, however, that try to replace the unavailable software, but many of them are insufficient and lack some functionalities of the original apps.

Ways of Using Windows Privately

If the use of Windows is required, there are a few ways of doing it securely. If Windows needs to be the main operating system on the machine, we recommend minimising telemetry data sent as much as possible. If that's not necessary, you could employ Windows Virtual Machines, use a separate drive specifically for it, or perform a dual-boot, which is not recommended because it can lead to issues like Windows updates interfering with the dual-boot setup.

Custom ROMs

Google Play Services, deeply integrated into stock Android, operate with elevated system privileges that go far beyond those of regular apps. While technically sandboxed, their status as system-level software grants them access to sensitive APIs and resources, making them indispensable for many features like location tracking, push notifications, and account synchronization. However, this level of access comes with significant drawbacks, like extensive data collection that includes location, app usage, and device activity. This level of privileges also poses a threat, because in case of a successful compromise, the attacker would get access to sensitive system-level data.

Unfortunately, it is pretty difficult to switch to an alternative in the Android ecosystem because many apps rely on Google Play Services for their functionality. For users seeking greater privacy, security, and control, custom ROMs offer an alternative by allowing the removal or replacement of Google Play Services with solutions like microG, which we will talk about in later chapters. A custom ROM is a modified version of the Android operating system developed by independent developers. They can be installed on older devices that aren't supported by the newest versions of Android. This ensures the latest features and security patches. Most of them were created with privacy in mind and do not contain telemetry data collection, and give their users control over installed services, which include those from Google.

There are also some downsides to using custom ROMs. There are some devices that aren't supported by the custom ROM Maintainers. Some applications that require strict security measures may not work on them. The last one is the knowledge required to set up a system like this. This may be quite intimidating for beginners, and they may even brick the phone (render it unusable).

A few popular options include LineageOS, which is known for its stability, privacy, and device support; crDroid, which offers extensive customization and performance tweaks; and GrapheneOS, which is especially interesting because of its sandboxing capabilities regarding Google Play Services. Thanks to this approach, a lot of Google Services no longer have root access to the device, which makes a few of the issues mentioned previously obsolete.

IOS Ecosystem Misconceptions

iOS takes a fundamentally different approach compared to Android and custom ROMs. As a closed-source operating system, it limits user control and customization but centralizes updates and security under Apple's control. Apple's business model, focused on hardware and services rather than advertising, reduces its reliance on extensive data collection. The argument can be made that Apple is less privacy-intrusive than Google, as its revenue is not solely tied to targeted advertising.

However, Apple's partnership with Google, particularly its deal to make Google the default search engine on iOS, complicates the picture. While Apple itself is unlikely to share user data with Google, Google can still collect data from iPhone users through its apps (e.g., YouTube, Gmail, Google Maps) and services (e.g., Google Search). This means that even on iOS, users are not entirely shielded from Google's data collection practices if they rely on its ecosystem.

Unlike custom ROMs, iOS does not allow users to remove or replace core system components, making it less flexible for those seeking full control over their devices. While its centralized design can simplify security and privacy management, it also locks users into Apple's ecosystem, leaving little room for alternatives.

4. Network

VPN

A VPN is a set of technologies that encrypt and tunnel your internet traffic. It does it by creating a virtual, encrypted tunnel between two devices. This makes it impossible for unauthorized entities to listen to our traffic. There are many types of VPNs, but in this document, we are going to focus on the 3 most relevant, in our opinion. Those are site-to-site VPNs, remote-access VPNs, self-hosted VPNs, and consumer VPNs.

Site-to-Site VPNs

Site-to-Site VPNs work by securely connecting entire networks, like branch offices to a central headquarters. These are generally used by organisations and aren't meant for day-to-day, private use. They use protocols, such as IPsec and IKEv2, to establish secure tunnels between routers or other dedicated VPN gateways. First, the gateways authenticate each other using certificates or pre-shared keys. Then the data packets are encrypted with algorithms like AES. Next, the encrypted packets are sent over the Internet to the receiving gateway. After arriving at the intended gateway, the data is then decrypted and forwarded to the destination network.

Remote-Access VPNs

Remote-Access VPNs are used mainly for remote work. It allows workers to securely connect to the company network over the Internet. They work similarly to the Site-to-Site VPNs, but they use different protocols, like OpenVPN, WireGuard, and do not require both parties to have a dedicated VPN gateway, only a specific application. First, the user has to authenticate with credentials or some form of multi-factor authentication. Then a secure tunnel is established, and once connected, the user's device behaves as though it's part of the private network, allowing access to internal resources. Remote-Access VPNs also allow for additional features like "Split tunneling", which means that only corporate traffic goes through the VPN and the rest goes through the Internet, and "Kill-switch" that automatically disconnects internet access if the VPN connection drops.

Self-hosted VPN server

Creating your own VPN server is another possibility for users who value privacy and control. This option removes the need to rely on and trust third-party vendors. However, it

requires significant technical knowledge and monetary investment. Because of that, it does not meet our agreed-upon criteria of ease of use, so we won't pursue this idea any further.

Consumer VPN Services

Consumer VPN Services are designed for private use. They enhance individuals' privacy while accessing the Internet. Before establishing a connection, users are required to download a VPN client on their devices. Once installed, the client connects to one of the provider's servers. Similar to previously mentioned VPNs, this one also encrypts data using protocols like OpenVPN, IKEv2/IPsec, or WireGuard. Then the data is forwarded to the destination address, but it appears as if it originates from the provider's server and not the user's device. This is the case because the provider acts as an intermediary, offering a proxy-like functionality.

Pros and Cons of VPN usage

Because the first two types of VPN are purely for corporate use, and the self-hosted VPN doesn't meet the ease-of-use criteria, we will focus on the Consumer VPN Services for the purpose of this document.

VPN usage comes with significant benefits to our privacy, by masking our internet traffic using encryption and bypassing geo-restrictions that allow users to access content unavailable in their region. However, they do not come without downsides. VPN services slow down our internet connection to some degree, and many of them are not as private as they claim they are. Some VPN providers are compelled by their respective governments to still collect data from users, and many of them also do it out of their own accord. For example, there were some controversies with IPVanish and PureVPN that showed that they did log user data and hand it over to the authorities despite claiming otherwise.

Current encryption methods face a long-term threat from quantum computing. Adversaries may be engaging in a strategy known as "Harvest Now, Decrypt Later", where encrypted data is harvested today with the intention of decrypting it once quantum computers become powerful enough to break traditional algorithms. To address this risk, some providers, like Mullvad, are adopting post-quantum cryptography protocols [61]. These methods are designed to resist quantum attacks, ensuring that data encrypted today remains secure even in the future. Generally speaking these techniques are categorized into a few families: [63], such as:

- **Lattice-based Cryptography**: These algorithms are based on hard mathematics, lattice, worst-case problems. [64]
- **Hash-based Digital Signatures**: Because of the nature of the hash functions, it is hard to crack them even for quantum computers, and these algorithms use that fact to create a secure digital signature technique.
- **Code-based Cryptography**: "It is based on error-correcting code. Error correction codes are codes used widely in communications to correct transmission errors. To send a message, the text is sent into an error correction code. Then to the output, a few errors are randomly introduced and sent."[63]
- **Multivariate Cryptography**: It employs a mathematics problem of solving a system of multivariate polynomials. This problem is considered a hard problem even for quantum computers. [63]

Before going deeper into specific implementations, its also worth mentioning that classic symmetric algorithms are fairly resistant to quantum computing, while using larger key-sizes like 256 bytes. Although Grover's quantum algorithm can still hurt their security. There is a lot of post-quantum algorithms, but we will focus only on ML-KEM and Classic McEliece, because those two are used by Mullvad VPN for key exchange. Regularly the key exchange is performed using an asymmetric key algorithm that isn't quantum resistant, like Diffie-Hellman. They aren't used for payload encryption, because of their bad performance, caused by the need for bigger keys. In Mullvad VPN, the payload is encrypted using the ChaCha20-Poly1305 algorithm, which is quantum resistant, because of its symmetric nature.

ML-KEM is a member of "Lattice-based Cryptography" family and is based on the presumed hardness of the so-called Module Learning with Errors (MLWE) problem, which is a generalization of the Learning With Errors (LWE) problem. LWE is based on the idea of representing secret information as a set of equations with errors. In other words, LWE is a way to hide the value of a secret by introducing noise to it. [65].

Classic McEliece is a member of "Code-based Cryptography" family. Its an asymmetric algorithm that is still safe, because it uses error-correction code as keys. Correctly working code is used in place of private key and the defective code is used as the public key. [63]

Mullvad allows user to choose, if he prefers to use one or both of these algorithms. This makes Mullvad a well-suited choice for users concerned about long-term data security.

Choosing a VPN

Choosing a VPN is also an important factor in terms of the level of our privacy. As we showed before, some VPN providers have lower standards for the privacy of their users. While picking a VPN Service, we should take into consideration the jurisdiction of the particular VPN Service. There are alliances between countries like:

- Five Eyes - Australia, Canada, New Zealand, the United Kingdom, and the USA
- Nine Eyes - The Five Eyes plus the Netherlands, Norway, Denmark, and France
- Fourteen Eyes - The Nine Eyes plus Italy, Germany, Belgium, Sweden, and Spain

that we want to avoid, because of their proclivity for data-sharing agreements. Another thing to keep in mind is whether the VPN service is regularly undergoing third-party audits. A lot of reputable VPNs publish transparency reports and have open-source apps to build trust. The next important thing to take into consideration while picking a VPN is the differences between free and paid services. There are some speculations that we want to address, that free VPN services sell user data to third parties to cover the operating costs. On the other hand, paid VPNs are generally more reliable and privacy-focused. They generally collect user data at lower rates, sometimes without even collecting any at all. Some of them go as far as to just create an account number during registration with no need for an email address. Another thing to take into consideration is the payment methods. Many reputable VPNs offer payment methods that enhance user anonymity, such as cryptocurrency payments, gift cards, and cash payments in the case of Mullvad VPN. These options make it harder to link your identity to your VPN account.

Taking everything we just talked about into consideration, it's quite clear that the Mullvad VPN is one of, if not the best, option available. It provides complete transparency, cash payments, quantum-resistant encryption, and supports regular third-party audits.

