



SCUOLA di
DOTTORATO

Dottorato in Ingegneria dell'Informazione

Direttore della Scuola di Dottorato
Prof. Felice Arena

Collegio dei docenti
Dottorato di Ricerca in Ingegneria dell'Informazione
XXXVIII ciclo

Tommaso ISERNIA (Coordinatore)

ARANITI Giuseppe
BUCCAFURRI Francesco Antonio
CAMPOLO Claudia
CAROTENUTO Riccardo
COPPOLA Giuseppe
COTRONEI Mariantonia
CROCCO Lorenzo
DALLEY Dominique
DE CAPUA Claudio
DELLA CORTE Francesco Giuseppe
DI DONATO Loreto
FAGGIO Giuliana
FAILLA Gioia
GIUFFRÈ Sofia
GROZA Voicu
LAX Gianluca
LAY-EKUAKILLE Aimè

LOHAN Elena Simona
MANZONI Pietro
MESSINA Giacomo
MISTRETTA Marina
MOLINARO Antonella
MORABITO Andrea Francesco
MORABITO Francesco Carlo
MORELLO Rosario
MUNTEAN Gabriel-Miro
MUSOLINO Giuseppe
PEZZIMENTI Fortunato
PRATICÒ Filippo Gianmaria
RAO Sandro
ROMANO Maria
ROSACI Domenico
RUGGERI Giuseppe
RUSSO Mariateresa
VINEL Alexey



SCUOLA di
DOTTORATO

SCUOLA DI DOTTORATO
Università Mediterranea di Reggio Calabria

DIPARTIMENTO
Dipartimento di Ingegneria dell'Informazione,
delle Infrastrutture e dell'Energia Sostenibile (DIIES)

DOTTORATO DI RICERCA
Ingegneria dell'Informazione

S.S.D. IINF-05/A
XXXVIII CICLO

Software and Hardware Architectures for Enhanced User Privacy

Dottorando
Aurelio Loris Canino

Tutor
Prof. Gianluca Lax

Coordinatore del Dottorato
Prof. Tommaso Isernia



Finanziato
dall'Unione europea
NextGenerationEU



To my cousin

Acknowledgments

This work was supported by a PhD scholarship funded by the Italian National Recovery and Resilience Plan (PNRR), Mission 4 'Education and Research', Component 1, Investment 3.4 (Action A), pursuant to Ministerial Decree No. 351/2022, funded by the European Union – NextGenerationEU.

Abstract

The digital ecosystem has evolved into a pervasive, interconnected environment in which personal data is continuously collected, processed, and shared across heterogeneous platforms. This evolution, while enabling unprecedented levels of personalization and efficiency, has also intensified users' privacy concerns, exposing individuals to risks of surveillance, profiling, and unauthorized access. This thesis investigates these concerns through a series of applied studies and technical contributions aimed at strengthening privacy in six contexts of information systems: (i) privacy-preserving attribute certification, introducing a blockchain-based model to enforce GDPR's data minimization principle; (ii) decentralized spyware detection through distributed signature sharing; (iii) privacy-preserving lawful interception using Self-Sovereign Identity (SSI) for verifiable access control; (iv) router-based parental control systems ensuring local and auditable enforcement of policies; (v) architectural weaknesses in IoT safeguard environments enabling covert interceptions and privacy violations; and (vi) privacy audits of Generative AI browser assistants, revealing new dimensions of user tracking and profiling.

The proposed models and implementations aim to reconcile functionality, transparency, and accountability, contributing to a privacy-by-design paradigm that addresses both human trust and technological robustness. Across these studies, this thesis demonstrates that privacy cannot be treated only as a regulatory requirement or an add-on to security, but as a foundational design principle that must be systematically embedded into the architecture of digital systems.

Sommario

L'ecosistema digitale si è evoluto in un ambiente pervasivo e interconnesso in cui i dati personali vengono continuamente raccolti, elaborati e condivisi attraverso piattaforme eterogenee. Questa evoluzione, pur consentendo livelli senza precedenti di personalizzazione ed efficienza, ha anche intensificato le preoccupazioni degli utenti in materia di privacy, esponendo gli individui a rischi di sorveglianza, profilazione e accesso non autorizzato. Questa tesi indaga tali preoccupazioni attraverso una serie di studi applicati e contributi tecnici volti a rafforzare la privacy in diversi contesti dei sistemi informativi.

La ricerca si articola in sei aree principali: (i) certificazione degli attributi a tutela della privacy, introducendo un modello basato su blockchain per applicare il principio di minimizzazione dei dati previsto dal GDPR; (ii) rilevamento decentralizzato di spyware tramite condivisione distribuita delle firme; (iii) intercettazione legale a tutela della privacy mediante l'utilizzo della Self-Sovereign Identity (SSI) per il controllo verificabile degli accessi; (iv) sistemi di controllo parentale basati su router che garantiscono un'applicazione locale e verificabile delle policy; (v) debolezze architetturali negli ambienti IoT di salvaguardia che consentono intercettazioni occulte e violazioni della privacy; e (vi) audit sulla privacy degli assistenti browser basati su intelligenza artificiale generativa, che rivelano nuove dimensioni del tracciamento e della profilazione degli utenti.

I modelli e le implementazioni proposti mirano a conciliare funzionalità, trasparenza e responsabilità, contribuendo a un paradigma di privacy-by-design che risponde sia alla fiducia degli utenti sia alla robustezza tecnologica.

Attraverso questi studi, la tesi dimostra che la privacy non può essere trattata esclusivamente come un requisito normativo o un complemento alla sicurezza, bensì come un principio fondamentale di progettazione che deve essere sistematicamente integrato nell'architettura dei sistemi digitali.

Contents

1	Introduction	3
1.1	Outline	6
2	Background	9
2.1	Foundations of Information Security	9
2.1.1	The CIA Triad	9
2.1.2	Cryptography	13
2.1.3	Cryptographic Hash Functions	15
2.1.4	Digital Signature	16
2.1.5	Secure Socket Layer	19
2.2	Blockchain	20
2.2.1	Core Concepts of Blockchain Systems	20
2.2.2	Consensus Mechanisms: Proof of Work and Proof of Stake	22
2.2.3	Ethereum	23
2.3	Self-Sovereign Identity	26
2.3.1	Verifiable Credentials	26
2.3.2	Issuers, holders, and verifiers	26
2.3.3	Digital Wallet	27
2.3.4	Digital Agents	27
2.3.5	Decentralized Identifies	27
2.3.6	Verifiable Data Registry	28
2.3.7	Governance Framework	28
2.4	Lawful Interception	28
2.4.1	ETSI Handover Interfaces	30
2.5	Generative AI	31
2.5.1	GenAI Systems and Browser Assistants	32
2.5.2	Traditional vs. GenAI Browser Extensions	33

3	Privacy-preserving Attribute Certification	35
3.1	Introduction	35
3.2	Related Work	36
3.3	Certificate Model	38
3.4	Proposed Scheme	39
3.5	Discussion	43
3.6	Conclusion	43
4	Spyware detection	45
4.1	Introduction	45
4.2	Literature Review	48
4.2.1	Different types of blockchains	52
4.2.2	Comparative study	53
4.3	Proposed Work	55
4.4	Implementation	58
4.4.1	Cost and latency	61
4.4.2	Smart Contract Source Code	63
4.5	Security Analysis	66
4.5.1	Threat Model	67
4.5.2	Security Properties	67
4.5.3	Attacks	68
4.5.4	Mitigations	68
4.6	Discussion and Conclusion	70
5	Privacy-preserving Lawful Interception	75
5.1	Introduction	76
5.2	Motivations and goals	78
5.3	Proposed approach	80
5.3.1	Setup	81
5.3.2	External agency integration	81
5.3.3	Credential revocation	82
5.4	Design of the SSI stack	83
5.4.1	Layer 1	83
5.4.2	Layer 2	85
5.4.3	Layer 3	86
5.4.4	Layer 4	86
5.4.5	Stack Overview	86
5.5	Proof of Concept	87
5.5.1	Technologies	87

5.5.2	Case study	88
5.6	Security Analysis	92
5.6.1	Threat Model	92
5.6.2	Security Properties	94
5.6.3	Attacks and Mitigations	95
5.7	Related Work	97
5.8	Conclusion and Future work	98
6	A Router-based Parental Control Tool for Safe Social Network Usage	101
6.1	Introduction	101
6.2	Related Work	102
6.3	System Architecture	104
6.4	Implementation	107
6.4.1	Network Infrastructure	107
6.4.2	Modules	107
6.5	Conclusion	108
7	An attack on privacy in IoT Safeguard Environments	111
7.1	Introduction	111
7.2	Related Work	112
7.3	Problem Formulation	113
7.4	Design and Architecture of the System	114
7.5	Proof of Concept	118
7.6	Conclusion	124
8	Privacy issues in Generative AI Browser Assistants	125
8.1	Introduction	125
8.2	Related Work	127
8.3	Threat Model	129
8.4	Methodology	130
8.4.1	Selection of GenAI Browser Assistants	130
8.4.2	Crawling Infrastructure	132
8.4.3	Auditing User Tracking	133
8.4.4	Auditing Profiling and Personalization	134
8.4.5	Semi-automated Analysis	136
8.5	Results	137
8.5.1	Architecture of GenAI browser assistants	137
8.5.2	User tracking by GenAI browser assistants	140
8.5.3	Profiling and Personalization	147

XII Contents

8.6	Concluding Remarks	149
8.7	Prompting Framework	151
8.8	System Prompts	152
8.9	Conclusion	154
9	Conclusions	161
9.1	Summary of Contributions	161
9.2	Future Research Directions	162

List of Figures

2.1	Comparison between Proof of Work and Proof of Stake consensus mechanisms.	23
2.2	ETSI-defined LI architecture.	29
3.1	Example of the model of a certificate.	39
3.2	Sketch of the code of the smart contract.	40
3.3	Example of certification.	42
4.1	Architecture of the solution.	57
4.2	Certification request.	59
4.3	Software for signature generation.	60
4.4	Cost of each operation.	62
4.5	Latency of each operation.	62
4.6	Sketch of the smart contract (part 1).	64
4.7	Sketch of the smart contract (part 2).	65
4.8	Sketch of the smart contract (part 3).	66
5.1	Sequence diagram of the Setup in the LI procedure.	81
5.2	Sequence diagram of the EA integration in the LI procedure.	82
5.3	Comparison between the general SSI stack (a) and its instantiated version (b).	84
5.4	Registration of LEA, LEMF, and NA Agency.	88
5.5	Fields of NA Agency registration.	89
5.6	Overview of the Employee Credential (a) and X Credential Verification (b).	90
5.7	Overview of QR Code (a) and Verifiable Presentation (b).	91
5.8	Overview of the X Credential issuance process (a) and its storage in Bob's wallet (b).	93

XIV List of Figures

5.9 Overview of the policy for the Y Credential (a) and the Y Credential itself (b).	94
5.10 Indy Scan.	95
6.1 Environment considered in the study.	104
6.2 Example of the architecture considered in the study.	105
7.1 Example of the environment considered in the study.	115
7.2 Architecture of the proof of concept.	119
7.3 Components used in the proof of concept.	120
7.4 Packet flow during covert channel: red packets are injected by Raspberry and removed by Provider.	122
7.5 Traffic flow at the Provider.	124
8.1 Threat model of GenAI browser assistants.	130
8.2 Experimental design of our auditing framework for AI browser assistants.	133
8.3 Sankey diagram depicting third-party flows of different assistants.	145

List of publications

1. G. Lax, A. Canino, and L. Musarella, "A blockchain-based approach for certifying information at different levels of granularity according to the data minimization principle of the gdpr," in *2024 8th International Conference on Computer, Software and Modeling (ICCSM)*. IEEE, 2024, pp. 31–35.
2. A. L. Canino and G. Lax, "A cost-effective solution leveraging public blockchain for massively sharing malware signatures," *Journal of Information Security and Applications*, 2025.
3. F. Buccafurri, A. L. Canino, V. De Angelis, A. Laurenda, and G. Lax, "An ssi-based solution to support lawful interception," *Applied Sciences*, vol. 15, no. 4, p. 2206, 2025.
4. A. L. Canino, V. De Angelis, and G. Lax, "A router-based parental control tool for safe social network usage," in *International Conference on Advances in Social Networks Analysis and Mining*. Springer, 2024, pp. 91–100.
5. A. L. Canino and G. Lax, "Enabling lawful interception in environments protected by iot safe-guard," in *International Conference on Electronic Government and the Information Systems Perspective*. Springer, 2024, pp. 139–153.
6. Y. Vekaria, A. L. Canino, J. Levitsky, A. Ciechonski, P. Callejo, A. M. Mandalari, and Z. Shafiq, "Big help or big brother? auditing tracking, profiling, and personalization in generative ai assistants," in *USENIX Security '25*, 2025.

Introduction

Privacy is a fundamental principle in modern digital society. It encompasses not only the protection of personal data but also the design and management of digital systems. While data protection traditionally focuses on preventing data breaches and adopting security best practices to avoid unauthorized access, privacy extends further: it includes the right to confidentiality, freedom from tracking, and the ability to disclose only the minimum amount of personal information necessary for legitimate purposes.

A first domain in which these principles are particularly significant is that of digital certificates. Systems that issue certificates without embedding privacy-by-design mechanisms risk disclosing more information than strictly required, thus violating the GDPR's principle of data minimization. The challenge lies not only in securing the data, but also in assessing the necessity and relevance of the attributes being disclosed. Platforms for managing digital certificates must ensure that only strictly essential personal data are disclosed, and solely for their intended purpose. To address this, a specific line of research proposes an architecture that enables the selective disclosure of certified attributes. Leveraging decentralized technologies such as blockchain empowers users to reveal only essential information, without requiring the reissue of new certificates. This research is published in [1].