The Dark Web

Another domain in the topic of confidentiality that we would like to address is the "Deep Web" and especially the "Dark Web". It is no secret that the Dark Web is home to a lot of criminal activities. There are a few other groups of people who put as much consideration into their confidentiality on the internet as criminals do. That's why we decided to take a bit from the entire dark web ecosystem.

First of all, let's clarify what the "Deep Web" and the "Dark Web" are. The Internet that we know of and that we have access to through classic web browsers like "Google", "Microsoft Edge", and "Firefox" is called the "Surface Web". That is an indexed portion of the internet that we can easily access, which consists of only around 4% of the internet. There is another 96% that is unindexed, and therefore, we cannot access it with our classic web browsers. That portion is called the "Deep Web"; it hosts things like academic records, government records, and financial records. Around 6% of the "Deep Web", where mostly criminal activities take place, is called the "Dark Web". That's the place where criminals conduct their illegal activities, like selling drugs, weapons, people, and ordering murders. It's important to point out that the Dark Web isn't all bad. It was created with good intentions - to give people anonymity and a platform for free expression, for example, for people from countries where free speech isn't a standard. Because of the anonymous nature of the Dark Web, almost all the transactions are conducted using cryptocurrencies, but we will talk about that in the later chapters.

The Onion Routing

But how do we access the Dark Web? We are going to need something called the "The Onion Routing (TOR) Browser". The TOR concept goes back to the mid-1990s, when the lack of security on the internet and its ability to be used for tracking and surveillance were becoming clear. This motivated the U.S. Naval Research Lab (NRL) Researchers and later MIT graduates to create TOR and later in 2008, the TOR Browser.

But how does TOR work? It encapsulates your standard TLS traffic in 3 layers of encryption with three separate symmetric keys. That's what we call a "cell", which is a basic unit of communication on the TOR channel. The cell is then sent to the recipient through 3 intermediary proxy servers called relays or nodes. Our Client PC has all three sets of symmetric keys, and each node has a different key; therefore, each node can decrypt only one layer of encryption from our traffic. It's like peeling an onion off its layers. Exchange of the symmetric keys is performed using the Diffie-Hellman algorithm. Most cells are of fixed length and therefore none of the nodes knows how many layers it has left and henceforth, which node in order they are. TOR always has 3 relay points called Guard Node, Relay Node, and Exit Node.

- **Guard Node:** The first intermediary, also called an entry node. This node knows who the sender is and where the second relay point is, but doesn't know the final destination
- **Relay Node:** The second intermediary, which only knows where the first and last proxies are. All the traffic going through it is encrypted.
- **Exit Node:** The last intermediary that knows the location of the second proxy and the destination server. It decrypts traffic with the last key and sends it to the destination server in the form of TLS traffic, or in the older versions, plaintext.

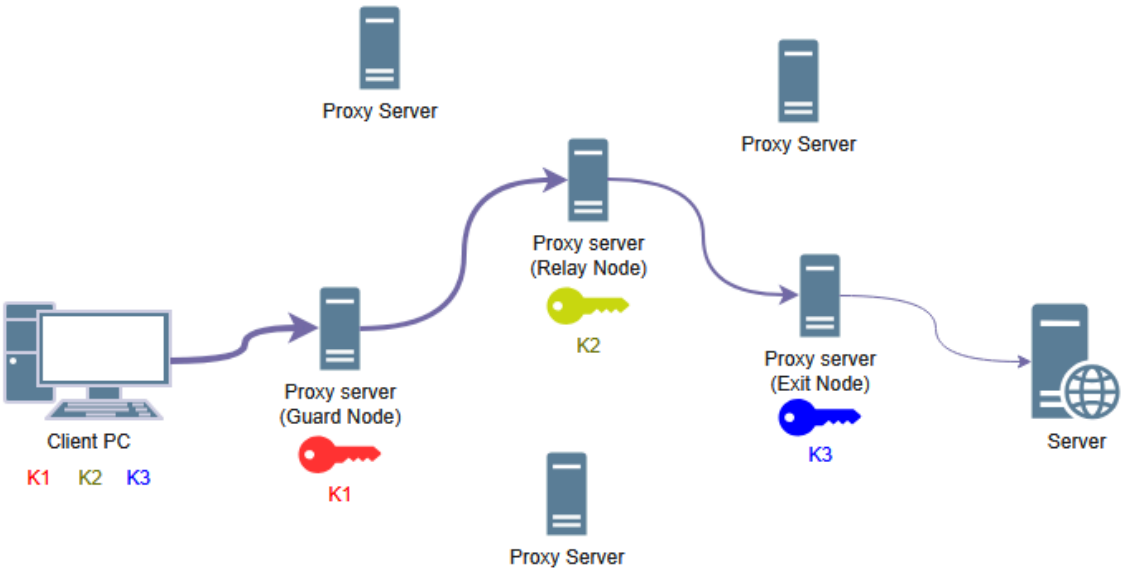

TOR Browser

Now that we have a clearer picture of how the TOR works, we can talk about the TOR Browser. The TOR Browser is made of 3 elements: Modified Firefox ESR, TOR Network, and Privacy Extensions.

- **Modified Firefox ESR:** The Tor Browser is built on a customized, in terms of enhanced privacy, version of Firefox ESR. Because of that, in the TOR Browser files, there are the same database files as in Firefox. That includes the files that hold browsing history, but they aren't updated and do not keep track of or store any browsing information.
- **TOR Network:** This is a network made of relay servers that we talked about in the previous chapter. It is responsible for anonymity.
- **Privacy Extensions:**
  - **NoScript:** Blocks JavaScript to prevent potential security issues. Users can create whitelists of servers and selectively allow JS.
  - **HTTPS Everywhere:** Is an extension that enforces HTTPS when possible.
  - **Tor Button:** Enables TOR connection modification and viewing current settings

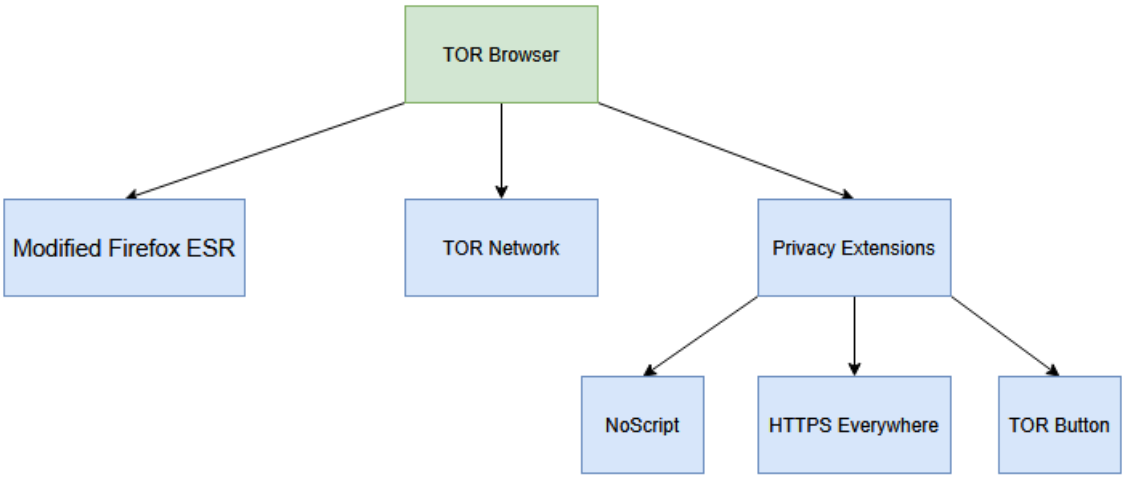

VPN

With the functionality of the TOR Browser covered, we have to address one more confidentiality issue in this model. The content of the first connection made with the Guard Node, of course, is not visible, but your ISP (Internet Service Provider) can see that you connected to the TOR Network. It is something that hinders our confidentiality. In 2013, a group of FBI agents managed to track a culprit who made bomb threats over TOR to Harvard University. Turned out, there was only one host that connected to the TOR Network at that time from the campus wifi, and this person later confessed to the crime. [14]

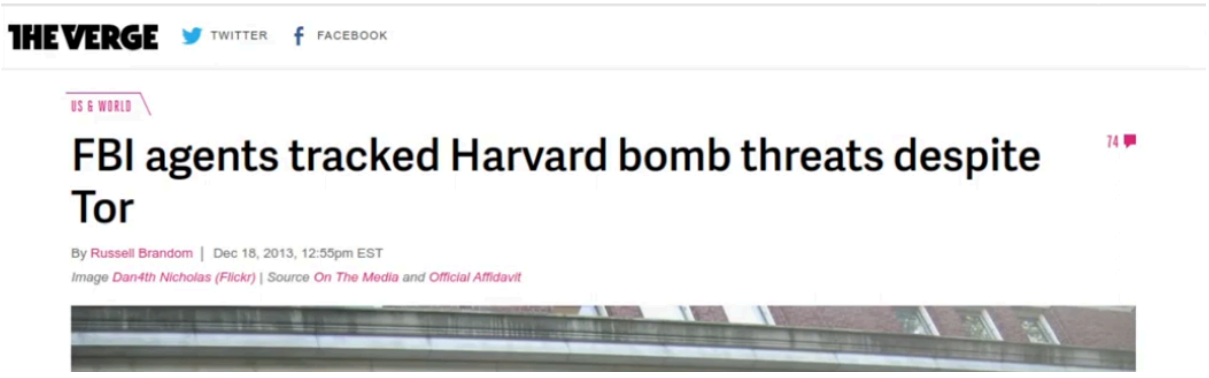
THE VERGE TWITTER FACEBOOK

US & WORLD

FBI agents tracked Harvard bomb threats despite Tor

74

By Russell Brandom | Dec 18, 2013, 12:55pm EST

Image Dan4th Nicholas (Flickr) | Source On The Media and Official Affidavit

In order to hide ourselves from our ISP and manage this confidentiality issue, we can use a VPN. Our ISP will still see that we connect somewhere, but the VPN will successfully prevent it from seeing where we are connecting to.

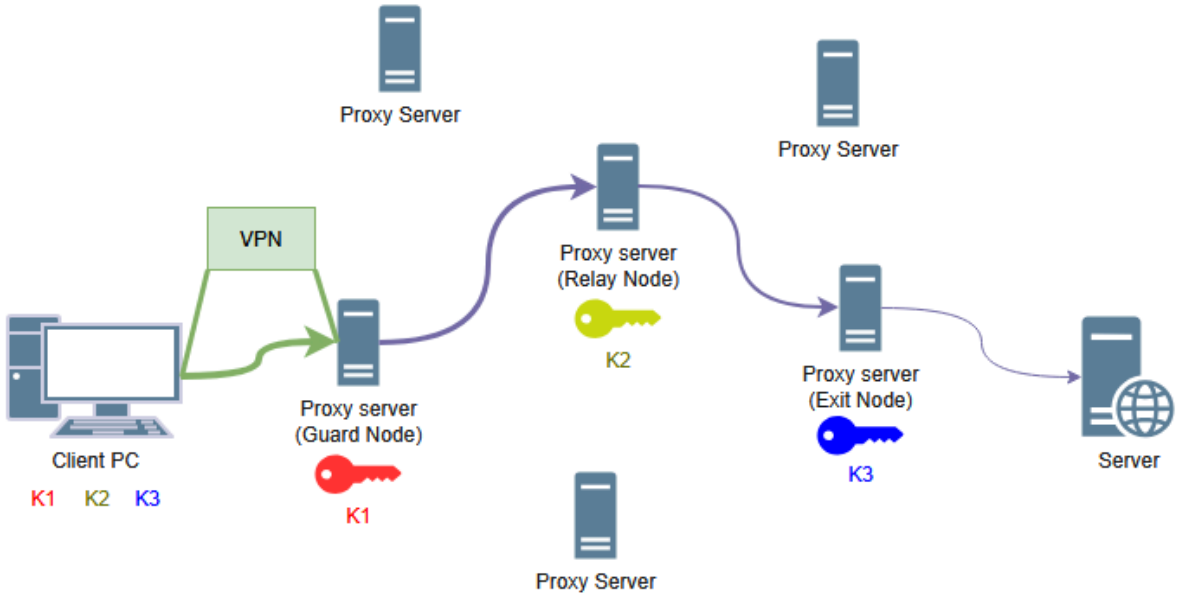

Host Issues

Now that the connection is fairly confidential, we are going to look into the problems connected to our operating system. TOR Browser, like any other browser, leaves a trace on our system in the form of various artifacts that are in our system's memory even after uninstalling the TOR Browser. Those artifacts can be accessed through means of memory forensics and registry forensics. In an article [12], the authors conducted a memory dump when the TOR Browser was open, closed, and uninstalled. Later, they analyzed those files using HxD software and found various data, which consisted of, for example, onion URLs of visited sites and timestamps.

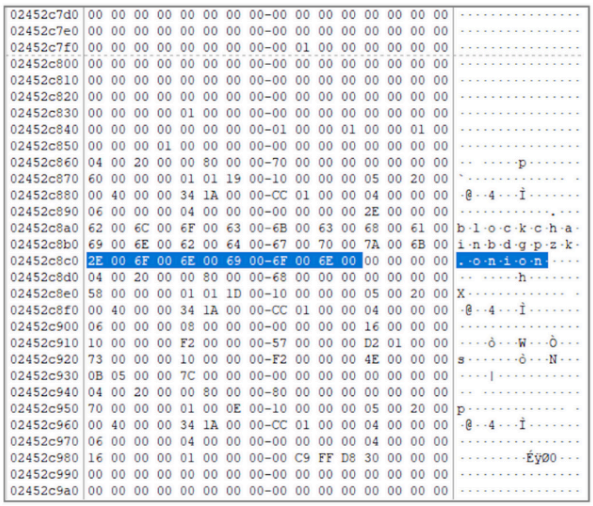

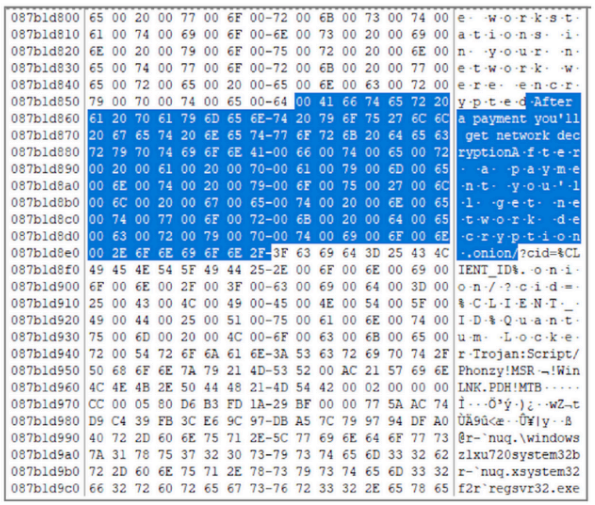

Table 2. Results obtained after memory forensics.

Tor Browser Stage	Tool Used	Visited Sites	tor.Exe	Tor Launcher	Timestamps	Registry Values
Open	FTK Imager + Bulk Extractor	Yes	Yes	Yes	Yes	Yes
Closed	FTK Imager + Bulk Extractor	Yes	Yes	Yes	Yes	No
Uninstalled	FTK Imager + Bulk Extractor	Yes	Yes	Yes	Yes	No

The authors of this article conducted more actions like registry forensics, storage forensics, and network forensics, and found quite a bit of information about our TOR Browser. This is another threat to our confidentiality that we should take care of. The most popular mitigation to that threat is using the TAILS (The Amnesic Incognito Live System) Linux OS that we covered in previous chapters.

As outlined above, TOR Browser is a great option if we want to maximize our confidentiality. We also need to mention that it comes with some downsides, because the traffic is conducted through three relay points each time, this browser is quite slow. This influences the “ease of use” aspect of the process. It is also important to mention that it is not a foolproof method, and it is possible to get to the end user through the TOR Network. It just makes it extremely hard. We also recommend using a VPN alongside the TOR Browser in order to hide ourselves from our ISPs. The safest way of using the TOR Browser, as outlined earlier, is via the Tails OS. This operating system, unfortunately, is a significant detriment to the “ease of use” aspect because of its amnesic nature, so we would recommend using it only if needed.

Browsers

A key aspect of internet activity is a web browser. Many web browsers differ in their approach to user privacy, and that is why we are going to talk about some of them. There are a few of them, like Firefox, that prioritize the ease of use, but also have features implemented for increasing the user's privacy. A step up from there, without losing too much of the ease of use, is the Brave Browser that blocks ads and trackers, automatically upgrades connections to HTTPS, and offers advanced, randomized fingerprinting protections. Brave also has a built-in Tor mode for an extra layer of privacy. Ungoogled Chromium is another privacy option that works by removing all traces of Google services from the browser, action so called degoogling, which we will be talking about in future sections.

The most private browsers that are available are Tor Browser and Mullvad Browser. These two browsers represent the gold standard for privacy and are based on the same foundation. They are both built from a hardened version of Firefox by the Tor Project and share the same advanced anti-fingerprinting technology. The former was already described, in detail, in the previous section, so we will focus on Mullvad. It’s the result of a collaboration between Mullvad VPN and the Tor Project, created with the privacy by default philosophy.

The simplest way to describe it is that it is the Tor Browser but without the Tor network. It is a fork of the Tor Browser, engineered to be used with a trustworthy VPN instead. The core part of the strategy is to make all users look identical to prevent fingerprinting. Mullvad browser also operates exclusively in a private mode. No browsing history, cookies, or site data are saved between sessions. All forms of data collection and reporting have been stripped out, and Ublock AdBlocker, which we will talk about later, is enabled by default. Similarly to Tor Browser, it provides the user with the ability to disable JavaScript. We also want to point out that it is recommended not to add extensions or change the configurations, as any modification can increase the possibility of fingerprinting. By not using the Tor Network, the Mullvad browser is significantly faster than Tor Browser, but still is slower than regular web browsers.

The reason we didn't mention Google Chrome in this chapter are numerous privacy issues related to data collection and tracking. First is the obvious connection to Google, which is known for user privacy violations. Second is the implementation of Manifest V3 that limits the functionality of ad blockers, making it harder to block ads and trackers effectively.

Search engines

Before going further, it is important to establish clear definitions for some key terms. The first concept is crawling, which is an automated process of browsing the internet to discover and index web pages. Second is indexing, which is generally defined as "collecting, parsing, and storing of data to facilitate fast and accurate information retrieval".

Privacy-focused search engines have emerged as alternatives to mainstream options like Google and Bing. These search engines prioritise user anonymity by not tracking users or creating user profiles. However, they come with their own set of advantages and limitations. They come in many variations, differing in user experience, data collection, and outside resources. The first engine we're going to cover is DuckDuckGo. It's one of the most well-known privacy-focused search engines in the world. Like other similar services, DuckDuckGo claims not to collect or store user data or profile users. The biggest downside in terms of privacy is its reliance on Microsoft Bing. Despite its reputation as a privacy-focused company, it was revealed in 2022 that DuckDuckGo did not block Microsoft trackers in its browser due to a contractual agreement tied to its use of Bing's search index [62]. An ad-based business model has its drawbacks. The company's data collection policies could be greatly influenced by the advertisers' need for better ad targeting. Next is Brave Search, which differentiates itself from DuckDuckGo by having a fully independent search index. They also offer a Brave Search Premium subscription that provides an ad-free experience, but also requires you to create an account. Brave, however, based on user experience, has slightly worse search results than Google. Another search engine we are going to mention is Startpage, which is a private way of using Google Search and Bing. It acts as an intermediary between you and other entities. Startpage submits your query to Google and Bing anonymously on your behalf, then returns the results to you. The last search engine that we will cover in this chapter is Kagi. This is the only purely subscription-based service that we discuss here. Because of this business model, they do not need to collect user data to earn profit. In order to minimize dependence on big tech companies, they also have their own web index, internally named "Teclis". A significant downside is its paid nature and limited searches on cheaper subscriptions.

5. Payments

Crypto

Another idea that's important to mention in the context of internet privacy is currency and the implications of spending it online. All transactions made online using traditional currency are recorded. So that means that banks can see all purchases we make, which clashes with the notion of maximising privacy. There are a few ways to counteract it, however, and these will be explored in this section

First, there is a pretty well-known option, cryptocurrency. Used mostly as an investment and notorious for quick get-rich money schemes, cryptocurrency can also be used as a secure and private option, with a few caveats that will be mentioned later, for customers to purchase goods and services, but before all of that, we would like to explain a few terms related to cryptocurrency.