Moving from digital identity to cybersecurity, another domain in which privacy is structurally challenged is spyware detection. Spyware constitutes one of the most intrusive threats to user privacy, as it is designed to collect personal information and send it to third parties without the person's consent. Depending on its capabilities, it may only track browsing history, or it can go as far as intercepting login credentials, financial data, and sensitive documents, enabling hidden surveillance, identity theft, and invasive profiling. In this context, delays in distributing signatures for new variants of malware not only increase the risk of systems being compromised, but also expose users to systematic violations of their private lives. Traditional anti-malware solutions rely on proprietary, centralized, and vendor-managed databases

that inhibit rapid signature updates, aggravating privacy risks from emerging spyware threats. To counter this, this research line introduces a decentralized system leveraging a public blockchain, such as Ethereum, for transparent and verifiable sharing of malware signatures among certified agencies and open verification by any user. This permissionless approach ensures the dissemination of threat intelligence without registration barriers, while smart contracts enforce agency certification via oracles and cryptographic checks to prevent tampering or Sybil attacks. By distributing signatures immutably and cost-effectively, the framework enhances real-time detection, preserves user privacy through anonymous verification, and demonstrates superior pervasiveness over private blockchain alternatives. The outcomes of this work are published in [2].

A particularly delicate balance between investigative needs and privacy protection arises in the domain of lawful interception (LI), where highly invasive investigative needs are closely linked to the necessity to safeguard users' privacy and the probative integrity of intercepted communications. In this context, existing standards precisely define the interfaces between the Network Operator (NO) and the Law Enforcement Monitoring Facility (LEMF), but do not regulate access by judicial authorities or delegated agencies to data held within LEMF systems. This leads to varied procedures, often manual, difficult to automate, and based on authentication mechanisms that are hard to verify. This architectural gap opens the door to structural risks such as interoperability problems, operational inefficiencies, insider threats, and uncontrolled data storage, which become even more severe in cross-border scenarios where several parties must cooperate on the same intercepted information flows. To bridge this gap, this research proposes an access control system based on Self-Sovereign Identity (SSI), where authorizations, delegations, and revocations are managed as verifiable credentials exchanged on decentralized infrastructures that follow the Sovrin/ToIP stack. This approach standardizes the interaction between Law Enforcement Agencies (LEA), LEMF, and external agencies through secure peer-to-peer channels, makes authorizations easy to automate with machine-readable formats, and allows precise control over revocation. As a result, it reduces the exposure of data to the minimum needed and fills the regulatory and technical gap between LEA and LEMF. This case study demonstrates that privacy can be treated as an architectural principle even in domains traditionally hostile to privacy, showing how decentralized paradigms can combine accountability, global interoperability, and investigative capabilities without relying on centralized identity providers or ad hoc procedures that are difficult to audit. This contribution is presented in [3].

Privacy challenges also emerge in domestic contexts, particularly regarding parental control over minors' online activities. In this thesis, the issue of parental control over

underage users' online activities is examined in light of the 'privacy by design' principle. Most current parental control systems rely on third-party apps and infrastructure, which are installed on individual devices and often use centralized cloud services. This model exposes minors' sensitive data and online behaviour to real risks of profiling, personalization, and constant monitoring by entities outside the family. These architectures also lack local and transparent control over the information collected and processed by the system. To address these issues, an alternative approach is proposed, focusing on integrating parental control mechanisms directly into the home router. The central role of the router in the home network makes it possible to monitor, inspect, and manage Internet traffic from any connected device, applying control policies in a consistent and verifiable way. This strategy significantly reduces exposure of data to third parties, allows privacy-friendly traffic inspection, since it does not require access to the content of communications, and ensures that rules set by parents are enforced locally, can be audited, and are transparent. As a result, this model provides effective online protection for minors without sacrificing their privacy rights, offering a practical example of how security, control, and privacy can be combined in modern digital architectures for families. This study is published in [4].

However, even architectures intended to enhance privacy can themselves become attack vectors. Another research line examines how defense mechanisms designed to protect privacy in domestic environments can be circumvented and turn into structural weak points themselves. The focus is on IoT safeguards integrated within home routers, which are intended to monitor network traffic and block anomalous connections. The proof of concept shows how an attacker can exploit compromised IoT devices to carry out a stealth lawful interception in environments protected by IoT Safeguard. By defining an architecture consisting of interceptors, concentrators, hidden exfiltration channels, and Blockchain-based mechanisms for preserving evidence, the case study demonstrates that current defence solutions are unable to prevent targeted interceptions by covert channels within regular traffic. This contribution highlights that privacy remains vulnerable to advanced forms of surveillance unless it is treated as a core design principle from the outset in the development of routers, IoT devices, and monitoring systems. These findings are detailed in the publication [5].

Finally, the rapid rise of browser assistants powered by generative AI is bringing new privacy concerns to everyday web browsing. Recent advances in generative AI have led to the fast development of browser assistants using large language models, such as OpenAI's ChatGPT and Google's Gemini. Unlike traditional browser extensions, these assistants can leverage natural-language understanding to offer features

such as conversational interaction, webpage summarisation, automatic form filling, coding support, and even explanations of online videos. While such capabilities can greatly enhance the browsing experience, they also give rise to important privacy concerns. Because generative AI browser assistants operate within the browser itself, they may collect personal information as users navigate both public and private spaces of the internet. Their ability to access page content, monitor user activity, and leverage LLMs opens the door to detailed profiling and personalized responses. Moreover, the lack of clear information about whether, how, and with whom user data is shared, including potential third-party advertisers and analytics services, creates further concerns. This thesis addresses these concerns by investigating whether current GenAI browser assistants collect and share user information, what types of personal data they gather, and whether they build behavioural or demographic profiles to personalize user interactions. To answer these questions, we introduce a novel prompting framework and perform a network traffic analysis of nine popular browser assistants, each downloaded by millions of users from official extension stores. Our audit examines their design, behaviour, data flows, and profiling mechanisms. The findings highlight the critical need for privacy-centric design principles, particularly privacy by default, in the development of GenAI browser assistants. They also underscore the importance of clearer, more explicit privacy policies that disclose what data is collected, what is transmitted to first-party servers, and which third parties receive user information. Ultimately, this study shows that as GenAI becomes increasingly embedded in everyday browsing, robust privacy safeguards are not optional, but essential. This extensive audit is published in [6].

1.1 Outline

This thesis is organized as follows:

- **Chapter 2 — Background.**

This chapter is devoted to discussing foundational topics that provide the background knowledge assumed in the subsequent chapters. The discussion is organized into five main thematic areas: Foundations of Information Security; Blockchain technology and consensus mechanisms; Self-Sovereign Identity (SSI) and Verifiable Credentials; the architecture of Lawful Interception standards; and the technical foundations of Generative AI, focusing on LLM-powered browser assistants.

- **Chapter 3 — Privacy-preserving Attribute Certification.**

In this chapter, we highlight how issuing certificates that are too detailed may lead to a violation of the principle of data minimization stated by GDPR. The problem lies in the design of the certification, which does not allow for the limitation of the personal attributes disclosed.

- **Chapter 4 — Spyware detection.**

A specific class of malware, known as *spyware*, is designed to steal personal or confidential information from infected systems, posing a direct threat to users' privacy. Traditional anti-malware solutions rely on centralized and proprietary databases of malware signatures, each managed by individual vendors. This model can slow down the distribution of newly discovered signatures, delaying the update of antivirus software and exposing users to emerging threats. To address these issues, we propose a decentralized approach based on a public blockchain, enabling transparent and verifiable sharing of malware signatures among users. This ensures rapid dissemination of threat intelligence while preserving user privacy.

- **Chapter 5 — Privacy-preserving Lawful Interception.**

In this chapter, we analyze how the absence of robust standards that define access by Law Enforcement Agencies (LEAs) or delegated authorities to intercepted data on Law Enforcement Monitoring Facility (LEMF) systems constitutes a structural software issue that can be exploited by malicious actors, unfaithful employees, or unauthorized users to compromise the privacy and integrity of investigations, as well as resulting in manual and insecure procedures for exchanging sensitive data, often based on unverifiable authentication mechanisms. The chapter proposes a Self-Sovereign Identity (SSI)-based solution that introduces a decentralized access-control mechanism designed to enforce strict access controls and close the gap between LEAs and LEMFs.

- **Chapter 6 — Router-based Parental Control for Safe Social Network Usage.**

In this chapter, we highlight and critically assess current parental control systems, which are single-device-based and managed by third parties; these architectures pose a serious software-level privacy risk that can result in the profiling, personalization, and continuous monitoring of the minor. The proposed approach, centered on the home router, mitigates this risk by centralizing control locally, minimizing data exposure, and ensuring privacy-preserving traffic inspection alongside verifiable, auditable enforcement of control policies across all devices subject to the parental control system.

- **Chapter 7 — An attack on privacy in IoT Safeguard Environments.**

In this chapter, we highlight how an architectural weakness in IoT safeguard mechanisms can be exploited to bypass their control, enabling an attacker to

conduct covert interceptions through compromised devices. This represents a privacy attack involving both hardware and software layers, exposing the structural limitations of domestic defense systems.

- **Chapter 8 — Privacy issues in Generative AI Browser Assistants.**

In this chapter, we analyze how GenAI browser assistants collect and transmit to remote servers the entire content of visited web pages, form inputs, and meta-data. The problem lies in the centralized and opaque design of the browser extension software, which implements profiling and tracking functions without local user control.

- **Chapter 9 — Conclusions.**

This chapter summarizes the thesis's contributions, highlighting how the proposed architectures advance the state of the art in privacy-preserving systems. Finally, it asserts the necessity of treating privacy as a non-negotiable architectural requirement for a more resilient digital infrastructure.

Background

This chapter is devoted to discussing foundational topics that provide the background knowledge assumed in the subsequent chapters. The discussion is organized into five main thematic areas:

- Foundations of Information Security;
- Blockchain;
- Self-Sovereign Identity;
- Lawful Interception;
- Generative AI.

2.1 Foundations of Information Security

Information security provides the conceptual and theoretical foundations for understanding how information systems can be designed and analyzed to ensure protection, correctness, and trustworthiness, as well as how security properties are defined, achieved, and evaluated in modern systems.

2.1.1 The CIA Triad

The CIA triad is a fundamental model in cybersecurity that defines and guides the main objectives of information protection.

The three core attributes of security help us distinguish among the different security goals to be achieved, as well as the various types of attacks that can compromise them:

- Confidentiality;
- Integrity;
- Availability.

These attributes form what is known as the CIA triad. They must be ensured both for data and services.

These attributes are not independent from one another. For instance, an attack on the integrity of a database may also compromise the availability of the service that relies on it, since corrupted data could prevent the service from functioning correctly.

Integrity

Integrity concerns the prevention of unauthorized modifications to data, services, and the underlying hardware and software infrastructure. Authorized modifications to a database are not considered anomalies; the anomaly arises when modifications are made without authorization, which is referred to as data tampering. Service compromises can also occur.

Access control mechanisms are designed to ensure the integrity attribute. A classic example is that of managers within an organization who have permission to access certain information. Therefore, the integrity of information systems refers to the necessity of maintaining their correct operation: functions must perform exactly as they were designed to.

Confidentiality

Confidentiality refers to the requirement that information must remain private and accessible only to authorized entities. The set of authorized subjects is defined in advance, and through specific mechanisms, only these subjects and no others are allowed to access the protected information.

Confidentiality does not simply rely on denying access to information. Even if an attacker gains access, confidentiality ensures that the attacker cannot interpret or understand the information content. Access control mechanisms also contribute to confidentiality by reducing the attack surface.

Confidentiality applies to:

- **Data:** ensuring that certain data remain confidential and defining the set of subjects for whom those data are not confidential;
- **Physical and legal entities:** physical entities are individuals, while legal entities include companies, public administrations, and states;
- **Organizations:** organizations must protect data such as trade secrets, sensitive information, and judicial data.

Organizations have a strong interest in protecting data, as such protection is linked to the risk of losing a competitive advantage. For example, industrial product data must be transmitted while ensuring confidentiality through appropriate

mechanisms. In general, organizations classify information according to levels of confidentiality (e.g., low, medium, high).

To ensure confidentiality, two different (and possibly combined) approaches can be used:

1. Making the information unintelligible to unauthorized subjects;
2. Hiding the information itself.

For the first approach, cryptography can be used. Cryptography is not limited to confidentiality; it can also serve to guarantee integrity. However, cryptography and confidentiality are not synonymous: cryptography is a security mechanism used to achieve the security property of confidentiality. In this chapter, we will discuss mechanisms such as symmetric and asymmetric cryptography.

The second approach is known as information hiding. Information hiding refers to situations where information exists within a communication session between two entities in a network, but an unauthorized third party cannot perceive it because it is well hidden. This technique requires a cover object within which the secret content is embedded.

The general term for implementing information hiding is the creation of a covert channel. The main techniques for achieving information hiding are steganographic techniques, although others exist as well. A hidden message might be embedded within a plaintext communication between two users. The recipient can reconstruct the secret message at the destination because both sender and receiver share a secret that allows them to recompose the original content.

Regarding steganographic techniques, there are two main approaches to creating a covert channel:

- Injection: bits are injected into another binary content without compromising the cover object. Typical file formats supporting injection include PDF, HTML, and TXT;
- Substitution: instead of adding bits, bits are replaced. This method is generally applied to multimedia files (audio, video, bitmap), where bit substitution often does not affect usability. The replaced bits usually correspond to noise or low-significance content.

Steganography can also be used for legitimate purposes, particularly in the field of watermarking.