Before diving deep into cryptocurrency, we have to talk about the idea of blockchain. Blockchain is a distributed ledger, which means that the ledger is not stored on one system, but distributed and synchronised across multiple different systems, with growing lists of records (*blocks*) that are securely linked together via cryptographic hashes. Each block contains a cryptographic hash of the previous block, a timestamp, and transaction data (generally represented as a Merkle tree, where data nodes are represented by leaves). Since each block contains information about the previous block, they effectively form a *chain*, with each additional block linking to the ones before it. The idea is to create a decentralized database. In this context, a database of cryptocurrency transactions. So what is cryptocurrency?

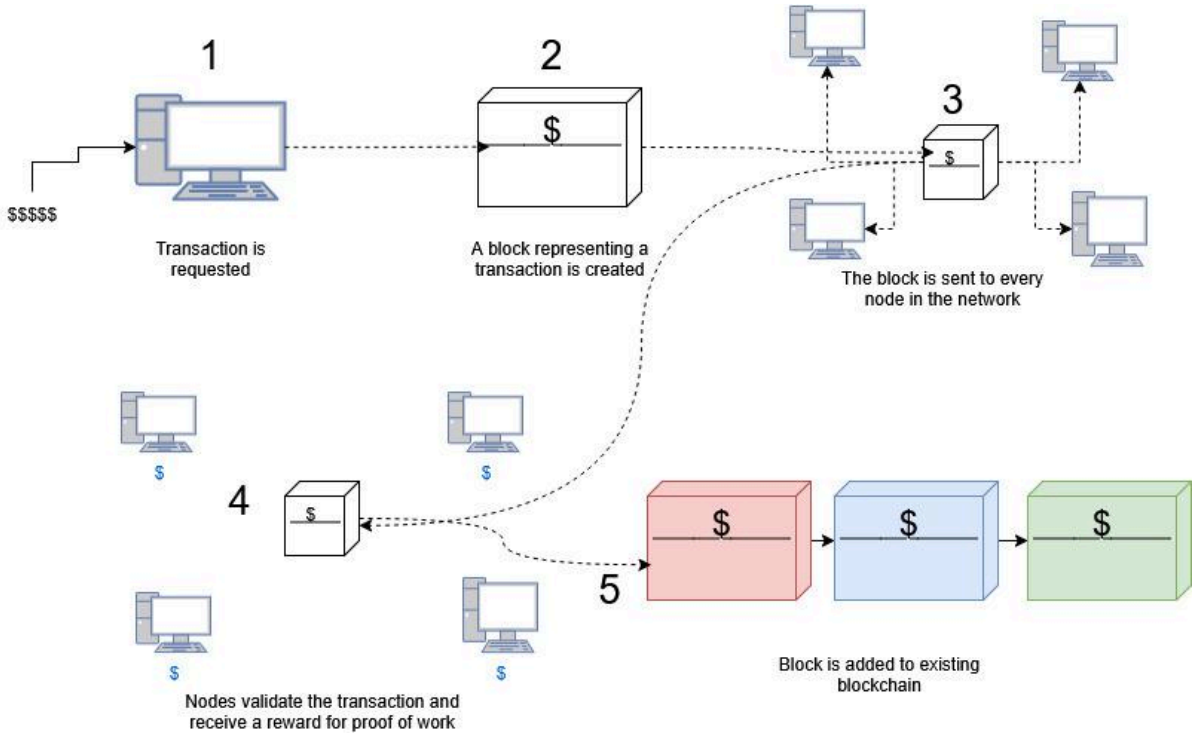

Cryptocurrency is simply digital currency that was created to work in this type of system that is not reliant on any central authority, such as a government or bank, to uphold or maintain it. The first implementation is Bitcoin, invented in 2008 by an unknown entity, Satoshi Nakamoto. Important ideas needed before talking about the privacy of cryptocurrencies are addresses, transactions, and cryptowallets. Addresses act like a bank account number; they tell you where or from where funds are being transferred, and they are created from the owner's public key. What can be found in transactions depends on the coin used; the most important ones from the perspective of privacy are the sender's and receiver's addresses and the sender's digital signature. And last are cryptowallets, which act as key managers, by which we mean they store owners' private keys, which in turn are used to create public keys.

First, we need to clean up some possible misconceptions about cryptocurrency. Most of them are not perfectly private and are considered pseudonymous. No actual names are used during transactions, but that does not mean that they can't be traced back to the persons involved in a few ways, like IP address tracking, reused addresses, and blockchain analysis tools. These issues stem from the fact that many cryptocurrencies, like Bitcoin, operate on a public ledger system, meaning that every transaction is recorded and visible to anyone. Another one is KYC or Know Your Client, which is related to anti-money laundering (AML), where for the purpose of determining the probability their customers poses money laundering risks, the customers personal data is enquired and compared to their official government-issued identification, such as a passport or state-issued driver's license, and proof of residence, such as a utility bill. The good thing is that this only applies to VASPs (Virtual Asset Service Providers) and other types of centralized exchanges.

The growing need for privacy in cryptocurrency has led to the development of specialized security controls designed to enhance user anonymity. One example is the idea of Zero-Knowledge-Proof, or ZKP for short. It is a cryptographic method by which one party can prove to another party that a given statement is true, without conveying any information apart from the fact that the statement is indeed true. The "prover" does not reveal any information about the transaction. At a high level, a zero-knowledge proof works by having the verifier ask the prover to perform a series of actions that can only be performed accurately if the prover knows the underlying information. If the prover is only guessing as to the result of these actions, then they will eventually be proven wrong by the verifier's test with a high degree of probability. This greatly hinders all types of investigations based on information present in the transaction. ZKPs can come in two forms, interactive or non-interactive. An interactive ZKP requires the prover to repeat the process for individual verifiers each time. Meanwhile, a non-interactive ZKP allows the prover to generate a proof that can be verified by anyone who has knowledge of the same proof. Three most defining characteristics of ZKP are:

•Completeness – If the information or statement provided, and the actions taken, are correct, then the verifier can determine that the prover possesses sufficient knowledge.

•Soundness – If the information provided is false and the prover is attempting to gain unauthorized access, then the verifier can easily determine that the prover does not know the correct input.

•Zero-knowledge – In this case, the prover can provide a true statement; however, this does not provide the verifier with any additional information regarding the correct input."

The most popular ZKPs include:

1. PLONK – An acronym of "permutations over Lagrange-bases for oecumenical non-interactive arguments of knowledge." PLONK is regarded as one of the most trusted and ubiquitous ZKP setups and is compatible with any program, while also able to include a large number of participants.

2. ZK-SNARKS – A "succinct non-interactive argument of knowledge" is small-scale and simple to verify. This type of ZKP generates cryptographic proof using elliptical curves and requires fewer computational resources compared to a hashing function.

3. ZK-STARKS – A "scalable transparent argument of knowledge" is a ZKP that involves a low level of interaction between the prover and the verifier, resulting in faster speed.

4. Bulletproofs – This type is a short, non-interactive zero-knowledge proof that does not need a trusted setup and is suitable for private transactions for cryptocurrencies.

Another way of securing cryptocurrency is a ring signature. It is a type of digital signature that can be performed by any member of a set of users who each have keys. It is used to obfuscate who exactly sent the transaction. The only thing the ring signature proves is that the signature came from a certain group of people, but it doesn't say from whom it came exactly. To create that signature, the system would need decoy outputs and public keys associated with them. The system then uses them in conjunction with the sender output and their key to the final signature. Now, instead of the sender's address and digital certificate in the saved transaction, only the ring signature is seen, and only the sender knows that the transaction is theirs, but there's still one issue. The receiver's address is still visible on the blockchain. To remediate this, cryptocurrencies use a technique called stealth addresses.

Stealth addresses are a technique that lets senders create a unique, one-time address for each transaction, even if the recipient uses the same base address. To create a stealth address following steps need to occur:

1. The receiver creates a private "spending key", which ensures that only they can access and use the funds sent to them.
2. Using the "spending key" receiver creates the "meta-address", which is used to signal their willingness to engage in stealth address transactions
3. Sender generates an ephemeral key, which is a single-use random number, used alongside the receiver's public key to create the Stealth Address
4. At the end, the sender includes the ephemeral key in the transaction data on the blockchain, which allows the receiver to detect that transaction as his

All of these steps allow for complete obfuscation of the receiver's address and for an additional layer of protection.

In practice, these security measures matter in terms of choosing which cryptocurrency to choose. Coins like Bitcoin offer little to no protection. Examples of so-called “privacy coins” are:

Monero (XMR): Utilizes ring signatures and stealth addresses to conceal transaction details.

Zcash (ZEC): Employs zero-knowledge proofs to enable private transactions.

Dash (DASH): Offers an optional privacy feature called PrivateSend.

MimbleWimble Extension Blocks (MWEB): A privacy feature implemented in the Litecoin blockchain, allowing users to opt-in to privacy-preserving transactions.

All of these coins provide exceptionally better privacy than standard cryptocurrency, and our pick for this project is Monero. It is one of the most recognised and respected privacy coins that boasts a really good level of data and privacy protection. It was also tested in February of 2024 by Cas Cremers, Julian Loss, and Benedikt Wagner for the security of transactions, which concluded in a positive finding for Monero. Summarizing:

No forgery or undetected double-spending without key knowledge.
Linkage detection when a key is used twice.
Robust composability even with complex interactions inside the protocol

Unfortunately, there are still many issues regarding privacy coins and crypto in general. One of the biggest ones clashing with our planned Monero usage is the legal implications of privacy coins. All of these legal hurdles revolve around Anti-Money Laundering, or AML for short. It is an international web of laws, regulations, and procedures aimed at uncovering money that has been disguised as legitimate income, and unfortunately, cryptocurrency and especially privacy coins were and still are used for just that, a convenient way to store and move money for criminals. We are already seeing the result of that. In 2023, the exchange Binance banned all privacy coins on their platform. Eu is planning to ban all privacy coins in 2027. Not only that, they started the Anti-Money Laundering Authority, AMLA for short, specifically for the purpose of ending any possibility of money laundering. It has been operational since January 1 of this year and is already focused on cryptocurrency and privacy coins, calling them “an immediate priority”. Previously mentioned KYC is also part of AML and acts as another roadblock for widespread usage of privacy coins. It is important to remember that tax fraud, which crypto could be used for, is another concern of AML.

Unfortunately, another major regulation poses a significant obstacle to the adoption of cryptocurrencies as a whole: GDPR. GDPR or General Data Protection Regulation is a European Union regulation on information privacy in the European Union (EU) and the European Economic Area (EEA), and cryptocurrency violates, as shown in this 2022 paper [59], notable rights given by this regulation, mainly ‘Right of erasure’ and ‘Right to rectification’. Before going to the specifics, let us explain what these rights mean.
‘Right of erasure’ mandates that controllers delete data in certain cases. According to Article 17, data subjects are granted the right to request the removal of all related personal data. It also means that, according to Article 6(1), it must be erased.

‘Right to rectification’ mandates that data subjects have the right to make a request to have their inaccurate personal data rectified, or completed if it is incomplete. Here, rectification means that data is updated to be accurate. These are simplified overviews of these rights, but they are crucial in understanding how cryptocurrency might make exercising them difficult.

Let’s start with ‘Right to erasure’, because it is easiest to explain. The main part of the reliability and trust of blockchains, and therefore cryptocurrency, is the integrity of individual blocks. The issue comes from the fact that if the data subjects (in this context, a person who is part of the transaction) request an erasure of their data, the whole blockchain loses consistency, which would impact the previously mentioned reliability and trust. In the paper there were few solutions to this problem, like the Hashing-out method where only hashes of data are stored on the chain and the data itself is stored off-chain or Pruning where old transactions and blocks are deleted after a predefined amount of time and only old block header with blocks hash are stored, but all of them suffer from some glaring issue, for Hashing-out, need for certain degree of control given to single-centralised party which clashes with original motivation behind blockchains and for Pruning issues like sacrificing security for the sake of scalability and privacy and adding an expensive overhead, leading further inconsistencies and scalability issues.

In the case of ‘Right to Rectification’, the issue is very similar; both rights require changes to be made to the blocks and blockchain as a whole, which reduces reliability and trust. It is also almost impossible to ensure coordination between nodes for the goal of data change, even if the data subject were to identify all the nodes or identify enough of them (over 50%) as claimed by Diogo Duarte in his paper [60].

There are also concerns regarding:
‘Right to Object’-allows individuals to stop or prevent organizations from processing their personal data under specific circumstances, which is not easy to meet for public blockchains due to their permanent nature.
Ideas of Lawfulness, Fairness, and Transparency, which are challenged by the decentralised nature of blockchain networks

Last but not least, there is the issue with the volatility of the cryptocoins. Significant changes in prices, sudden ups, followed by downs, are not rare occurrences in crypto. Cryptocurrency is still, relatively speaking, a new phenomenon. Not only is it a new market for investors, but even older cryptocoins like Bitcoin still reach a new price discovery phase in each market phase. Another thing caused by the age of crypto is the stigma related to it and the sentiment that changes after every new scandal, which are prominent in the crypto world, but why talk about this in the context of our system? Like we said at the beginning, we also try to include “ease of use” in every concept, and drastic price swings can make cryptocurrency extremely difficult to use as a form of currency.

At the end, it has to be mentioned that even if cryptocurrency and specifically privacy coins can help with achieving better privacy, issues stemming from difficulties complying with the law, resulting lack of vendors who allow payment with crypto and inherent risk with keeping capital in volatile currency make crypto and privacy coins bad option for consistent use in everyday life. That’s why we are also going to take a look at different secure and private payment options.

Virtual cards

Another digital payment method is virtual cards. These are digital versions of traditional credit cards that do not tie your bank account to every purchase. They are created under your main bank account, but can be created and deleted very easily. Services like Revolut let you create such cards in their app. The details of the main card do not get exposed to merchants. It is important to point out that your bank still sees every purchase. That means that they aren't fully private, but they still increase your overall privacy.

Prepaid cards

Another often talked about private form of payment is prepaid cards. They're a type of payment card with a preloaded balance that can be used to make purchases or pay bills, similar to a debit card. According to the Consumer Financial Protection Bureau (CFPB), "prepaid cards" can refer to more than just reloadable cards. It could also be a reference to:

- Payroll cards that employers use to deliver paychecks
- Government benefit cards used by some agencies to pay benefits like unemployment insurance
- College ID cards can also function as prepaid cards.

The advantage of them in relation to debit cards is the fact that they are not connected to consumers' bank accounts. It allows for increased privacy and a lack of tracking. But how do they work exactly?

They are similar to gift cards, where they have a certain amount of money and consumers are allowed to use it as much as they want until the money runs out. They differ from gift cards in the fact that they can be "reloaded", which means that funds can be added back to the card. They can also be used online, unlike traditional physical cash. They come in handy in many ways, but they are not the perfect solution.

First, many of them have imposed limits on how much funds can be loaded into the card. They might be applied to a specified period that may be daily or per month. There can also be a limit to how little can be loaded.

Another issue is the fees related to the cards. They come in the form of monthly fees, transaction fees, ATM fees, reloading fees, and foreign transaction fees. There's also the possibility of fees for checking the amount of funds currently on cards. This could highly limit the possibility of using them effectively.

They are also less flexible as the only form of payment. It relates to the already mentioned limits opposed by prepaid card issuers. They also need constant reloading if used as a primary payment option. Some prepaid cards might have restricted access functionality based on country or goods and services that the consumer is trying to purchase.

Last is refusal due to security measures. Some vendors automatically flag certain types of prepaid cards transactions due to the risk associated with credit card fraud. Some vendors will block cards that have no cardholder name. Another often occurring issue with fraud detection concerns recurring payments, though this can depend on the card company and the specific conditions set by the card provider.

Prepaid cards offer a practical way to enhance privacy and can help bridge the gap with online vendors that do not accept cryptocurrency.

Cash

Last but not least is the traditional way of paying for goods and services, physical cash. For day-to-day purchases, it offers good privacy because it is not tracked. The only people who are aware of the purchases are customers and shopkeepers, but even they, assuming they don't know customers' identities, only know items purchased and for how much. And if the shop is keeping track of purchases, only the aforementioned information will be documented, no PII (Personally Identifiable Information).

Unfortunately, there are downsides. We live in a society more and more focused on digital, cashless transactions. More vendors gravitate to online spaces, and even some physical stores don't accept cash as a form of payment. Using primarily cash significantly limits options available and doesn't allow us to purchase online in most cases. There is a "loophole" in a sense, where you can buy products with physical cash over long distances. It is done by, as mentioned in previous chapters, Mullvad. They allow you to buy their services by making you send them cash via mail. It allows for another layer of privacy, but this practice is not something that can be expected from most vendors.

Another limitation might be posed by government regulations. Depending on where the transaction is made and by whom, there might be a limit to how much cash can be spent before it is legally mandatory to use cashless forms of payment. In Poland, for example, the payments between businesses are limited to 15,000 Polish Zloty (which is around 3,267 euros) and no limit for private individuals, and in the Czech Republic, the universal limit is 270,000 Czech crowns per day.

6. Services

Digital independence

One of the biggest sources of privacy loss is day-to-day usage of all kinds of software. Most obvious examples are social media platforms, where millions of users daily talk about their lives, post pictures of their family and friends, but even inactive users have their data collected, for example, what kind of posts they like, which can tell a lot more about a person than most think. Unfortunately, it doesn't include just social media platforms, but email, shopping, streaming, and even digital maps. What's even worse is the fact that many of these services are controlled by a single entity, which gives them all the control over our data and what we see, through specially developed algorithms. For most ordinary people, it manifests in specifically tailored ads or content that they would most likely enjoy, but it can get

significantly more invasive. Depending on the amount of data an entity has about somebody, purely acquired from social media usage or shopping habits and tendencies, they can determine that person's gender, sexuality, and even whether they are pregnant before their family finds out. Our goal is to avoid that as much as it is possible. So, in this section, we would like to talk about all the ways to become independent from the biggest Tech companies.

First, let's discuss DeGoogling. It is a movement that advocates for consumers to stop using Google products entirely due to growing privacy concerns regarding the company. Google is one of the biggest tech companies in the world. They are currently fifth in the world and hold a $2.345T market cap, and it can be felt in almost every part of our digital lives. The most used software and services are owned by Google. They have the most used email service (Gmail), file-hosting platform (Google Drive), and social media and online video sharing site (YouTube), just to name a few, and, of course, the biggest search engine and browser, Google and Google Chrome. We already touched on these two in the previous segments and how to find more secure and private alternatives, so we won't go into the details here.

We are going to start with alternatives to YouTube. One interesting option, acting as an alternative free and open-source front-end to YouTube, is Invidious. It is available in public form, from a list of public Invidious Instances hosted by individuals, or for self-hosting for interested individuals. Invidious doesn't actually use any Google API, but actually scrapes the website (extracting the data) for the requested video and metadata, such as views and likes. The main benefit of this solution is that it decreases the amount of data searched with Google. Another upside of Invidious is the lack of ads and the possibility of exporting users' YouTube subscriptions. There are also 3rd party projects using Invidious as a frontend, like Freetube, a desktop libre software YouTube app for privacy, and Clipious, an unofficial Invidious client for Android. The main upside of using Invidious, compared to other alternatives, is the fact that users, in theory, have access to all content that is normally available on YouTube. Unfortunately, there are downsides. Publicly facing Instances are in danger of getting shut down by Google itself. In 2023, Google sent a cease and desist letter to Invidious, claiming that they are violating YouTube's API policy and demanding that the service be shut down within seven days. As of this date, nothing has come of this, but Invidious still might be at risk. Another issue comes from the fact that using public instances still requires someone to have access to users' data, IP addresses for example.

Another option is PeerTube. It is a free and open-source, decentralized, ActivityPub(a protocol and open standard for decentralized social networking) federated video platform, which means that it is part of a federated network known as the Fediverse. This design, similar to Invidious, allows individuals to host social media platforms using their own hardware while still allowing interaction with other instances, but there are still public instances available through the internet. As the name suggests, PeerTube uses peer-to-peer communication and streaming, and because of that, there's no database. All the videos are stored and streamed directly from the instances. That means that not all PeerTube instances have the same content, but there's still a possibility of mirroring individual videos or whole friend instances, creating an incentive to build communities of shared bandwidth. In theory, all of this would allow for the creation of giant networks of servers hosting PeerTube instances (also named federations), with all of them sharing their own videos with each other without the need for individual instances to hold all the videos themselves. The main issue

with PeerTube stems from the popularity of the service. It is based on user participation and how many videos they are willing to share. Many federations focus on specific areas such as sports, games, or music, but none can match YouTube in terms of the variety and sheer volume of videos uploaded. Another issue, similar to Invidious, is the fact that members of the network can see your public IP. Overall, this option is not ideal; there are just not enough people participating in most federations, so there are just not enough videos, and most of them are created by a few individuals.