Watermarks are steganographic contents (stegcontent) embedded within a cover object (digital content). They serve two main purposes:

1. Ensuring the integrity of digital content to prevent it from being tampered with (from the verb *to tamper*, meaning to alter or compromise);

2. Attributing authorship to digital content is an aspect of Digital Rights Management (DRM), which involves the protection of intellectual property rights, for example, enabling a photographer to prevent unauthorized use of their images.

Availability

The third element of the CIA triad is availability, which represents the degree to which a system or service remains operational and accessible when required for use. Formally, availability can be defined as the probability that a system operates correctly at a given point in time. It is often expressed mathematically as:

$$A = \frac{T_{total} - T_{downtime}}{T_{total}}$$

where T_{total} is the total observation time and $T_{downtime}$ is the cumulative time during which the system is not operational.

Availability ensures that authorized users can access the information and resources they need, whenever they need them. In modern cybersecurity frameworks, availability is also considered a key attribute of dependability, a broader concept encompassing the system's overall ability to deliver services reliably and safely.

Dependability

Dependability refers to the ability of a system to provide services that can be justifiably trusted. It includes several key attributes, typically identified as:

- Reliability: the continuity of correct service, usually quantified by the *Mean Time Between Failures (MTBF)*:

$$MTBF = \frac{T_{total} - \sum T_{downtime}}{N_{failures}}$$

Reliability measures how long a system can operate correctly without interruptions.

- Availability: the readiness of the system to deliver the intended service when requested. While reliability and availability are related, they are not identical. For example, a web service that fails briefly every few seconds may still be considered highly available (since the downtime is minimal) but not reliable, because the mean time between failures is very short.
- Maintainability: the ease with which a system can be modified, repaired, or restored to operational status. Although maintainability is rarely the target of direct attacks, it significantly affects long-term system dependability.
- Safety: the property ensuring that the system does not cause harm to the environment, users, or external entities during operation.

- Confidentiality and Integrity: these are also attributes of dependability, since a dependable system must protect data privacy and ensure correctness.

A secure system must also be dependable. For instance, a system designed to resist Denial of Service (DoS) attacks inherently possesses higher dependability, since it remains operational even under heavy load or abnormal conditions. Thus, measures that improve availability, such as redundancy, fault tolerance, and load balancing, also contribute to improving the overall dependability of the system.

In summary, availability is not only a core component of the CIA triad but also a critical pillar of dependability. Ensuring high availability means ensuring that the system continues to deliver its services correctly, reliably, and safely, even in the presence of faults, failures, or attacks.

2.1.2 Cryptography

Cryptography plays a fundamental role in information security, particularly within the domain of Protection. The National Institute of Standards and Technology (NIST) defines “Protection” in its Cybersecurity Framework as the set of actions and mechanisms aimed at safeguarding systems and information from cyber threats. Among these mechanisms, cryptography represents one of the most essential defensive tools.

While cryptography is typically associated with defensive applications, it is important to acknowledge that it can also be exploited by adversaries. For instance, many modern malware variants employ encryption techniques to conceal their payloads, thereby evading detection once executed on a victim’s system. Moreover, attackers often require strong expertise in cryptography in order to bypass or break existing cryptographic protections.

In this thesis, the discussion of cryptography is limited to conceptual and foundational aspects. The detailed study of specific symmetric and public-key algorithms is intentionally excluded, as the goal is to provide a conceptual framework for understanding how cryptography contributes to protection mechanisms within information security systems.

Basic Structure of an Encryption Scheme

An encryption scheme is formally defined by the following components:

- The set of all possible keys, called the *keyspace*, denoted by K ;
- The set of encryption functions, denoted by E_k , where each function operates under a specific key $k \in K$;

- The set of decryption functions, denoted by D_k , which are the inverses of the corresponding encryption functions.

Each encryption function E_k maps a message from the plaintext space to the ciphertext space, while each decryption function D_k maps ciphertexts back to plaintexts. For every key $k \in K$, there exists a unique decryption key $k' \in K$ such that:

$$D_{k'}(E_k(M)) = M$$

for every message M . In other words, each encryption operation has a corresponding inverse decryption operation that restores the original plaintext.

Symmetric and Asymmetric Encryption

The distinction between symmetric and asymmetric encryption schemes lies in the relationship between the encryption and decryption keys:

- In symmetric encryption, the encryption and decryption processes use the same secret key, or two keys that can be trivially derived from one another. Consequently, both parties must share this secret key securely before encrypted communication can take place.
- In asymmetric encryption (also known as *public-key cryptography*), the computation of the decryption key from the encryption key is computationally infeasible. The two keys, called the public key and the private key, play asymmetrical roles: the public key is openly distributed, whereas the private key must remain secret.

A key advantage of asymmetric encryption over symmetric encryption lies in the elimination of the key exchange problem. In symmetric encryption, communicating parties must first exchange the secret key through a secure channel before encrypted communication can occur. In contrast, in asymmetric encryption, a sender (e.g., Bob) can use the receiver's public key (e.g., Alice's) to encrypt a message. Only Alice, who holds the corresponding private key, can decrypt it, removing the need for a secure pre-shared channel.

Principles of Confusion and Diffusion

According to Claude Shannon's information theory, any robust encryption scheme should be designed according to two fundamental principles:

- **Confusion:** The relationship between the ciphertext and the encryption key should be as complex and unpredictable as possible. This ensures that even if an attacker analyzes the ciphertext, it remains infeasible to deduce the key.

- **Diffusion:** Each element of the plaintext should influence multiple parts of the ciphertext, distributing information across the entire encrypted message. As a result, the ciphertext should appear statistically random, preventing any observable structure that could aid cryptanalysis.

These two principles jointly ensure that an encryption algorithm produces a ciphertext that is indistinguishable from random data, making cryptanalysis significantly more difficult.

2.1.3 Cryptographic Hash Functions

Cryptographic hash functions differ from non-cryptographic hash functions. Like their non-cryptographic counterparts, they map an input domain of arbitrary size to a fixed-size codomain. Formally, given a message M belonging to a domain D , the hash function h produces a digest $d = h(M)$ belonging to a codomain C of fixed length k bits.

The essential property of any hash function is that the mapping from D to C is many-to-one, meaning that multiple inputs in the domain are mapped to the same output value. This is an inevitable consequence of the fact that $|D| \gg |C|$, since the set of all possible input messages is vastly larger than the finite set of possible digests.

A fundamental requirement for both cryptographic and non-cryptographic hash functions is that the mapping be uniformly distributed. This means that collisions instances where two distinct inputs map to the same digest should occur with equal probability across the entire domain. The uniform distribution ensures that no particular output value is more likely to appear than another.

Although collisions are mathematically unavoidable (as demonstrated by the *Pigeonhole Principle*), a good cryptographic hash function makes finding such collisions computationally infeasible. The security of a hash function does not rely on the absence of collisions, but rather on the practical impossibility of discovering them.

Properties of Cryptographic Hash Functions

A function $h : D \rightarrow C$ qualifies as a cryptographic hash function if it satisfies the following properties:

1. **One-wayness (Preimage Resistance):** Given a digest d , it must be computationally infeasible to find any message M such that $h(M) = d$. In other words, while it must be easy to compute the digest of a message, it should be infeasible to reverse the process. An attack that violates this property is called a *preimage attack*.

2. Second Preimage Resistance: Given a message M_1 , it must be infeasible to find a different message M_2 such that $h(M_1) = h(M_2)$. This property ensures that an attacker cannot replace a legitimate message with another that produces the same digest. Breaking this property corresponds to a *second preimage attack*.
3. Collision Resistance: It must be infeasible to find any two distinct messages M_1 and M_2 such that $h(M_1) = h(M_2)$. This is a stronger requirement than the previous one, as the attacker is free to choose both messages. An attack that violates this property is known as a *collision attack*.

Domain and Codomain Considerations

Cryptographic hash functions operate on input messages of arbitrary length without the need to define an upper bound. This is achieved through an iterative structure, in which the input message is divided into fixed-size blocks that are processed sequentially. Each block contributes to updating the internal state until the final digest is produced.

For example, in commonly used hash functions such as SHA-256, the codomain size is fixed at 256 bits, while the domain may include all possible binary sequences of any length. As a result, the number of potential input messages is astronomically larger than the number of possible outputs, implying that an enormous number of collisions exist in theory but are computationally unfeasible to find.

Comparison with Non-Cryptographic Hashes

It is important to distinguish between cryptographic and non-cryptographic hash functions:

- Hash functions for data structures (e.g., modular functions) are designed primarily for efficiency and uniform distribution, not for security.
- Weak hashes, such as cyclic redundancy checks (CRC), are used for error detection in data transmission and storage. They are not cryptographically secure and can easily be forged.
- Cryptographic hashes, on the other hand, are designed to withstand adversarial attempts to invert or collide them, making them suitable for digital signatures, integrity verification, and authentication mechanisms.

2.1.4 Digital Signature

A digital signature provides the security service of non-repudiation, has legal validity, and is cryptographically verifiable. Consider two entities, Alice and Bob. Alice

wishes to send Bob a digitally signed message. Both Alice and Bob possess a pair of cryptographic keys: a public key and a private key. The private key must remain secret, while the public key is openly distributed.

Signature Generation and Verification

Let M be the message to be signed, which can be any digital object (e.g., a document or file). Alice first computes the cryptographic hash of M using a secure hash function such as SHA-512, obtaining a fixed-length digest $d = h(M)$. Alice then signs the digest d using her private key, that is, Alice encrypts d with her private key, producing the digital signature σ .

Alice sends to Bob:

- the original message M ;
- the digital signature σ ;
- the identifier of the hash function used (e.g., SHA-512).

Upon receiving these elements, Bob verifies the signature as follows:

1. Bob applies the same hash function to the received message, obtaining $d' = h(M)$;
2. Bob decrypts the digital signature σ using Alice's public key, recovering the digest d originally computed by Alice;
3. If and only if $d = d'$, the signature is valid and the message is guaranteed to be authentic and integral.

If even a single bit of M has been altered during transmission, the two digests will differ, invalidating the signature. Authenticity is ensured because only Alice, who possesses the private key, could have generated the signature; integrity is ensured by the hash comparison; and non-repudiation is achieved because Alice cannot deny having signed the message.

Functions of a Digital Signature

A digital signature fulfills two key functions:

- Indicative function: the signature identifies the signer's identity;
- Declarative function: the signature expresses the signer's explicit approval of the content being signed.

The declarative function is achieved through the cryptographic signing process, while the indicative function relies on the Public Key Infrastructure (PKI).

Public Key Infrastructure and Certificates

A PKI enables the validation of public keys through digital certificates. Digital signature certificates follow the international standard X.509 [7], which defines the structure and attributes of certificates issued by a trusted Certification Authority (CA).

A CA is a trusted third party responsible for verifying the identity of the certificate holder before issuing the certificate. In Italy, for instance, CAs must be qualified and accredited by the *AgID* (Agenzia per l'Italia Digitale). The CA generates the key pair (public and private) for the subject, includes the public key within the certificate, and signs the certificate using its own private key to ensure authenticity and integrity. Each CA also possesses its own key pair; its public key is included in the certificates it issues.

Structure of an X.509 Certificate (v3)

- V: Version;
- SN: Serial Number;
- AI: Algorithm Identifier;
- ISSUER: Certification Authority (CA);
- VALIDITY: Time period (not before / not after);
- SUBJECT: Certificate owner;
- SUBJECT PUBLIC KEY INFO: Algorithm and public key value;
- ISSUER ID / SUBJECT ID: Identifiers;
- EXTENSIONS: Additional fields such as key usage;
- CERTIFICATE SIGNATURE: Signature of the certificate by the CA.

Among the extensions, the Key Usage field is particularly relevant, as it defines how the key can be used. Some common values include:

- Digital Signature: Used for authentication and data origin verification;
- Non-Repudiation: Ensures that the signer cannot deny having signed;
- Key Encipherment: Used in protocols that encrypt session keys (e.g., SSL/TLS);
- Data Encipherment: Used to encrypt data other than keys;
- Key Agreement: Used in key exchange protocols such as Diffie–Hellman;
- Certificate Signing: Used by CAs to sign other certificates;
- CRL Signing: Used to sign Certificate Revocation Lists.

Digital Signature Formats

Two widely adopted digital signature formats are CAdES and PAdES.

CAdES (CMS Advanced Electronic Signature)

CAdES can be viewed as a container or envelope that includes all the elements required to verify a digital signature:

1. The original document;
2. The signature (the digest encrypted with the signer's private key);
3. The cryptographic hash function used;
4. The signer's certificate containing the public key;
5. (Optionally) a trusted timestamp.

The resulting file conforms to the PKCS#7 standard and typically uses the `.p7m` extension (e.g., `document.pdf.p7m`). This format is enveloping, meaning it can contain documents of various types (PDF, JPEG, TIFF, etc.).

PAdES (PDF Advanced Electronic Signature)

PAdES is specifically designed for PDF documents and cannot be applied to other file types. It supports the same verification elements as CAdES but preserves the PDF file format, producing a signed document with the same `.pdf` extension.

2.1.5 Secure Socket Layer

Secure Socket Layer (SSL) and Transport Layer Security (TLS) are two cryptographic protocols that allow the creation of a secure connection between a client and a server. The main features of TLS are:

1. parameter negotiation between the client and the server involved in a connection;
2. server authentication at the client;
3. confidential communications;
4. data integrity protection.

In the stack ISO/OSI, TLS is placed between the Application Layer and Transport Layer. HTTP over TLS is named HTTPS (Secure HTTP) and is generally available on the 443 standard port. HTTPS uses a TLS certificate, which follows the X.509 standard and associates the domain name of the server with the public key of the server. The TLS certificate is signed by a trusted Certification Authority (CA), and most common client browsers have pre-installed several certificates issued by trusted CA that are globally recognized. After the connection request of the browser, the server sends to the client its TLS certificate, and the client uses the server's public key to encrypt a challenge that only the server can decrypt with the private key. Moreover, the client inspects the certificate to verify that the public key is associated with the correct domain name.

2.2 Blockchain

Blockchain is a peer-to-peer network that is part of Distributed Ledger Technology (DLT). A blockchain enables the creation of transactions based on a distributed consensus protocol among network participants without a central authority. The designer of blockchain technology is Satoshi Nakamoto, who made public the blockchain Bitcoin [8]. Unlike traditional databases, for data integrity, blockchain uses a Merkle tree, a hierarchical data structure that enables efficient integrity checking of large datasets by using hash values without the need for storing or transmitting the entire dataset [9]. Information consistency is guaranteed using consensus algorithms, such as Proof of Work (PoW) or Proof of Stake (PoS).

A blockchain can be:

- Public: information is openly accessible, and anyone can participate.
- Private: the network is managed by a centralized authority that restricts participation.
- Consortium: a hybrid solution in which consensus is controlled by a group of organizations.

2.2.1 Core Concepts of Blockchain Systems

A blockchain is composed of several fundamental elements that together define how the system records, validates, and protects information:

- Nodes. Nodes are the participants of the blockchain network, usually represented by the servers of each participant. Each node maintains a copy of the ledger and contributes to validating transactions.
- Transactions. A transaction consists of data that represents the values being exchanged and must be verified, approved, and eventually stored in the blockchain. Transactions are the atomic units of information.
- Blocks. A block is a collection of transactions grouped together to be verified, approved, and archived. Each block contains its own hash, a pointer to the hash of the previous block, a nonce, and metadata such as timestamp and block number. This structure ensures immutability and preserves the chronological order.
- Ledger. The ledger is the public register in which all transactions are recorded in an ordered, sequential, and immutable fashion. Every participant of the blockchain has access to the ledger, and in public blockchains, it can also be inspected by external users.
- Hash Functions. A cryptographic hash function transforms an input of arbitrary length into a fixed-length output (the digest). In blockchains, hash functions serve multiple purposes:

- they are used in digital signatures to guarantee authenticity;
- they link each block to the previous one by including the previous block's hash in the header;
- they allow efficient verification of transactions inside a block via Merkle trees.

The security of the blockchain depends on the properties of the hash function, namely pre-image resistance, collision resistance, and second pre-image resistance.

These components are tied together by cryptographic algorithms. Distributed Ledger Technologies (DLTs) rely on the use of public-key cryptography, where each user is identified by a public/private key pair. The private key is used to sign transactions, while the public key is shared with the network to allow other nodes to verify the authenticity of the transaction. Importantly, the blockchain does not provide a direct association between a public key and a real-world identity: a node may know another node's public key, but not the identity of the entity controlling it. This pseudonymity is one of the core features of blockchain-based systems.

Each block is linked to the previous one using a security hash, creating a chain of blocks that is resistant to tampering. The block header typically contains the nonce, the hash of the previous block, and the Merkle root of all transactions included in the block. The hash of the current block then becomes part of the next block, guaranteeing both the integrity and the temporal order of all transactions.

The transaction itself generally includes:

- the sender, who initiates the transaction and authorizes it by digitally signing it with their private key;
- the receiver, who is the intended recipient of the transaction and whose public key (or address) identifies the destination;
- the content, which may be a cryptocurrency value or any other type of asset;
- the encryption code or digital signature, used to ensure authenticity and prevent tampering.

Once validated, multiple transactions are grouped into a block, and the block is appended to the chain.

From these design choices, blockchain ensures several fundamental properties:

- Immutability: once information is stored in the blockchain, it cannot be altered or removed.
- Transparency: the ledger can be inspected by all participants, and in public blockchains, even by external users.

- **Traceability:** through the hash mechanism, every input and output associated with a transaction can be traced back, enabling full accountability.
- **Security:** cryptographic techniques protect transactions and block data against tampering and forgery.

Finally, the process of mining (in Proof-of-Work systems) or validation (in alternative consensus protocols) serves to verify and confirm the transactions within a block. The actors performing mining are called miners. Through this process, the block is added to the chain, and the ledger is updated across all nodes.

2.2.2 Consensus Mechanisms: Proof of Work and Proof of Stake

Consensus mechanisms ensure that all participants in a distributed ledger agree on the state of the blockchain, preventing conflicts such as double-spending and ensuring the immutability of transactions. The two dominant approaches are Proof of Work (PoW) and Proof of Stake (PoS), each implementing consensus through fundamentally different principles of trust and validation. A visual comparison of the two mechanisms is presented in Figure 2.1.

Proof of Work (PoW)

In PoW, each node, called a miner, competes to solve a cryptographic puzzle by finding a hash value that satisfies specific conditions (typically a number of leading zeros). The solution requires high computational power and energy consumption, since miners must iteratively vary a parameter known as the nonce until a valid hash is found. Once a miner discovers a valid block, it broadcasts the block and its proof to the entire network. The other nodes independently verify the validity of the proof-of-work, and, if the verification succeeds, the block is accepted and appended to the blockchain. The miner receives a reward, consisting of newly minted cryptocurrency (the *coinbase*) and transaction fees. PoW provides strong security guarantees because modifying past blocks would require redoing all subsequent computations, making attacks computationally infeasible. However, it is energy-intensive and can lead to centralization, as miners with more powerful hardware have higher chances of success.

Proof of Stake (PoS)

PoS replaces computational competition with an economic model. Nodes, called validators, do not mine but stake a certain amount of cryptocurrency as collateral to participate in block validation. A validator is chosen pseudo-randomly often based

on the amount of stake and the duration it has been locked to propose and attest new blocks. If a validator acts maliciously (e.g., approving invalid transactions), part of its stake is forfeited, a process known as slashing. This mechanism aligns economic incentives with network security: the higher the stake, the greater the responsibility and the potential loss for dishonest behavior. PoS drastically reduces energy consumption and increases throughput, but introduces challenges such as stake concentration and the so-called “nothing at stake” problem, where validators could theoretically support multiple chains at once.

Comparison

PoW secures the network through computational effort, while PoS relies on economic incentives. PoW imposes costs in terms of energy and hardware resources, while PoS shifts the cost to the amount of capital staked. PoW guarantees security through physical work, PoS through financial accountability. For these reasons, several modern blockchains, including Ethereum have transitioned from PoW to PoS to improve scalability and sustainability.




Fig. 2.1: Comparison between Proof of Work and Proof of Stake consensus mechanisms.

2.2.3 Ethereum

Ethereum [10] is a platform for developing decentralized applications (DApps) that require fast development, security, and interoperability. To support computa-

tion within the network, Ethereum introduced its own cryptocurrency, *Ether* (ETH), which is used to pay for the computational resources necessary to execute programs or smart contracts.

DApps in Ethereum are written in Solidity, a Turing-complete programming language designed to allow anyone to create smart contracts. Smart contracts are self-executing pieces of code stored on the blockchain, which run deterministically when invoked by users or other contracts. These contracts are executed in the Ethereum Virtual Machine (EVM), a decentralized virtual machine that ensures all nodes in the network agree on the outcome of contract execution. The development pipeline involves writing Solidity code, compiling it into EVM bytecode, and deploying it to the blockchain, where it is executed by Ethereum nodes.

Ethereum as a State Machine

Ethereum can be understood as a global state machine. At any given time, the state of Ethereum is represented by a collection of accounts, each identified by a 20-byte address derived from a public key. Accounts can be of two types:

- Externally Owned Accounts (EOAs): controlled by users through private keys, able to initiate transactions.
- Contract Accounts: controlled by smart contract code, activated when receiving a transaction or message.

The state transition function updates the global state after each transaction, moving from the old state S to the new state S' .

Transactions and Messages

A transaction is a signed data packet sent from an EOA, which includes the recipient address, an optional value in Ether, a data field (for contract interaction), and two parameters: `STARTGAS` (the maximum number of computational steps allowed) and `GASPRICE` (the fee per step the sender is willing to pay). The product of these two values determines the cost of the transaction, denominated in gas. Every operation in Ethereum, whether sending Ether or executing a smart contract, requires gas to be paid.

Contracts can also send messages to each other. Unlike transactions, messages are not recorded on the blockchain directly, but exist only within the EVM execution context. They allow contracts to call other contracts, transfer Ether, or trigger new contract code.

Accounts and State

Each Ethereum account maintains four fields:

1. a *nonce*, used to ensure transactions are processed once,
2. the account's *current balance*,
3. optional *contract code* (for contract accounts),
4. *storage*, a persistent data area associated with the account.

State updates result from transactions and contract executions, with each block recording the new state root hash to ensure integrity.

Consensus and Mining

At the time of its launch, Ethereum relied on Proof of Work (PoW), where miners competed to validate blocks and were rewarded with newly minted Ether plus transaction fees. More recently, Ethereum has transitioned to Proof of Stake (PoS), where validators stake ETH to propose and attest blocks, drastically reducing energy consumption.

Exploring Ethereum

Blockchain explorers such as Etherscan allow inspection of the Ethereum ledger. By opening a block, one can observe its timestamp, the number of transactions, internal contract calls, the miner or validator, associated fees, difficulty, gas usage, and gas limits. By opening a transaction, it is possible to view sender, receiver, amount, execution cost, and event logs generated by smart contracts.

Performance and Limitations

Ethereum provides relatively fast block confirmation times, which makes it significantly quicker than Bitcoin's block interval. However, faster confirmation introduces security challenges, such as block propagation delays and potential forks due to network latency. These are mitigated by Ethereum's consensus protocols and eventual finality mechanisms.

Overall, Ethereum supports secure and transparent exchange not only of currency but also of logic and data through smart contracts and decentralized applications.

2.3 Self-Sovereign Identity

Self-Sovereign Identity (SSI) offers a new way to manage digital identities based on decentralization. The main innovation of this paradigm is that the locus of control shifts to the owner of an identity (e.g., person, animal, organization, or object) who can interact with everyone else as a full peer. SSI can be defined as a set of technologies that are put together to create a new model for digital identity management. In this section, we discuss the most relevant features of each building block of SSI that have been used in our proposal.

2.3.1 Verifiable Credentials

The Verifiable Credential (VC) is a set of information that some authority claims to be true about a specific entity that is the subject of the credential. In SSI, the authority is called *issuer*, while the subject of the VC is called *holder*. The holder can prove something about themselves by presenting their VC to a verifier. At a high level, the structure of the VC includes four sections:

- Identifier of the VC;
- VC metadata (e.g., issue date, expiration date, state of the credential);
- VC's claims about the subject;
- Proof, which is the digital signature of the issuer of the credential. It can be verified through the blockchain.

2.3.2 Issuers, holders, and verifiers

Issuers, holders, and verifiers are the three roles involved in the SSI system, also known as the Trust Triangle. The issuers are the source of VCs, and they are usually some authorities, such as universities or companies. The holders are the subjects of the credential who are responsible for storing the VCs in a digital wallet. The holders can be individuals, but also organizations or objects. Lastly, the verifiers are the entities that offer a benefit to the holder after the execution of a verification process of one or more credentials. The verifiers can be individuals or organizations. The three actors are involved in several interactions, such as requesting, issuing, and verification of the VCs. One of the most relevant interactions among SSI roles is the verification of one or more VCs presented by the holder to the verifier. The verifier must define an access control policy for each benefit they offer. In the SSI world, this policy can be seen as the set of VCs that the holder must present to obtain access to the service. After the presentation process, the verifier starts the verification, which is the union of several steps:

- VC's state verification: the verifier must check that the VC is still valid so that it has not been revoked.
- VC's claims verification: the verifier must check that the information contained in the VC satisfies the access control policy.
- VC's issuer verification: the verifier must check the authenticity of the VC's issuer. The VC includes the issuer's Decentralized Identifier (DID), allowing the verifier to retrieve the issuer's public key from the blockchain to validate the VC's digital signature.

2.3.3 Digital Wallet

The digital wallet is the location where the VCs are stored securely. It is a software application installed in a generic device such as a smartphone, tablet, or laptop. The owner of the digital wallet is the holder of the VCs. Strong authentication processes must be implemented to access the wallet because it contains private information. The digital wallet is configured to maintain:

- Verifiable Credentials (VCs) issued to the holder of the digital wallet.
- Peer Decentralized Identifiers (Peer DIDs) used to establish the peer-to-peer connections.

2.3.4 Digital Agents

The Digital Agent is a software module that helps the user during interactions with their digital wallet. The digital agent executes several actions on behalf of the owner of the wallet. As an example, the digital agent generates a new peer DID when the owner wants to establish a new peer-to-peer connection, or it stores the VCs in the wallet when the owner accepts it.

The peer-to-peer connections, which have been mentioned above, are widely used in our proposal. A peer-to-peer connection is established off-chain after some actions are taken by the controller (owner of the wallet) as a QRCode scan or a click of a link. If the controller authorizes the connections, the digital agent will prepare all the cryptographic materials necessary to establish the connection. On these connections, the VCs can be exchanged between the holder, issuers, and verifiers.

2.3.5 Decentralized Identifies

Decentralized Identifiers (DIDs) are a type of URI that uniquely identifies an entity in the SSI system. Each DID resolves to a DID Document, which includes details about the identified entity (e.g., the public key for the digital signature verification).

There are two types of DIDs:

- **Ledger-based DIDs:** this type of DID involves a blockchain or a generic distributed ledger. It is typically the blockchain address. It is created, updated, and deactivated by performing a transaction.
- **Peer-DID:** this type of DID does not involve a blockchain. Each endpoint of a peer-to-peer connection generates a Peer DID that is securely shared with the other peer entity using a specific communication protocol.