The last two options we'll cover were created by YouTube creators. First is Floatplane, created on April 3, 2017, by Linus Sebastian and Yvonne Ho as part of their digital media entertainment company, Linus Media Group Inc. It is an online streaming service that offers creators a platform to upload and monetize their content behind a paywall. The creator Linus Sebastian is known for advocating for privacy and even created a video series on his YouTube channel about degoogling, so there's a lot of trust in his platform. The second one is Nebula.tv, founded by Dave Wiskus on May 23, 2019. It is a video-on-demand streaming service provider, and similarly to Floatplane, it requires a subscription to access the video content. Both of these platforms contain videos from specific content creators, in the case of Floatplane, mostly Linus Tech Tips, and contain both videos available on YouTube and special ones only available on the specified service. In terms of privacy, both platforms collect users' data, and both in similar ways:

> Contact details, such as your first and last name, email and mailing addresses, and phone number.
> Account information, such as your username (email address) and password, watch history, preferences, and other details about your use of the services.
> Payment and transactional data, such as the information needed to complete merch orders and subscriptions on or through our Platform, and records of merch and subscriptions purchased.
> and more.

Both Nebula.tv and Floatplane claim that they don't sell collected data, and to this day, there's no specific reason to distrust that claim. On the other hand, they still need to follow laws in the continents and specific countries in which they operate. That included GDPR mentioned in the cryptocurrency portion. A big difference between these two options is the hosted content. Because of the sheer amount of content creators on Nebula.tv and their variety, it is, for most people, a significantly better YouTube alternative than Floatplane, mostly specialized in tech and gaming content. In the end, both options primarily serve to reduce the amount of customer data Google collects and shares with third parties; they don't stop the data collection altogether.

Next, let's mention an alternative to Google Play Services, micro, which is a free and open-source implementation of proprietary Google libraries for the Android operating system. MicroG allows Android apps to access replica application programming interfaces (APIs) that are provided by Google Play Services, including the APIs associated with Google Play, Google Maps, and Google's geolocation and messaging features. The main benefit of this service is that, unlike Google Play Services, it does not track user activities. It also allows for customization of which feature of the API is enabled or disabled. The project is still being developed, with a dedicated community on many social media platforms. There are issues, however. There might be some issues with certain apps, and recently, some Huawei users reported issues with microG after the latest OS update, but because we plan on using custom

ROMs for mobile devices in our final system, this should not affect us. This is even more important because for microG to work fully and work for most apps, signature spoofing, which is a mechanism that allows microG to pretend to be Google Play Services by faking its digital signature. It is crucial for apps that perform strict signature checks, because otherwise, they will not accept microG as a replacement. Fortunately for us, many ROMs support signature spoofing; similarly to the previous issue, we will not be worrying about that.

Next, let's talk about alternatives to Gmail. It is almost impossible in most cases to truly remove yourself from it. Many institutions use it for their professional emails, but for this analysis, we would like to focus on personal usage. The first option is Tuta, formerly known as Tutanota. It is an open-source, end-to-end encrypted email app and a freemium, with both free and premium options, and a secure email service. It allows for one address with 1GB of space in the free plan and up to 500GB of space, 30 addresses, and 10 personal domains for premium users. Other pros include encrypted messages, potentially even when sending to non-users, encrypted subject lines, end-to-end encryption, encrypted data at rest, stripping IP address from emails, two-factor authentication (2FA) support, doesn't sell users' data, usage of proprietary client-server protocol to increase security and privacy. The whole Tuta system works by encrypting almost every part of communication and only storing hashed values sent by the user.

First, for every new account combination of public and private keys is generated. The email and public key are stored in plaintext. The password, however, is combined with a generated salt and then hashed using the Argon2 algorithm. This results in an "AES password key" used for encrypting and decrypting the private key. After that, it is hashed once, using the sha256 algorithm, to create a "password verifier" that then is sent to the server, and only then is it stored after another hash, sha256, has been applied, creating a "hashed pw verifier". This makes it impossible to use data in rest for logging usage. The private key is, as mentioned before, encrypted using the "AES password key" and it is stored locally and in the Tuta database, but because it is encrypted, only users' clients can have access to it and use it to secure their emails. All of this communication is also under TLS encryption, so that third parties don't have access to event-encrypted and hashed values in plaintext. It is also important to mention that Tuta is planning on implementing Post-Quantum encryption of part of their systems. They already added Kyber encryption during the password and key encryption phase and are planning on adding Post-Quantum encryption to data at rest. Unfortunately, there are downsides to using this service. They don't support classic email protocols like IMAP/POP3 and other email encryption methods like PGP and S/Mime, and because of that, use cases for Tuta are decreased. That doesn't mean that communications with certain people are impossible, but, for example, it is impossible to use Tuta with Outlook. And in the same vein, Tuta requires the use of specialized applications. There is also no way of importing existing accounts.

Another very popular option is Proton Mail. It is an open-source, Swiss end-to-end encrypted email service launched in 2014, owned by the non-profit Proton Foundation. Similar to Tuta, it operates with free and paid plans. Free plan gives customers 1GB of storage with one address and user, and paid ones up to 1TB of storage, 3 custom email domains, 2 users, 30 extra email addresses, and other options like unlimited hide-my-email aliases and access to their VPN, Proton VPN. Proton Mail allows for very similar benefits as

Tuta, which are end-to-end encryption, sending encrypted emails to non-users, two-factor authentication (2FA) support, and zero-knowledge storage, which works in a very similar way to Tuta's standard for storing user data. Both store the private key encrypted using the customer's password. The main difference between Proton and Tuta is what protocols they support and the extent of encryption they provide. Proton Mail supports PGP, which allows users to export their Proton Mail PGP-encoded public key to others outside of Proton Mail. Not only that, but there's also support for IMAP/SMTP and, in turn, many popular email apps, including Outlook, Thunderbird, and Apple Mail. All of this means that Proton Mail is significantly easier to use in many different situations, but also provides worse encryption. There's no encryption for subject line encryption and only partial for metadata, unlike Tuta.

Last, there's the possibility of hosting your own email server. This requires significantly more expertise and knowledge about how email systems operate and what they need to work. There are a few "out-of-the-box" solutions like Mail in the box, but it still will probably require a Virtual Private Server (VPS) and potential experience when issues appear.

The next important service that we would like to talk about is the file-hosting service Google Drive. It is well known that Google can "see" what you're uploading to your drive. To fight that, one popular option, developed by previously mentioned Proton Foundation, is Proton Drive. Similar to Google Drive, it is a cloud-based file storage. Allows for safe and encrypted storage, where only customers have access to their data. The whole system works pretty similarly to Photon Mail; that's because it uses the same idea as Photo Drive, aka Zero-knowledge storage. Unfortunately, the dedicated app is only available for Windows and macOS devices. The only way to interact with Photon Drive via Linux is using the "rclone" command-line program. This clashes with our pick for a private operating system, so Photon Drive will not be taken into consideration.

A different option is Nextcloud. It is a suite of client-server software for creating and using file hosting services. The main idea behind Nextcloud was self-hosting, where individuals would, using their own hardware or VPS, host a Nextcloud instance that only they or other people allowed would have access to, but now there are multiple public Nextcloud installs. They allow anybody to create an account for free or for a fee. It supports logging and monitoring, permissions, and multi-layered encryption (SSL\TLS for data in transfer and AES for data at rest). Unfortunately, most of these security benefits mostly apply to administrators, so they can't be configured when using public instances. Another issue with public installs is that Nextcloud doesn't use zero-knowledge storage, so admins can see all the files. Only sending already encrypted files would help, but would significantly reduce ease of use, and some of them don't even support full end-to-end encryption. Nextcloud is mostly an option for self-hosting, and to make it easier, but slightly less private, the use of a VPS is possible.

Next is Tresorit. It is a cloud storage platform that offers functions for administration, storage, synchronization, and transfer of data using end-to-end encryption. Allows for 3GB of storage on the free plan and up to 4TB of storage and a maximum of 10GB file on paid. Tresorit supports zero-knowledge, cryptographic key sharing, and client-side integrity protection, which guarantees that a file's content cannot be modified without the customer's knowledge, even if somebody hacks Tresorit's system. It supports 2FA and allows for Gmail and Outlook integration. It uses "multi-level encryption", which means that

encryption is applied on each level from tenant (or subscription) level(provides administrators additional capabilities allowing them to access the content of managed users, reset their passwords), to account level (ensures only owner of the data can access it), shared folder level (only intended recipients can access shared folders) and last file level (all encryption and decryption operations are completed client-side so the content of files never leaves an end-user's device in unencrypted form). Creation and storage of keys are very similar to Tuta. Tresorit supports Windows (10 or later), macOS (10.15 (Catalina) or later), Linux (Ubuntu 16.04 or later, and all distributions with a window manager from the last 4 years), iOS (15 or later), and Android (5.0 (Lollipop) or later). An important thing to mention in relation to privacy is the location of the organisation, Switzerland, because of that, Tresorit handles data under Swiss privacy laws that provide more substantial protection than similar laws in the US or even the EU. Tresorit also doesn't use non-convergent cryptography (algorithms that always produce the same output, given the same input), which makes it impossible to determine when your content matches others' content in the cloud. Downsizing mostly comes from the price of the service, which can be steep for private individuals, and some difficulty in usage. Some users also report issues with Linux support.

The term self-hosting has been thrown around a few times in this section, and we would like to dig a little bit deeper into what it means exactly and what security and privacy benefits might come from it. First, the term self-hosting refers to the practice of hosting and managing applications on your own server instead of consuming from public providers. It allows for full control over the configuration of services, full control over who has access to the services, and who sees all the data coming from and to it. With proper setup, it minimises and sometimes completely removes any personal data being sent to third parties, but it does require a significant level of expertise and knowledge. To set up whole working and private systems, many times requires setting up a personal proxy server, domain, DNS, and more. It also costs a considerable amount of money to even start, without having any old laptops or other types of hardware.