2.3.6 Verifiable Data Registry

A Verifiable Data Registry is used in SSI to obtain the property of decentralization to create an identity system that supports privacy without a central authority. The verifiable data registry is typically a blockchain used to manage the DIDs and the public keys. The blockchain must be public [11] to enable access to everyone, and can be permissioned or permissionless. It is important to store the DID of the issuers on the blockchain to facilitate the verification process carried out by the verifiers.

2.3.7 Governance Framework

The Governance Framework is the last block of the SSI architecture and involves social and human aspects. Indeed, the Governance Framework simplifies the job of the verifier. A governance authority publishes a list of trusted issuers. When a verifier receives a new VC issued by a certain authority, they can check the presence of that issuer in the list. If they are present, then they are trusted.

2.4 Lawful Interception

Lawful Interception is a procedure in which a network operator provides a law enforcement agency with intercepted communications of private individuals. In the description of this section, we make reference to the ETSI standard [12] adopted in many countries of the world. Moreover, the considerations emerging in this thesis also apply when considering other standards such as CALEA [13] (adopted in the USA).

The LI procedure involves three main actors: the Law Enforcement Agency (LEA), the Law Enforcement Monitoring Facility (LEMF), and the Network Operator.

- The LEA is the Law Enforcement Agency that is involved in a certain investigation about one or more individuals;
- The LEMF (Law Enforcement Monitoring Facility) is the collection of the technical equipment used to receive the result of the interception from the network operator;




Fig. 2.2: ETSI-defined LI architecture.

- The Network Operator is the telecommunication operator offering the network platform to the individuals under investigation and performs the interception in accordance with legal mandates.

This activity is strictly regulated and must comply with both national and international standards, ensuring that interception is verifiable, auditable, and limited to authorized cases.

Figure 2.2, adapted from the ETSI standard [12], illustrates the overall LI architecture. It distinguishes two primary domains: the Network Operator Domain managed by the Network or Access Provider (NWO/AP/SvP) and the LEA Domain, operated by the LEMF.

Within the operator's domain, the innermost circle represents the Network Internal Function (NIF), which contains all network functionalities such as routing, switching, and session management. From this network core, two types of data are generated:

- **Intercept Related Information (IRI):** metadata describing the intercepted communication (e.g., caller and caller identifiers, timestamps, and protocol-specific information);
- **Content of Communication (CC):** the actual payload of the intercepted communication (e.g., audio, video, messages, or IP traffic).

Both IRI and CC originate from the Internal Interception Function (IIF), which exposes these data to the Internal Network Interface (INI). The INI serves as a bridge between the IIF and the mediation devices. Two dedicated mediation functions then process and transmit the information:

- the IRI Mediation Function, responsible for encoding and transmitting IRI data to the LEMF through the HI2 interface;
- the CC Mediation Function, responsible for the handling and transmission of CC data via the HI3 interface.

The LEMF collects, stores, and analyzes both IRI and CC as a result of the interception. Depending on the national legal framework, the LEMF may be managed directly by the LEA or delegated to an external entity acting under its authority.

2.4.1 ETSI Handover Interfaces

Three standardized handover interfaces: HI1, HI2, and HI3, define how administrative, signaling, and content data are transferred between the Network Operator and the LEMF. Each interface serves a specific function in the interception workflow.

HI1 - Administrative Interface

HI1 is bidirectional and connects the LEA to the operator's Administration Function (AF). It is used to exchange administrative and judicial data, such as:

- activation, suspension, or termination of an interception order;
- legal authorizations (start and end dates);
- identification of the target (e.g., MSISDN, IMSI, IMEI, email address, or IP address).

This interface ensures that interception activities are formally authorized, traceable, and auditable.

HI2 – Intercept Related Information (IRI)

HI2 carries metadata related to intercepted communications from the operator to the LEMF. IRI parameters are encoded using *ASN.1* and *BER* encoding, as defined in ETSI[12]. Transmission can occur:

- via secure FTP over TCP/IP (ASN.1 formatted files); or
- via IP streaming within a dedicated VPN port (e.g., port 12500).

Each IRI record contains session identifiers, timestamps, and other service-specific information about intercepted connections.

HI3 – Content of Communication (CC)

HI3 transports the actual intercepted communication content from the operator to the LEMF. The CC may be transmitted as a clear-text copy or, if encrypted, accompanied by the decryption key sent via HI2. Because communications are bidirectional, two distinct channels (toward and from the target) are typically established. Examples include telephone audio, SMS, MMS, VoIP traffic, and web sessions.

Table 2.1: Summary of ETSI handover interfaces in Lawful Interception systems. Each interface defines the type and direction of information exchanged between the Network Operator and the Law Enforcement Monitoring Facility (LEMF).

Interface	Type of Information	Direction	Protocol/Encoding	Example Use
HI1	Administrative data (activation, deactivation, extensions)	Bidirectional	TCP/IP, XML-based messaging	Judicial orders, target identification, service activation/deactivation
HI2	Intercept Related Information (IRI): metadata of intercepted communications	Operator → LEMF	ASN.1 BER encoding over secure FTP or IP streaming	Transfer of session identifiers, timestamps, service-specific metadata
HI3	Content of Communication (CC): actual intercepted data (voice, video, IP traffic)	Operator → LEMF	IP streaming or dedicated channel; clear-text or encrypted copy	Delivery of intercepted content and associated decryption keys

2.5 Generative AI

Artificial Intelligence has a history spanning over 50 years, dating back to the 1960s when rule-based systems were first introduced [14]. Over the years, rule-based models were improved using statistical machine learning approaches and multi-layered neural networks that allowed for the development of natural language processing applications. Traditional language models relied on sequential processing that made it difficult to capture long-range dependencies in data. The advent of transformer architectures in 2017 [15] addressed this limitation by allowing all elements of the sequence to be processed simultaneously, enabling the model to capture relationships between words regardless of their position in the sequence. This innovation, coupled with massive datasets, improved computational power, and model architec-

ture advancements led to the development of large language models (LLMs) like the generative pre-trained transformer (GPT) [16].

2.5.1 GenAI Systems and Browser Assistants

Generative AI (GenAI) systems typically rely on large language models (LLMs), which operate by iteratively predicting the next token in a sequence based on the context of preceding tokens [17]. These models are built on transformer architectures, where a self-attention mechanism enables the model to assign varying importance to different parts of the input sequence [15]. When a user inputs a query, the model encodes the input into an embedding space and generates a probability distribution over possible next tokens. It then selects a token based on this distribution and repeats the process until the output is complete. This approach enables GenAI systems to provide coherent and context-aware responses to user input.

GenAI systems face several challenges that affect their performance, usability, and broader applicability. Understanding these limitations is essential for evaluating how these systems function in practice. The first fundamental issue is their *probabilistic* nature. As described earlier, GenAI systems generate the next token based on a probability distribution, which can lead to varying outputs even for the same input. This variability is influenced by the *temperature* parameter, which controls the randomness of the output. Higher temperature values encourage the selection of less probable tokens, often producing more creative but potentially less coherent responses, while lower values lead to more deterministic and focused responses [18]. The second challenge is the resource-intensive nature of training these models [19]. This means that these models cannot incorporate up-to-date information that was not part of the training process without access to live search capabilities. Finally, while general-purpose models are versatile, adapting them to domain-specific use cases, such as personalized assistants is non-trivial. These typically have access to only the user-submitted queries within the platform, which limits their context to what is explicitly included in the user prompt and makes personalization challenging.

Gen AI browser assistants aim to leverage the capabilities of GenAI models to personalize the user experience on the web. Implemented as browser extensions, these assistants function as wrappers around GenAI models such as OpenAI’s ChatGPT [20], Google’s Gemini [21], Meta’s Llama [22], etc. They monitor user activity on the web and use this context to query backend GenAI APIs and return personalized responses. Unlike foreground web applications, these assistants operate through background service workers that run continuously while the extension is active. They can inject *content scripts* into every page that a user visits. These content

scripts may load additional libraries (e.g., jQuery) to facilitate actions like DOM manipulation, and can communicate with the background script to log and share user activity. This browser extension-based architecture effectively grants the assistants broad access to everything a user does in the browser. By monitoring a user's browsing activity across websites, browser assistants can offer in-the-moment suggestions tailored to specific browsing contexts. In summary, these assistants address many of the limitations of general-purpose GenAI systems and offer a more personalized, context-aware experience.

GenAI browser assistants offer a range of features that we group into three broad categories: search-based assistants, content-based assistants, and automation assistants. Search-based assistants offer features such as natural language responses to search queries, support web page summarization, interactive Q&A with web pages, and text highlighting. Content-based assistants assist with content-specific tasks such as writing tasks, social media post generation, meeting transcript summarization, YouTube video explanation, SEO generation, etc. Automation assistants focus on streamlining workflows by supporting tasks such as auto-filling forms, coding and debugging agents, task scheduling, automated data extraction and scraping, voice-enabled assistance, etc.

2.5.2 Traditional vs. GenAI Browser Extensions

While browser extensions have long been supported in browsers, the integration of GenAI in browser extensions introduces new risks through expanded data collection and inference capabilities. To understand these risks, we highlight three key differences between traditional and GenAI-powered browser extensions:

GenAI extensions rely on third-party GenAI models to analyze a user's browsing activity. This allows them to build rich user profiles by combining observed browsing behavior with inferred characteristics such as interests and intent. Distinct from traditional extensions, GenAI extensions can personalize responses based not only on what users explicitly do, but also on what the LLM infers about them through reasoning.

Traditional extensions typically collect limited, task-specific data such as URLs or selected DOM elements. For instance, a price comparison extension may only extract product names and prices from searches on a specific site like Amazon. On the contrary, GenAI extensions often collect and transmit full-page content or the entire DOM.

Many GenAI browser assistants include a memory component that persists across navigations, sessions, or tabs. This enables longitudinal tracking and cross-context profiling, which is not common in traditional extensions.

Privacy-preserving Attribute Certification

In today's age, ensuring the integrity and authenticity of digital information is an important aspect in various domains such as healthcare and education. Unlike traditional digital certificates that primarily authenticate identity, attribute certificates provide additional information about a subject's attributes, such as roles, qualifications, or membership status. The General Data Protection Regulation (GDPR) underlines the principle of data minimization, emphasizing the collection and processing of only necessary personal data for specific purposes. In this chapter, we deal with the need to certify attributes and satisfy the data minimization principle, thus designing a technique for certifying information at various levels of detail. The proposed framework utilizes blockchain smart contracts to automate the certification process, ensuring the authenticity and integrity of certified data. A proof of concept of the technique is presented in the context of certification in online education.

3.1 Introduction

In an era where digital information reigns supreme, ensuring its authenticity and integrity has become a concern across various domains ranging from financial transactions to healthcare records. Information certification provides a robust solution to guarantee the trustworthiness of data.

Differently from digital certificates that primarily focus on authenticating the identity of the subject, attribute certificates provide additional information about the subject's characteristics. Indeed, an attribute certificate is a digital certificate that associates a set of descriptive data items with a subject name or an identity certificate [23]. In this context, attribute certification refers to the process of verifying specific characteristics (such as role, membership status, and qualifications) associated with an individual or entity and issuing a certificate attesting to those characteristics.

Once the characteristics are verified, typically a trusted entity issues an attribute certificate containing the verified information along with cryptographic information

used to ensure integrity and authenticity of the certificate. Furthermore, attribute certificates have a finite lifespan and may need to be revoked if the claimed characteristics change or are missing.

When dealing with certified information, it is worth recalling that the General Data Protection Regulation (GDPR), which is the data protection law in the European Union [24], established the principle of data minimization as one of the key principles in data processing. Data minimization is aimed at ensuring that organizations collect and process only the personal data that is necessary for the specific purpose for which it was obtained. Thus, organizations should limit the collection of personal data to what is adequate, relevant, and necessary for the intended purpose [25], [26].

In this chapter, we propose a technique for certifying information at different levels in such a way that the user can select the detail level of information to disclose without the need of obtaining or creating a new certificate. We reach this goal through the utilization of the blockchain technology (see Section 2.2), a promising solution capable of providing a decentralized and immutable ledger for recording transactions and data interactions. Blockchain's inherent characteristics, including immutability, decentralization, and transparency, offer a robust foundation for building a secure and tamper-proof certification system [27], [28]. The proposed framework leverages blockchain smart contracts to create a distributed and tamper-resistant ledger for certifying information. Different details of information to be certified are implemented by using a different random secret for each degree of sensitivity and importance associated with the information. Smart contracts are employed to automate the certification process, providing a seamless and efficient mechanism for verifying the authenticity and integrity of certified data. Furthermore, the chapter shows a proof of concept of the technique in the context of certification in online education.

The chapter is organized as follows. In the next section, we survey related work. Section 3.3 defines the model of certificates. Section 3.4 presents the scheme that exploits a blockchain for certifying information at different levels of granularity. Section 3.5 discusses the scalability and complexity properties of the proposed scheme. Finally, Section 3.6 draws the conclusions.

3.2 Related Work

Blockchain technology has been adopted in several protocols and proposals in recent years, from healthcare to digital advertising, from digital identity to supply chains [29], [30].

Furthermore, blockchain technology has almost immediately aroused interest regarding the possibility of being used to certify information and records [31]. In [32], the authors proposed a framework based on the private blockchain Hyperledger to exchange and approve students' certifications electronically. In particular, their proposal needs web services to interface and manage the requests of a given information with the blockchain layer.