One of the rising services being self-hosted are AI models. AI is everywhere these days, and nowhere is it more apparent than in a flux of AI chatbots like ChatGPT, DeepSeek, or CoPilot. Unfortunately, most of them are heavily monitored and tracked. Recently, users have been complaining that their chats with ChatGPT are seen in Google search results. Deepseek's privacy policy discloses that it collects a significant amount of personal data from users, all of which is stored on servers located in China. This data includes user-generated content such as chat prompts, conversations, and file uploads; profile information like username, date of birth, email address, phone number, and password; as well as automatically collected device and network details, including IP addresses and device identifiers. There exist private public options like duck.ai, created by DuckDuckGo, that claim to be a private alternative to regular AI chatbots. Unfortunately, we can never be too sure, and because of controversies related to the DuckDuckGo browser, the only option that can be 100% trusted, assuming we know the black box model doesn't contain malware, is self-hosted AI models, and there are many options for that.

One of the most noteworthy is DeepSeek. It is important to mention that running models used in public chatbots requires resources not available to individuals, so we will focus on distilled models. There are many options like DeepSeek-R1-Distill-Qwen-7B and DeepSeek-R1-Distill-LLaMA-70B that differ by the number of parameters, in other words,

internal variables that the model learns during training to perform a specific task. More parameters generally mean more resource usage, but more accurate results, if the parameters are chosen correctly. It is an option for more tech-savvy people with unused hardware. There's even the possibility of changing the model to fit your preference, if needed.

The last idea we wanted to mention in this section is data removal. Even after obtaining digital independence from big tech companies, customer data collected beforehand might still be there. The best and easiest way is to pay for a data removal service. There are a lot of them on the market, and all of them are best on regulation regarding privacy. For example, in Europe, it is previously mentioned, GDPR and more specifically, "Right to erasure". A data removal service, in this context, acts on behalf of the customer and requests that services and organizations delete the customer's personal data. This process is significantly faster than manual requests due to the automation of tasks performed by the services. Some issues arise because most of these services are not free. They also require the customer to give the data removal service some of their PII, so it still requires some loss of privacy.

Most notable of these services are Deeper Dive and Incogni. Deeper Dive finds and removes personal data from hundreds of brokers, offers free detailed DIY data removal steps, links directly to found personal data, and provides detailed verification of removals. There is also optional AI processing for greater accuracy. Cons are that it doesn't distinguish between removed data from never-found data; in other words, there might be false positives. Incogni, on the other hand, finds personal profiles on people search sites, repeats removal requests as needed, registers you to suppress data gathering, and allows for custom removal requests. Unfortunately, unlike Deeper Dive, there is no verification of removed personal data and no details about found personal data. The issue with both of these options and services alike is the need for constant subscription. Without it, there's nothing stopping the same third parties from accessing our data again.

Browser Extensions

Another topic we would like to discuss is browser extensions, especially cookie blockers and ad blockers.

Cookies play an important role in enhancing user experience through personalized interactions on various sites. That is why some sites may not work as intended when we don't accept them. That being said, tracking cookies can actually hinder our privacy and confidentiality by collecting data about users' behaviors on the internet. Thankfully, thanks to GDPR (General Data Protection Regulation), websites are forced to give the user a choice to accept or reject cookies. Our main objective while writing this document is to maximize the confidentiality of the user; that's why we recommend that we reject all the cookies, and when the site forces us to accept some, we find another source on another website. There are some browser extensions that could make our lives easier, like "I still don't care about cookies". This extension hides the form, giving us an option to accept or reject cookies, but it comes with a downside. When there are cookies necessary for a website to work, this extension automatically accepts them. This extension certainly saves us a lot of clicks, providing bigger "ease of use", but clashes with our "maximum confidentiality" factor by

automatically accepting cookies that are necessary, without giving us an option to leave that site and look for another.

Another type of browser extension that we would like to cover is AdBlockers. Adblockers are a very popular type of extension that, as the name suggests, block “ads” (advertisements) and trackers. They aren’t perfect, and sometimes something slips by. More interesting are not the domains that the AdBlocker extensions block, but the additional domains that are present only when we use an AdBlocking service. In the paper [23] authors discovered that there are, in fact, domains like that. For example, mixpanel.com, which is a Google domain that tracks visits by AdBlock users, and stripe.com, which is an online payment platform through which users can support AdBlock by donating. There are different domains present based on what AdBlock we are using. Only with Ublock (An AdBlocking extension), the authors didn’t notice any significant increase in additional domains. This paper shows that AdBlock extensions, while blocking some trackers, introduce some of their own.

Another concern that we may look at, though not interfering with our confidentiality, may hinder the “ease of use” aspect of the system. We are talking about the toll that AdBlockers take on the performance of our system, though it isn’t significant, we want to mention it. In this aspect, as shown in the paper [23], “Ublock” also performed the best.

The last aspect of confidentiality concerning the browser extensions is the way we use these extensions. Since AdBlockers can be customized to a significant extent, by, for example, changing filter lists, this allows advertisers to fingerprint the user for cross-site tracking. It was shown in the study [32] that “users with advanced blocking are more susceptible to fingerprinting.”. The process of applying more aggressive filter lists may, in fact, paradoxically have detrimental effects on our privacy. According to this paper, customization of the filter lists with just 45 rules can reduce the anonymity sets of users that in fact care about their privacy to less than 48, which in their dataset consists of 0.2% of subjects. The question that presents itself is, why won’t we just disable the ability to customize filter lists? If we did that, it would have some unwanted consequences. For example, the potentially bad rules in the filter list would affect all the users, and we can’t forget that there are also filter lists tailored for specific needs, and using them would be rendered impossible. The lesson to take from this is that using more tools doesn’t necessarily mean more privacy.

As outlined by us earlier, in order to maximize our privacy, we recommend declining all the cookies and using Ublock as our AdBlocker extension, because it seems not to introduce trackers and it is the most effective performance-wise. We would also recommend not customizing the browser extensions too much, because it increases the possibility that we will be fingerprinted.

Disposable Emails and Phone Numbers

The next topic we are going to discuss is disposable emails and phone numbers. An email address is one of the most important PII (personally identifiable information). That is why many people would prefer to keep it private, and that is where the disposable email

technology comes into play. It became quite popular over the years because of its spam-preventing capabilities. This technology allows us to generate a temporary email that will disappear after a short while, which is usually around 30 minutes. People use it to create accounts on services without worrying that this service will send email messages to their main email address. It also makes sure that our main email address isn't in this service's database, which means that the owner of this service does not have our email address and that in case of data leaks, malicious actors won't get our real email address.

Of course, this doesn't mean it comes with no downsides. Let's start by mentioning that some disposable email services keep the emails for up to 30 days in spite of the claimed 25-minute expiration time. Disposable emails are also susceptible to email tracking, which could hinder our confidentiality by providing the sender with information about where and when the email was read. Another downside is the fact that the disposable email's inbox is public, which means, any user can see other users' temporary inboxes. If someone is using an example@ex.com email address, another person can also access this address's inbox at this moment. There are little to none disposable email services with implemented sandbox functionality for inbox isolation. They utilize browser cookies to distinguish each inbox. In order for other people to access our temporary email, they need to know it. There are two ways of creating a temporary email address. The first creates it on the basis of a keyword. The user provides the keyword "info" and the email generated would be info@xx.com. The other way creates a random email address slkexmtzh@xx.com. When we first look at it, the second option seems safer, but let's take into account the length of the randomly generated email. We could make it even safer if we used the 1 method and, as a keyword, gave a random hash. Now the email address is randomly generated and pretty long, which lowers the possibility that someone would find our inbox via a brute force attack.

It is advised not to create accounts using temporary email addresses, because those accounts can be easily hijacked through the password reset. Users should also avoid providing PII on the websites they access via a disposable email address, because they can then send us an email with our PII that anyone can read.

Another disposable media that we will mention is a disposable phone number. In the era where 2FA becomes more widely used, phone numbers are quite commonly requested by services to verify their users via a time-bound code sent in an sms. As a disposable media, it is very similar to disposable email addresses, so we won't talk about it.

Social Media Applications

Another topic in our project focuses on social networking applications. Sharing information about ourselves with the world through these apps isn't exactly helping our privacy. These apps profit from advertisements, and for those ads to generate profit, they need to be well-targeted. That is why these apps collect a lot of information about each user - to improve their advertising targeting. The more we use them, the more information about us those apps have. They know our age, occupation, romantic status, how many kids we have, and many, many more. The most basic recommendation about the usage of those sites that would increase our privacy is to not post any sensitive information on those pages for everyone to see. There have been multiple criminals caught because they posted something on social media. Some criminals who use social networking apps through the TOR Network use something called the "sock puppet tactic" to avoid compromising their identity. A sock

puppet is a fake identity or a person who uses multiple usernames to avoid detection. It's not a bad tactic if you are using those apps through the TOR Network, but as with everything, it's not foolproof, and other people can still deduce your identity, as shown in this paper [11].

The best way would be not to post anything or scroll. Then other people won't have any information about you, and the app won't be able to collect any information about you. As we see, the solution that would provide us with the most confidentiality would be not to use those apps at all.

7. Summary

With all the key aspects related to privacy established, we could now start building the systems. Before getting into specifics, we're going to explain our methodology. For each entry, we're going to discuss two or more possibilities, the first one will be focused on maximum privacy, and the other one will take a reasonable level of ease of use into consideration. It's also worth mentioning that both systems are required to have an internet connection.

The first thing to discuss is operating systems. As mentioned in the related chapter, the best option for maximum privacy is Tails OS with its amnesic nature and a lot of tools implemented for the sole purpose of user anonymity. This solution, however, is very inconvenient. With that in mind, for the second system, we are going to choose Mint Linux or Ubuntu Linux. Both systems are open source, so they are transparent by nature and do not force data collection. For those people who have to use the Windows system, we recommend removing all the telemetry.

Next, we're going to discuss privacy on mobile devices. For maximum confidentiality, we recommend custom ROMs because most of them are open source and are built with privacy in mind. Our custom ROM of choice will be Graphene OS, because of how it handles Google Play by sandboxing it, but LineageOS and crDroid are good options as well. For less technical people, standard OSs like IOS and Android with removed telemetry should be sufficiently private.

Another important tool for ensuring our privacy is a VPN. The best choice in our opinion is Mullvad VPN, which provides the user with the highest level of anonymity out of all the other options. However, if we would like to increase our confidentiality in multiple fields, we would recommend Proton, which also provides other private services like mail and drive.

We can not talk about internet privacy without mentioning internet browsers. The most confidential browser that we talked about is the TOR Browser. It provides enough privacy to safely access the dark web. It doesn't come without some downsides, because it is very slow and it doesn't come with a VPN built in, like our next option, the Mullvad Browser. It's still slow, but significantly faster than TOR Browser, while still providing almost full anonymity when accessing the surface web. The last option that provides the best balance

between ease of use and confidentiality is the Brave Browser. It doesn't affect web usage while blocking trackers and providing fingerprinting protection.

After browsers, let's mention search engines. We recommend that users test the previously mentioned options themself and choose the option that best suits them.

Digital payments are one of the places where we may want to preserve our privacy. The first thought that comes to mind is cryptocurrencies, but while researching them, we found that they are not viable solutions for privacy, because of various restrictions placed on them by different countries and states. Most of them require the information about the customer to be revealed in some form, which makes the private coins an unviable option. Which leaves us with cash, only available in non-digital payments, virtual cards, which only makes customers invisible to the merchant (uses a different card number), not the bank, and prepaid cards that are the best option for online purchases, but many vendors don't support them. We recommend using a combination of prepaids when it is available and virtual cards when needed.

Next, let's talk about independence from big tech companies like Google. Their services collect a lot of data about users to enhance their ad targeting. To minimize it, we found alternatives to the most commonly used services. First, let's tackle the Google Mobile Services, like Google Play. The main one we want to propose is microG. It uses Google libraries to mimic traditional Google Services, allowing for the usage of applications that rely on them, without collecting our data. Another service that we found a replacement for is YouTube, and it is Invidious. It acts as an alternative Frontend for YouTube, where, instead of using the YouTube API to show content, it dynamically scrapes it for you. This gives Google significantly less data about users. We also recommend all apps that use Invidious as part of their system. A paid option, that might act as an alternative, is Nebula.tv, which hosts both regular content and premium one of the few YouTube content creators. The different service we wanted to replace is Gmail. The most secure and private alternative we found is Tuta. It encrypts every part of the service possible, from user emails to the end-to-end transferring process. Another one, previously mentioned, is Proton. This time, Proton Mail. Similarly to Tuta, it allows for encryption of customers' emails and the transfer itself, but it doesn't allow for full encryption. However, it allows for the usage of different email service apps like Thunderbird and Outlook, with the Proton Bridge functionality. The next service that we found an alternative for is Google Drive. The most private option is Tresorit, because of its zero-knowledge storage, multilayered encryption, and compliance with Switzerland's privacy policies. Downsides come in the form of its price, which can be a bit steep for individuals, and the difficulty of use. There is a cheaper option in the form of Proton that we mentioned a couple of times before. It is also characterized by its zero-knowledge storage, but the encryption is on a lower level. AI chatbots became more common over the years, and that is why we searched for more confidential alternatives to, for example, ChatGPT. The most confidential is a self-hosted AI Chatbot, which unfortunately suffers from severe downsides in ease of use. It requires significant resources of computational power, technical knowledge, and is still going to be slower and less complex than the mainstream alternatives. There are public alternatives, focused on privacy, like duck.ai and Luma AI, but by the nature of self-hosting, they are still less private and require users to trust the provider. The last topic we covered in the chapter about digital independence is data erasure. Our two main

recommendations are Deeper Dive and Incogni, because of their proficiency and quality of erasing data.

Browser extensions are a crucial element for tracker blocking. Different Adblocker-type extensions block trackers, but the only one that doesn't put some trackers of its own is Ublock, which is our recommendation. It is also the best performance-wise. The extension that could possibly increase our ease of use is "I still don't care about cookies" which hides the cookie consent widgets, which saves us a lot of clicks, but unfortunately, it also automatically accepts the cookies that are necessary for the website to function properly.

There is one technology that seems to increase our confidentiality, but in reality, it does not. We are talking about temporary emails and phone numbers. While being quite efficient with spam reduction, it creates another vector for our data to leak into unauthorized entities. Because of the public inboxes, it can not be considered a confidential alternative to regular email.

The last topic we will discuss is social media usage. In order to maximize our anonymity and minimize our vulnerability to OSINT, we recommend not using them at all, but if needed, we should at least apply social media hygiene rules. Those consist of things like not posting any sensitive information there.

The maximally confidential system is composed of the following components:

- OS - Tails OS
- Mobile OS - Graphene OS
- VPN - Mullvad
- Browser - TOR Browser
- Payments - Prepaid Cards / Virtual Cards
- Service Alternatives
    - Google services - microG
    - YouTube - Invidious
    - Gmail - Tuta
    - Google Drive - Tresorit
    - AI Agent - Self-Hosted
- Browser Extensions - Ublock
- Social Media - Do not use

The maximally confidential system, accounting for a reasonable level of ease of use, is composed of the following components:

- OS - Ubuntu/Mint
- Mobile OS - IOS/Android with disabled telemetry
- VPN - Mullvad/Proton
- Browser - Mullvad/Brave
- Payments - Prepaid Cards / Virtual Cards
- Service Alternatives
    - Google services - microG
    - YouTube - Invidious/Nebula
    - Gmail - Proton Mail
    - Google Drive - Tresorit/Proton Drive + rclone

    - AI Agent - Duck AI/Luma AI
- Browser Extensions - Ublock/I still don’t care about cookies
- Social Media - Apply proper usage hygiene

The differences between the search engines were so subtle that you should choose the one you personally prefer.

In this document, we provided some options for maximal privacy and other options for maximal privacy, paired with a reasonable level of ease of use.

References

[1] https://www.privacyguides.org/en/basics/why-privacy-matters/
[2] https://learn.microsoft.com/en-us/windows/privacy/
[3]https://threecats.com.au/comparison-of-custom-alternative-android-os-roms-grapheneos-divestos-calyxos-iodos-eos-lineageos-stock-android-aosp
[4] https://cybelangel.com/harvest-now-decrypt-later-hndl-attacks/
[5] https://www.privateinternetaccess.com/blog/security-audit-2024/
[6] https://mullvad.net/en/browser/mullvad-browser
[7]
https://www.researchgate.net/publication/380558771_Cybercrimes_in_the_Darknet_and_Their_Detections_A_Comprehensive_Analysis_and_Future_Directions
[8] https://ieeexplore.ieee.org/abstract/document/9197590
[9]
https://www.forbes.com/sites/kellyphillipserb/2019/10/16/irs-followed-bitcoin-transactions-resulting-in-takedown-of-the-largest-child-exploitation-site-on-the-web/#327343231ed0

[10] https://www.sciencedirect.com/science/article/pii/S1877050915008406
[11] https://www.sciencedirect.com/science/article/abs/pii/S0950705112002365
[12] https://www.mdpi.com/2078-2489/15/8/495
[13] https://www.youtube.com/watch?v=8ZX7kG2-aQY
[14]
https://www.theverge.com/2013/12/18/5224130/fbi-agents-tracked-harvard-bomb-threats-across-tor
[15] https://ieeexplore.ieee.org/document/668972
[16] https://medium.com/@balasubramanya.c/tor-browser-forensics-d11dfa13a063
[17]
https://www.researchgate.net/publication/315449078_Tails_Linux_Operating_System_Remaining_Anonymous_with_the_Assistance_of_an_Incognito_System_in_Times_of_High_Surveillance
[18] https://arxiv.org/abs/1906.11078
[19] https://kyc-chain.com/intersection-kyc-data-privacy/
[20] https://www.ulam.io/blog/is-cryptocurrency-anonymous#
[21]
https://www.ulam.io/blog/blockchain-analysis-in-action-real-life-use-cases-and-insights#

[22] https://www.alternativeairlines.com/cryptocurrency-vs-traditional-currency
[23] https://dl.acm.org/doi/10.1145/3653478
[24] https://eprint.iacr.org/2021/1054.pdf
[25] https://www.mdpi.com/1099-4300/25/9/1334
[26] https://chainstack.com/stealth-addresses-blockchain-transaction-privacy/
[27]
https://www.merklescience.com/blog/privacy-coins-legitimate-uses-and-illicit-risks-explained
[28] https://eprint.iacr.org/2023/321
[29] https://academy.bit2me.com/en/what-is-privatesend-dash/

[30] https://jewishpostandnews.ca/features/why-prepaid-cards-are-the-last-refuge-for-online-privacy-in-2025/
[31] https://arxiv.org/abs/1705.03193
[32] https://www.usenix.org/conference/usenixsecurity25/presentation/el-hajj-chehade
[33] https://mtlj.usc.ac.ir/article_215039.html?lang=en
[34] https://pmc.ncbi.nlm.nih.gov/articles/PMC10980991/
[35] https://dl.acm.org/doi/abs/10.1145/2065023.2065035
[36] https://arxiv.org/abs/1702.01807
[37] https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/
[38] https://docs.invidious.io/applications/
[39] https://joinpeertube.org/faq#what-are-the-main-advantages-of-peertube
[40] https://nebula.tv/privacy
[41] https://www.floatplane.com/legal/privacy
[42] https://tuta.com/pl/encryption
[43] https://cyberinsider.com/email/reviews/tuta-mail/
[44] https://proton.me/mail/security
[45] https://proton.me/support/proton-mail-encryption-explained
[46] https://workaround.org/ispmail-bookworm/creating-a-tls-encryption-key-and-certificate/

[47] https://proton.me/blog/business-cloud-storage
[48] https://nextcloud.com/secure/
[49] https://nextcloud.com/partners/
[50] https://cyberinsider.com/cloud-storage/reviews/tresorit/
[51] https://tresorit.com/security
[52] https://linuxblog.io/deepseek-local-self-host/
[53] https://growthscribe.com/why-duckduckgo-is-bad/
[54] https://arstechnica.com/tech-policy/2025/08/chatgpt-users-shocked-to-learn-their-chats-were-in-google-search-results/
[55] https://www.pcmag.com/picks/the-best-personal-data-removal-services#
[56] https://ieeexplore.ieee.org/document/1253566
[57] https://ieeexplore.ieee.org/document/8835374
[58] https://privacytests.org/
[59] https://arxiv.org/abs/2210.04541
[60] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3545331
[61] https://cryptodeeptech.ru/doc/Developing_Quantum-Resistant_Cryptography_Encryption_for_a_Post-Quantum_World.pdf
[62] https://www.bleepingcomputer.com/news/security/duckduckgo-browser-allows-microsoft-trackers-due-to-search-agreement/
[63] https://www.researchgate.net/publication/368727573_Post_Quantum_Cryptography_A_Review_of_Techniques_Challenges_and_Standardizations

[64] https://epubs.siam.org/doi/abs/10.1137/S0097539703440678
[65] https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.203.pdf