The authors of [33] propose a system based on Hyperledger using IPFS (Inter-Planetary File System) as storage to validate and collect educational certifications. In [34], the authors implement a solution based on Hyperledger as well. Their main aim is to provide a higher level of privacy to academic certificates. In particular, the proposal implements a solution that enables the owner of the certificate (e.g., a student) to control who has access to their academic certificates and when.

In [35], the authors propose VECefblock, a blockchain-based authentication approach for certificate authentication for the Vietnam educational system to counter fake certifications. This solution is based on Hyperledger and consists of three main layers involving both on-chain and off-chain steps: an application layer, a data modeling layer, and the Hyperledger layer. The authors of [36] propose another solution based on Hyperledger to store and manage digital certificates, comparing their Hyperledger-based approach with public blockchains like Bitcoin and Ethereum.

All the proposals presented above implemented their solutions with private and permissioned blockchains (e.g., Hyperledger) because public blockchains, like Ethereum, are designed to be transparent and immutable so that anyone can access the information stored [37]. On the one hand, this can represent a disadvantage because certificates could involve sensitive data. On the other hand, private blockchains need a central authority to coordinate and manage the information over the network other than well-suited strategies and access control policies. For this reason, the optimal solution would be to provide a mechanism based on public blockchains able to guarantee the confidentiality of such certificates. This is the main advantage of our proposal, which exploits a public blockchain, with respect to the techniques presented above.

In [38], the authors propose a solution based on the public blockchain Ethereum and smart contracts to manage certifications. In detail, their proposal splits and classifies the certificate data into three different types: certificate template, certificate receiver, and certificate information. This division helps increase efficiency when adding new certificates to the blockchain and/or the same entity delivers many certificates. However, such a solution has better performance when a trusted third party is involved. In contrast, our solution does not rely on a trusted third party.

The authors of [39] implement a web application that uses Ethereum and smart contracts to store and check Ph.D. certificates. In this proposal, the data is still stored on an off-chain DBMS, so that it loses all the decentralization and distributed characteristics of the Blockchain technology. In contrast, our proposal is fully decentralized.

3.3 Certificate Model

The purpose of this section is to model an attribute certificate in such a way that a system can manage certified information. We observe that there exist several standards that can be used to define the format of a certificate, and standard X.509 is surely the most known one [7] (see Section 2.1.4). As will be clear in the following, this type of certificate is oversized compared to our needs and, at the same time, it misses some features necessary to implement our proposal. Although we could extend standard X.509 to include the missing feature, we prefer to define a new model of an attribute certificate in order to simplify and make the presentation of the proposal self-contained.

We start the modeling of a (digital) certificate by defining the components of its structure and, then, we will present an example of an instance of the model.

Definition 3.1. *Given a certificate C , we define the issuer of C as the tuple $\langle \text{name}, \text{uri}, \text{address} \rangle$, where name is the name of the entity (typically an organization) that certifies the information, uri is a URL belonging to the domain of Web site of this entity, and address is a blockchain address.*

Definition 3.2. *Given a certificate C , we define a disclosed datum (DD) of C as a tuple $\langle \text{attribute}, \text{value} \rangle$, where attribute is the name of the certificated attribute and value is its value (for example, $\langle \text{date of birth}, "01/01/2000" \rangle$).*

Definition 3.3. *Given a certificate C , we define an optional datum (OD) of C as a tuple $\langle DD_1, \dots, DD_o, \text{nonce} \rangle$, where each DD is a disclosed datum and nonce is a random value.*

Definition 3.4. *A certificate C can be modeled as a tuple $\langle \text{issuer}, DD_1, \dots, DD_d, OD_1, \dots, OD_o \rangle$ (all the components have been defined above).*

To better explain how the model works, we show in Fig. 3.1 an example of a certificate that regards the university career of a person. Observe that the first column on the left has been added as a comment to mark each component (thus, this column is not represented in a real certificate); moreover, the included data are very simple and dummy for the sake of presentation.

<i>Issuer</i>	<University of World, https://www.world.uni/certificate address, 0x076A108AA9f3A1731Ac>
<i>DD₁</i>	<person name: "Gianluca Lax">
<i>DD₂</i>	<date of birth: "01/01/2020">
<i>DD₃</i>	<degree, "MSc in Computer Science">
<i>OD₁</i>	<degree date, "01/02/2020", efQ*W2T6h3h+w>
<i>OD₂</i>	<degree mark, "distinction", f-WfE9^ww>
<i>OD₃</i>	<passed modules, "Principles of Programming, Object-Oriented Programming, ...", Kf3)jfi(/fe>

Fig. 3.1: Example of the model of a certificate.

The first line refers to the *issuer* of the certificate: we find a dummy name of the university, a URL, and a Blockchain address (the reason why these fields have been included will be clear after the proposal formalization is presented).

The certificate includes three *disclosed data* reporting the name and birth date of the person and the obtained degree. Observe that *attributes* are in black and *values* are in green. These data are the minimum amount of information disclosed by the certificate and are used when the user needs to state to have an MSc in Computer Science (for example, when he applies for a job).

The remaining lines are three *optional data* and are used when the user needs to disclose other details about his MSc degree (specifically, the date, the mark, and the passed modules). Observe that *nonces of optional data* are in red and are encoded as strings.

3.4 Proposed Scheme

In this section, we present the scheme that exploits a blockchain for certifying information at different levels of granularity. In particular, given a certificate, the *disclosed data* are always certified, whereas the user can choose which combination of the *optional data* should be disclosed by suitably running the proposed scheme.

The main actors or components of our scheme are:

- a certification *authority*, which is enabled to issue certifications;
- a *user*, who needs to certify some personal attributes;
- a *verifier*, an entity that needs to verify a certification of a user;
- a *blockchain* supporting the execution of smart contracts.

The solution we propose is built on the following phases.

Phase 1: Setup. The solution we propose is based on the use of a blockchain that supports smart contracts. Indeed, a smart contract is used to certify *optional data* (see Definition 2 in Section 3.3. Such a smart contract is developed in this phase and presents three functions (named *add*, *verify*, *remove*) to add, verify, and remove a certification, which are described in the following. As an example of proof

of concept, we report in Fig. 3.2 a sketch of the code of the smart contract that will be used and commented in the following phases.

```

contract Verifier {
    // Map that associates a hash with a Blockchain
    address
    mapping (bytes32 => address) map;
    // Function to add the association between a hash and
    the address of the function caller
    function add(bytes32 hash) public {
        map[hash] = msg.sender;
    }
    // Function to verify whether a hash is associated
    with an address
    function verify(bytes32 hash, address
        address_blockchain) public returns (bool) {
        return map[hash] == address_blockchain;
    }
    // Function to remove a hash from the map
    function remove(bytes32 hash) public {
        if (map[hash] == msg.sender)
            map[hash]=null;
    }
}

```

Fig. 3.2: Sketch of the code of the smart contract.

Phase 2: Authority registration. An authority is the entity that issues certificates and should be trusted and identifiable. Observe that trust, intended as “*belief in honesty and goodness of an entity*”, is a subjective property, hence it is hard to capture in a quantified way [40]. Establishing whether an entity is trusted or not is out of the scope of this work, so we use the standard approach to deal with this issue. We assume that the authority has a website and each verifier can process the information reported on the website to establish whether such an entity can be considered trusted.

The second aspect we need to address regards how to identify an authority. For this purpose, we rely on a blockchain. Thus, we assume that an authority registers a blockchain address, which is associated with a pair of private/public keys. Then,

the authority publishes on its website the registered blockchain address in such a way the association between the authority and its blockchain address is publicly available.

Phase 3: Certification generation. Let us consider again the certificate example presented in Section 3.3 and reported in Figure 3.1. To generate the final certification, the issuer does the following operations:

1. For each *optional datum*, its cryptographic hash is computed;
2. Each *optional datum* is replaced by its cryptographic hash.

An example of a final certification is reported in Figure 3.3, where the cryptographic hash function used is Keccak-256. Observe that the transformed certificate reports in clear only *disclosed data*, whereas *optional data* are reported in an obfuscated way.

Finally, the issuer calls the function `add` of the smart contract for each generated cryptographic hash and for the cryptographic hash computed from the whole final certification. This function receives a cryptographic hash and maps this hash to the blockchain address of the caller. In the considered example, this function is called 4 times (3 times for the hashes of the 3 *optional data* and another one for the hash of the whole certificate). In this way, it is possible to associate each hash with a blockchain address.

Phase 4: Certification checking. In this phase, the user discloses some *optional data* to the verifier (recall that in the final certification, such data are represented in an obfuscated way). For example, suppose that the user wants to disclose OD_1 in the certificate of Figure 3.3. The user presents to the verifier both the whole certificate (i.e., the one reported in Figure 3.3) and the *optional datum* in plain text (i.e., $\langle \text{degree date, "01/02/2020", efQ*W2T6h3h+w} \rangle$).

The verifier performs the following operations:

1. Calculate the cryptographic hash of the whole certificate;
2. Calculate the cryptographic hash of each disclosed *optional datum*;
3. Extract from the certificate the `address` of the `issuer` (see definitions given in Section 3.3 and call the function `verify` of the smart contract for each cryptographic hash. It receives a cryptographic hash and a blockchain address and returns *true* if there is a mapping between this blockchain address and this hash (*false* otherwise). If even just one function call returns *false*, the verifier considers invalid the certification.

Finally, observe that some human checks should be performed. The first regards the identification of the user to verify that he/she is who claims to be. The second check consists of verifying that the blockchain address reported in the certificate is

actually of the authority. This can be done by accessing the authority website, where the blockchain address of the authority is publicly available (recall phase 2). Finally, the verifier should evaluate the trust in the authority that issued the certificate. As discussed above, this is a subjective evaluation that is out of the scope of the study.

Phase 5: Certification removal. In some cases, a certification (or an *optional datum*) can expire or be revoked (think of a driver's license, for example). Recall that each certification (and *optional datum*) is associated with a cryptographic hash, which should be removed from the mapping. Thus, in this case, the authority that issued the certificate should call the function `remove` of the smart contract. It receives a cryptographic hash and removes this hash from the mapping associated with the blockchain address of the caller. Thus, from now on, a verification of the certification associated with this hash will fail.

We can conclude this section with some final remarks.

<i>Issuer</i>	<University of World, https://www.world.uni/certificate_address , 0x076A108AA9f3A1731Ac>
<i>DD₁</i>	<person name: "Gianluca Lax">
<i>DD₂</i>	<date of birth: "01/01/2020">
<i>DD₃</i>	<degree, "MSc in Computer Science">
<i>OD₁</i>	22196f9da08f9d3f287c9385a1ddfc3d5aef35be4ad6358488ae799227a57748
<i>OD₂</i>	ca5afd723d1e2027ab0767305ab2969e1e435815b49b25b11e195d274f2814bc
<i>OD₃</i>	6ff574c8f0cf8985de8ff0d68f826e4bce71e944657fb570b75c1ab264f92afb

Fig. 3.3: Example of certification.

The use of a nonce in each *optional datum* is necessary to avoid that obscured values are guessed by a trial-and-error method. Indeed, in absence of a nonce, when the domain of the possible values is limited (e.g., the sex), an attacker could compute the hash for each possible value until the hash reported in the certificate is found. Using a nonce randomly generated in a suitable domain protects against this attack.

The use of the blockchain allows us to be certain that stored information (i.e., the hashes) cannot be compromised (see Section 2.2.1). Moreover, it is impossible for an entity to call a function from the blockchain address of another entity because each call is signed by a private key (we leverage the standard assumption that private keys are kept secret).

Note that the use of a blockchain allows us to avoid that certificates are signed, as typically occurs. Indeed, certificates are simple text documents and are not signed by any entity or authority. However, any change of the certificate after it is issued from the authority is detected. Indeed, after changing just even one bit of the certificate, the computed cryptographic hash is different from the one the authority has sent to

the smart contract. Consequently, the new hash will not be found when calling the function `verify`.

3.5 Discussion

The security of the proposed scheme relies on two standard cryptographic assumptions, whose formal definitions are provided in Section 2.1.3 of the Background chapter. First, the hash function Keccak-256 is assumed to be collision-resistant, ensuring that a forged certificate cannot match a legitimate hash stored on the blockchain. Second, Keccak-256 is assumed to be pre-image resistant, ensuring that undisclosed optional data cannot be recovered from their hashes, even when these are publicly visible on the blockchain. Both properties are widely accepted for Keccak-256 and are standard assumptions in the cryptographic literature. Regarding scalability, the primary concern under high transaction volumes is Ethereum network congestion, which can increase gas costs and confirmation latency. However, this is a property of the underlying blockchain infrastructure, rather than of the proposed scheme itself. Several established approaches mitigate this concern. Layer 2 solutions such as Ethereum Rollups bundle multiple transactions into a single on-chain operation, significantly reducing per-transaction costs and increasing throughput [41]. For deployments requiring higher transaction volumes, alternative public blockchains such as Solana [42] or TON [43] offer substantially higher throughput while maintaining the decentralization properties required by the proposed scheme. Furthermore, since certification operations are typically performed in batches by a certification authority rather than in real time by individual users, the system is naturally suited to off-peak scheduling, further reducing the impact of network congestion on practical deployments.

3.6 Conclusion

The evolving landscape of digital information requires robust mechanisms to ensure its authenticity, integrity, and privacy. While traditional digital certificates focus on authentication, attribute certificates are used for certifying specific characteristics associated with individuals or entities. The implementation of information certification must also adhere to the General Data Protection Regulation (GDPR), and, in particular, to principles like data minimization. This principle ensures that organizations collect and process only necessary personal data, safeguarding individuals' privacy rights.

In response to these challenges, this work proposed a novel technique for certifying information at different levels of detail, empowering users to selectively disclose information without the need for creating new certificates. The proposed solution leverages the blockchain technology and exploits smart contracts to automate the certification process and enhance efficiency and transparency.

We provided a proof-of-concept presented in the context of online education to underscore the feasibility and potential of the proposed technique in real-world applications. As digital ecosystems continue to evolve, innovative approaches to information certification, such as the one proposed here, will play a pivotal role in fostering trust, security, and privacy in digital interactions.

Spyware detection

In recent years, the proliferation of malware has reached unprecedented levels, leading to escalating cybercrime costs. Signatures extracted by static analysis of files have been widely adopted for malware detection: vendors maintain databases of known malware signatures that are shared with registered users. The recent literature has proposed the use of private and consortium (thus, permissioned) blockchains for spreading signatures among blockchain users. These approaches require controlled access to enhance trust and accountability but restrict the widespread sharing of up-to-date signatures because users must be registered. In this chapter, we present a novel technique that leverages a public blockchain to enable the massive dissemination of malware signatures among any users since a public blockchain is permissionless. On the other hand, the use of a public blockchain introduces new challenges related to security and data privacy, which our solution solves. The main benefit and outcome of our solution is that any users can securely access and verify malware signatures facilitating real-time detection of malicious files. We implemented our solution in Ethereum and exploited a smart contract written in Solidity to demonstrate that our approach is highly cost-effective.

4.1 Introduction

Malicious software (malware) is a major threat to computer security and has become increasingly sophisticated over the years [44]. Cybersecurity Ventures forecasts that cybercrime costs will grow by 15% per year and will reach 10.5 trillion dollars by 2025 [45]. Due to the relevant damage and cost caused by cyber crimes, the attention of the research community to malware detection is quickly growing [46], [47], [48].

Among the wide spectrum of malicious software, a particularly dangerous category is spyware, which is specifically designed to collect personal or confidential information from infected devices. The widespread diffusion of spyware directly threatens users' privacy, since stolen data can include credentials, personal docu-

ments, or financial information. Traditional anti-malware systems rely on centralized and vendor-controlled databases of signatures, which means that new spyware variants may remain undetected until each vendor updates its local repository. This delay not only increases users' exposure to privacy violations but also highlights the need for a transparent and privacy-preserving mechanism for rapidly disseminating verified signatures.

One of the most used approaches to detect malware is based on *malware signature* [49]. A signature-based malware detection system relies on a centralized database of known malware *signatures*, a set of bit strings. The detection involves scanning a file to see if the file matches these attributes. Malware signatures are found by vendors and can be publicly available or exclusive to the vendor.

To adequately protect computer systems, anti-malware solutions should detect the widest range of existing malware. For this reason, building a collection of malware signatures coming from various vendors can surely help to improve malware protection. An attempt in this direction is given by the service VirusTotal¹, where file checking is done with more than 40 antivirus solutions. However, the check is very slow because a request to each antivirus solution is done. Instead, having a common database of signatures would allow for enlarging signature sharing.

The problem we address in this study is making malware signatures publicly available to a wider number of users, yet guaranteeing that this set of signatures is not corrupted.

For this problem, recent literature has proposed the use of blockchain technology thanks to its advantages in terms of security, immutability, traceability, and transparency [50], [51], [52]. The peculiarity of blockchain is that it enables the creation of transactions based on a distributed consensus protocol among network participants without central authority. The decentralization introduced by blockchain mitigates single points of failure, enhancing the resilience and availability of the malware signature database.

In the literature, many of the approaches that use blockchain for improving malware detection (see for example [53], [54], [55], [56], [57], [58], [59], [60]) are based on *private* or *consortium* blockchains. A private blockchain operates under restrictions because only selected participants have permission to create and read blockchain data. Access is limited to a restricted set of users (typically belonging to a single organization) and the system is managed centrally, offering high control, privacy, and transaction rate. In a consortium blockchain, the management of the system is carried out by a group of organizations rather than a single entity, enabling a certain degree of decentralization. This decentralization reduces the transaction rate

¹ <https://www.virustotal.com>

and introduces some issues regarding access control and privacy. Again, a public blockchain is fully open, allowing anyone to join the blockchain, generate transactions, and view all data stored in the blockchain. This type of blockchain favors decentralization and transparency but strongly reduces access control, accountability, and privacy.

In the literature, approaches based on private or consortium blockchains have been preferred thanks to their scalability (because the number of nodes is controlled), low transaction costs (costs are predetermined), low energy consumption (because consensus protocols are less complex), more security (no risk of 51% attack [61]), no criminal activity (because users cannot be anonymous). Despite the significant advantages given by private and consortium blockchains, the need to register all users who need to access the blockchain strongly limits the pervasiveness of the solution. This mandatory registration restricts access to the blockchain, excluding the broader community from the potential benefits. Moreover, because each private or consortium blockchain typically operates independently, users must register to multiple blockchains to access all available malware signatures, fragmenting access and further complicating their experience [54]. For applications of malware signature sharing where open access and wide-reaching dissemination are crucial, this limitation undermines the system's scalability and its ability to foster timely.

To overcome this limitation, in this work, we designed a system for malware signature sharing that is based exclusively on a public blockchain, to have a more pervasive and reliable system of signature spreading. It is worth noting that while a public blockchain solves the accessibility limitations typical of private and consortium blockchains, its use introduces new concerns. Indeed, in public blockchains, any user can store information without restriction, which increases the risk that an attacker submits harmful data to compromise the integrity of the malware signature database. Moreover, since the access is not controlled, it becomes difficult to ensure the reliability of data sources, which is crucial for malware detection. Thus, designing an effective solution for malware signature spreading based on a public blockchain is challenging.

To the best of our knowledge, only one proposal based on public blockchain has been presented in the literature. Fuji Ryusei et al. [62] proposed a system based on a voting mechanism to determine whether a file is malicious or benign. When a user shares a file's signature on the blockchain, other users who encounter the same file can check the blockchain for existing votes (previous users' evaluations of the file) or submit their votes based on their detection systems' assessments. The system relies on a web server for user authentication and restricting voting access to approved users. This server introduces a critical vulnerability to the system, compromising

some of the core benefits of blockchain technology and creating a single point of failure. In contrast, our proposal presents a more advanced solution that leverages a public Ethereum blockchain to allow open access to malware signature verification, and a smart contract for access control, thus providing a fully decentralized solution. Moreover, the technique used in [62] to contrast Sybil attacks is weak because only the registration of consecutive addresses from the same IP address. These countermeasures may not be sufficient because IP addresses can be easily spoofed (e.g., using VPNs). In contrast, our solution provides a more robust countermeasure to Sybil attacks.

The main contributions of our research are summarized in the following (all these aspects are discussed in Section 4.2):

1. Differently from the state of the art in which private or consortium blockchains are used, the proposed system exploits a *public* (thus, permissionless) blockchain to store malware signature data across the network's nodes. The use of a public blockchain is a novelty our research introduces and allows everyone to check whether a file is malware, thus increasing the pervasiveness of the solution. On the other hand, the use of a public blockchain introduces new challenges related to security and data privacy that our solution solves.
2. We provide the implementation of our proposal on the public blockchain Ethereum to allow both reproducibility and comparison.
3. Our solution is shown to be cheaper than the state of the art (Section 4.4.1).

The rest of this chapter is organized as follows. Section 4.2 provides a review of the state of the art. Section 4.3 describes the proposed system in detail, including the signature storing and verification process. Section 4.4 presents the implementation of the proposal and its evaluation. The security analysis of our proposal is discussed in Section 4.5. Finally, Section 4.6 discusses some important aspects of the study and concludes the study.

4.2 Literature Review

Malware detection techniques aim to identify and mitigate the presence of malicious software, commonly known as malware, within computer systems and networks [63]. As malware continues to evolve in sophistication and complexity, a range of detection techniques have been developed to counter these threats. In the following, we provide an overview of the main types of malware detection techniques.

Signature-based detection is the most common approach, which involves comparing the digital signatures of files or code snippets against a database of known

malware signatures [64]. If a match is found, the file is flagged as malicious. This approach is effective against known malware but can be easily evaded by polymorphic or previously unseen malware [65], [66].

Dynamic analysis techniques involve executing suspicious files or code snippets within a controlled environment, such as a virtual machine, and monitoring their behavior in real-time [67], [68]. By observing the actions and interactions of malware during execution, dynamic analysis can reveal its intent and potential harm.

Behavior-based detection techniques focus on monitoring the behavior of programs or processes to identify suspicious or malicious activities [69], [70], [71]. These techniques establish a baseline of normal behavior and look for deviations or anomalies that indicate potential malware. Behavior-based detection can detect new and previously unseen malware but may produce false positives if legitimate software exhibits unusual behavior [72].

Heuristic-based detection involves using a set of predefined rules or heuristics to identify potentially malicious behavior [73], [74]. These rules are derived from the characteristics and patterns commonly exhibited by malware. Heuristics can include code analysis, system monitoring, and pattern matching. While heuristic-based detection can detect unknown malware, it may also generate false positives and false negatives due to the reliance on predetermined rules [49].

Anomaly-based detection techniques establish a baseline of normal system behavior and detect deviations that indicate potential malware [75], [76]. These techniques leverage statistical analysis, machine learning algorithms, or artificial intelligence to identify patterns that deviate significantly from the norm. Anomaly-based detection can detect unknown and zero-day malware but may produce false positives if benign but uncommon activities occur.

Machine learning techniques have gained popularity in malware detection due to their ability to analyze large amounts of data and identify complex patterns [77], [78]. Machine learning models can be trained on features extracted from known malware samples, allowing them to learn and detect similar patterns in new files or behaviors. This type of detection can be effective against both known and unknown malware, but it requires a robust training dataset and may be vulnerable to adversarial attacks.

Ensemble approaches combine multiple detection techniques to improve accuracy and reduce false positives. This can involve combining signature-based and behavior-based methods or integrating machine learning with traditional detection techniques. Ensemble techniques leverage the outputs of multiple detection systems to make a collective decision, enhancing overall detection performance.

From this overview, it emerges that each approach has its strengths and limitations. However, among all these approaches, the signature-based one is commonly employed alone or in conjunction with other approaches. Since the signature-based approach is widely adopted in current antimalware systems, in the following we will focus on this type of approach for malware detection and on the use of blockchain to support malware detection. Indeed, blockchain technology has been explored as a potential solution to the limitations of traditional malware detection systems. The decentralized nature of blockchain provides a more secure and reliable system for sharing and verifying malware signatures [79].

Homayoun Sajad et al. [53] propose a Blockchain-Based Malware Detection Framework (B2MDF) for detecting malicious mobile applications in mobile applications marketplaces (app stores). The framework consists of both an internal and an external private blockchain as well as a consortium blockchain. The internal private blockchain stores feature blocks extracted by both static and dynamic feature extractors, while the external blockchain stores detection results as blocks for current versions of applications. The blockchain provides the final decision by considering the results of several detection engines.

Martin George et al. [54] use a blockchain to control malware infections in mobile devices. Each node of the blockchain acts as an anti-malware agency, and a file is marked as malware by a majority voting procedure. Signatures of malicious programs are shared between all system's members.

A framework of two blockchains for Malware Detection and Evidence Extraction (CB-MDEE) in mobile devices is proposed in [55]. A blockchain is only accessed by a set of pre-selected users who can participate in the consensus process. Moreover, this blockchain stores the features of malicious codes to form a distributed malware fact-base. The second blockchain is open for any user who needs to provide detection and evidence services to join as a member node.

Golomb Tomer et al. [56] proposed *CIoTA*, a Collaborative IoT anomaly detection system via blockchain. The study demonstrates the use of the blockchain protocol to facilitate collaboration among devices of the same type, creating a trusted anomaly detection model. The system leverages public key infrastructure to identify and authenticate every device and uses also symmetric key encryption, ensuring that only devices with the correct shared secret can access and contribute to the blockchain.

Kumar Rajesh et al. [57] propose a framework that combines machine learning and blockchain technology to enhance malware detection for Android IoT devices. The solution incorporates both blockchain technology and machine learning techniques for detecting malware in IoT devices. Machine learning automatically extracts malware information using clustering and classification techniques, and

stores the information in the blockchain, which provides real-time information in a distributed malware database. This increases the speed and accuracy of malware detection by storing the extracted features in blocks to announce malware feature information to all users.

A system using Federated Learning for decentralized Android malware detection [58]. When an app is uploaded to the network, its version is stored in the InterPlanetary File System (IPFS), and a deep learning model extracts benign and malicious features. The updated model and its weights are stored in IPFS to reduce blockchain costs. The app's hash and model's decision are logged on a distributed ledger. Finally, during the app download process, a smart contract verifies the app's hash, notifying the user if the app is identified as malicious. The system exploits a smart contract to verify that apps uploaded to the network are either benign or malicious based on the aggregated models stored on the blockchain. The system also leverages the InterPlanetary File System to store app data, with the hashes of these files recorded on the blockchain.

Denysiuk Dmytro et al. [59] present an approach to malware detection by integrating blockchain technology with deep learning. They use a distributed network of computers that each analyze potentially dangerous code using machine learning algorithms. The results from the participants are stored and validated on a blockchain using a custom consensus mechanism named Proof-of-Action. Proof-of-Action uses a mechanism for validating results from a group of nodes that have completed the verification.

A signature-based system for malware detection is presented in [62]. When an executable file is downloaded by a user, the file signature is verified to check whether it is already registered as suspected malware on the blockchain. If the signature is not found and the antivirus system detects the file as malware, the file's signature is shared on the blockchain. In case the signature is already in the blockchain, the result of the local antivirus detection system (both malware and goodware) is sent to the blockchain. In this way, the blockchain indexes the files marked as suspected by at least one user and the number of times in which the file is both marked as malware and as goodware.

A blockchain-enabled security framework to defend against ransomware attacks in smart healthcare systems is presented in [60]. The framework leverages blockchain technology for secure detection and mitigation of ransomware. Although blockchain had not been widely applied in this context at the time, the framework's use of distributed ledger technology helps create and store digital signatures of ransomware attacks for future detection. Critical security alerts regarding potential vulnerabilities are stored on the blockchain, ensuring their immutability.

4.2.1 Different types of blockchains

Many of the techniques described above are based on the use of blockchain. The purpose of this section is to highlight the different types of blockchain and to remark on the challenges and advantages derived from the use of a public blockchain.

We have three types of blockchain: private, consortium, or public (see Section 2.2). In Private and Consortium Blockchain (e.g., Hyperledger), all nodes must be deployed and managed by the participants. The process for setting up nodes requires servers or cloud resources, along with a technical team for setup and support. When using cloud resources, costs can include fees for compute power, storage, and data transfer, all of which scale with network size. For example, on platforms like AWS or Azure, each node may need a virtual machine instance and dedicated storage, which increases monthly costs as storage demands grow. Costs include hardware, server space, and high availability considerations, all of which scale with the network size. Additionally, maintaining a network with multiple participants demands administrative coordination, configuration, and security protocols. Private and consortium blockchains have more stable and predictable costs, ideal for systems with a constant and controlled load, but they require initial infrastructure expenses. Every organization using a private or consortium blockchain must manage its own updates. Regular updates are mandatory not only for new features but also for critical security patches. Each participating node requires updates, with coordination among organizations. These updates demand significant planning and resources, increasing costs.

In private and consortium blockchains, network security is entirely dependent on the entity (or entities) running the nodes. This includes data encryption, secure access management, and protection against external attacks. With fewer nodes and more centralization, the network is often more vulnerable to attacks. Additionally, internal teams are required to manage access controls, set security policies, and monitor threats, which substantially increases costs.

In public blockchains, stored information is publicly accessible. The public blockchains are typically maintained by thousands of nodes distributed globally. Each node is operated by an individual, often miners are incentivized by transaction fees (gas) and block rewards. There are no costs associated with server maintenance, software upgrades, or security, thanks to the decentralized nature of the network, users pay exclusively for every transaction or smart contract executed. In public blockchains, both software and protocol upgrades are fully managed by the official community and development teams. This greatly simplifies management overhead and reduces associated costs. As a result, users do not have any responsibility for these matters. Users are not required to maintain servers, and they do not need to

Blockchain Type	Infrastructure Costs	Update Costs	Security Management Costs	Transaction Fees (Gas)
Public	Low	None for end-users	None for end-users	for each transaction
Non public	High	High	High	None

Table 4.1: Comparison between public and non-public blockchains.

handle the consensus protocol or maintain the security of the system. On public blockchains. The end-users do not face direct costs associated with infrastructure maintenance, hosting, or hardware. The blockchain's consensus mechanism (e.g., Proof of Work or Proof of Stake) secures the network, making it resistant to tampering, and end-users benefit without additional costs or intervention.

Table 4.1 summarizes the results of the comparison between using a public and a non-public blockchain. From the analysis of this table, we find that using a public blockchain gives advantages in terms of no need to set up, maintain, or manage the infrastructure. On the other hand, the small initial costs are compensated by the fees to pay at each transaction (this aspect is investigated in Section 4.4.1).

4.2.2 Comparative study

In this section, we compare our proposal and other similar techniques that utilize blockchain technology for malware detection. For each proposal, we report in Table 4.2 the following characteristics:

- Overview: a summary of the blockchain-based approach each study uses, highlighting the specific role of blockchain technology in the security solution.
- Case Study: the specific application or scenario where the solution is applied.
- Method: the type and number of blockchains implemented.
- Result: The key benefits or improvements achieved by each solution.

In the table, we report also the properties used to compare the various solutions, which are:

- Public: if the proposal uses exclusively public blockchains. This aspect is important as discussed in Section 4.2.1).
- Implementation (Impl.): if there has been proposed an implementation of the solution. This aspect is relevant to evaluate the cost and effectiveness of the solution.
- Pervasiveness (Pervas.): a solution is pervasive if any user can exploit it without the need to register in the system. This aspect is important because pervasiveness enlarges the audience of users of the service.

- Anti-Sybil: this characteristic evaluates the robustness of the solution when an attacker tries to subvert the system by creating a large number of pseudo-users who are used to gain a disproportionately large influence [80].

The four characteristics have a boolean value (i.e., they can be present or not).

The first considered technique ([53]) exploits private blockchains, thus an infrastructure should be built. Moreover, as the blockchain is private, only permitted users can exploit the solution.

The proposal presented in [54] exploits a consortium blockchain (thus, not a public one) and allows only users of the consortium to access malware signatures.

The solution proposed in [55] relies on a consortium blockchain for malware detection. Moreover, no implementation of the architecture is provided (only the algorithms for malware detection are implemented).

The solution proposed in [56] relies on a private blockchain to enable secure and trusted communication among a predefined set of IoT devices. Due to this, only known devices can join and participate in the network, which limits pervasiveness. However, *CloTA* is robust to Sybil attacks by combining the consensus protocol with strong encryption. Even if some nodes are compromised, the majority of honest nodes can reject malicious behavior, and the attacker would need to compromise a large number of devices to obtain their secret keys.

The multimodal malware detection technique for Android IoT devices, proposed by Kumar Rajesh et al. [57], uses a non-public blockchain, acting as a distributed malware database. Although the implemented framework provides strong experimental results, access is limited to authorized participants, and the solution is not pervasive, as it requires users to be part of a permissioned network.

The proposal presented in [58] relies on a private blockchain for its system. The paper does not explicitly address resilience to Sybil attacks. While private blockchains can be susceptible to such attacks if node participation is not properly controlled.

Denysiuk Dmytro et al. [59] implemented their solution by creating a network of 120 computers, divided into 12 subnets. The dataset used for the experiments is publicly available and contains 3214 samples of different malware classes. The use of the Proof-of-Action mechanism provides some level of protection against Sybil attacks.

The security framework proposed in [60] uses a private blockchain for healthcare data, limiting access to authorized users. The implemented solution also includes anti-Sybil protection via the Practical Byzantine Fault Tolerance (PBFT) algorithm and employs a private blockchain designed specifically for healthcare data, restricting access exclusively to authorized users, making it a non-pervasive solution.

The solution proposed by Fuji Ryusei et al. [62] relies on a centralized web server, which represents a critical aspect. To contrast Sybil attacks, the registration of consecutive addresses from the same IP address is blocked and the registration exploits CAPTCHA. These countermeasures may be not sufficient because IP addresses can be easily spoofed (e.g., using VPNs) and CAPTCHA systems could be bypassed [81], [82]. An additional limitation of this proposal is that any user can contribute to malware detection, thus allowing malicious colluding users to manipulate malware detection results. In contrast, in our proposal, only a *Certified Agency* can contribute to malware detection.

The improvement of our work with respect to the above proposals can be summarized as follows. Firstly, we use a public blockchain instead of private or consortium blockchains. This is important because reduces the implementation costs as no infrastructure has to be built (see Sections 4.2.1 and 4.4.1). Moreover, a public blockchain allows pervasiveness, since any user can verify whether a file is malware without the need to be registered to a private or consortium blockchain. In addition, we deployed our smart contract not only to facilitate the adding of signatures from certified agencies but also to establish a massively shared signature system where any user can interact with them for free to ensure that the file with which the user is working is not malware. Finally, the Anti-Sybil property is satisfied because each anti-malware vendor who wants to enroll in our system has to obtain both human trust and cryptographic trust during the registration process by interacting with the owner of the smart contract.

4.3 Proposed Work

In this section, we present the approach based on a public blockchain used to detect malware. The architecture of the solution, which is depicted in Figure 4.1, relies on the following stakeholders:

- a public *blockchain*. The function of the blockchain is to manage malware signatures. For this purpose, a suitable smart contract is deployed (an oracle is exploited by the smart contract);
- an *Owner*. This entity has the function of building and publishing the smart contract mentioned above;
- several *Agencies*. These entities publish on the blockchain the signature of malware detected by means of self-managed tools (we recall that how to extract malware signature is an orthogonal problem that is not the focus of our research). An Agency can become a *Certified Agency*.
- *Users*, who exploit the system to verify whether a file is malware.

Paper	Overview	Case Study	Method	Result	Public	Impl.	Pervas.	Anti-Sybil
[53]	A blockchain stores feature blocks, another stores detection results as blocks.	Malware detection in mobile app stores.	2 private blockchains, 1 consortium blockchain.	Distribution of detection results across multiple security providers in mobile app stores.	No	Yes	No	Yes
[54]	A blockchain for collective signature updates and majority voting.	Detecting malware on Android devices.	1 consortium blockchain.	Improved collaboration among anti-malware agencies.	No	Yes	No	Yes
[55]	A blockchain managed by distributed malware detection organizations, another blockchain that allows general users to request detection and evidence services.	Malware detection in Android mobile devices.	1 consortium blockchain, 1 public blockchain.	Enables secure and decentralized sharing of malware data among trusted nodes.	No	No	Yes	Yes
[56]	Lightweight framework that utilizes the blockchain concept to enhance the security of each device and the security of the network as a whole	Anomaly detection in IoT devices.	1 private blockchain.	Improving IoT network security and enabling scalability across large IoT networks.	No	Yes	No	Yes
[57]	A blockchain to store malware information.	Malware detection in Android devices.	1 private blockchain.	Improved ability to identify and respond to new malware.	No	Yes	No	Yes
[58]	A blockchain to store Android app features and a novel smart contract that enables aggregation process over blockchain platform.	Malware detection on Android devices.	1 private blockchain.	Security is improved by using blockchain to store Android app features and enable real-time malware detection.	No	Yes	No	No
[59]	A blockchain to store the results of analyses on potentially dangerous code.	Malware detection across computer networks.	1 private blockchain.	Improved scalability, enabling the integration of new malware detection methods.	No	Yes	No	Yes
[60]	A blockchain-based framework for ransomware detection in healthcare systems, providing secure storage for signatures.	Ransomware detection in smart healthcare.	2 private blockchains.	Enhanced protection against ransomware with secure data integrity.	No	No	Yes	Yes
[62]	Blockchain system for sharing malware signatures with decentralized access.	Malware signature sharing.	1 public blockchain.	Improved the sharing of malware signatures without the need for a centralized organization.	Yes	Yes	Yes	No
Our	Using Ethereum for sharing malware signatures	Malware signatures sharing and verification.	1 public blockchain.	Cost-effective, open, reliable verification for all users.	Yes	Yes	Yes	Yes

Table 4.2: Comparison with related work.

Now, we describe the steps done by these stakeholders in the considered solution.

A. *Blockchain Registration.* Since the solution is based on the use of the smart contract of a blockchain, the involved actors *Owner*, *Certified Agencies*, and *Users* must have the possibility to access the blockchain. Thus, the first step that these entities run is to generate a blockchain account. An account is composed of an address and a pair of private/public keys. Typically, the blockchain address is obtained from the public key, and the public key is obtained from the private key: in both cases, a cryptographic hash function is used for generation.




Fig. 4.1: Architecture of the solution.

- B. *Smart Contract Deployment*. A smart contract, whose definition will be presented in Section 4.4, is used to manage malware signatures. This smart contract is deployed by the *Owner*, which is the only entity authorized to modify an important parameter of the solution: the list of certified agencies. Such a list is initially empty and is enriched when an agency becomes certified (see next step).
- C. *Certification Application*. Each agency has some tools to detect malware. When an agency wants to share its knowledge with the platform, a certification request is generated. In particular, a certification request is a tuple $\langle B, u, e \rangle$, where B is the blockchain address of the agency, u is the URL of the website of the agency, e is the encryption of $\langle B, u \rangle$ done by the private key associated with the SSL certificate of the domain.
- D. *Certification Issuing*. Once the *Owner* receives a certification request from agency A , the following two checks are done, and both of them must succeed in order for the certification to be issued.
- In the first check, the *Owner* verifies that A is judged suitable for managing malware signatures. Observe that this check is critical because the *Owner* has to decide whether to grant the certification to the agency. This decision cannot be made automatically and requires the intervention of a human (this aspect is discussed in Section 4.6). If this check fails, the certification is not granted and the procedure stops.

- b) The second check is done by an *oracle*, a trusted Application Program Interface that provides the necessary data for the smart contract to complete off-chain operations [83]. The oracle verifies that the Certification Application comes from the correct company (this is done to contrast impersonation attacks – see Section 4.5). The *Owner* calls the smart contract to add A to the set of *Certified Agencies*. Now, the smart contract has to verify that the encryption e is valid, and an oracle is used to do this. By the oracle, the SSL certificate of the domain u is downloaded, and the corresponding public key is extracted. This key is used to decipher e and retrieve A and u . If the decryption is successful, then A becomes a *Certified Agency* and its blockchain address is saved.
- E. *Malware Addition*. Only a *Certified Agency CA* can mark a file as malware. When CA identifies a malware M in a file by means of its detection tools, this malware is reported to the blockchain. Specifically, the cryptographic hash digest d of the malware M is calculated and sent to the smart contract. The smart contract verifies if the agency is certified by checking that the blockchain address of the request sender is included in the list of the blockchain addresses of certified agencies. In this case, the smart contract stores that the digest d has been reported by the agency CA .
- F. *Malware Removal*. A signature can be removed only by the CA that included it previously. This is done by sending to the smart contract the cryptographic hash digest of the malware M . The smart contract first verifies that CA has previously included such a signature, and then removes the address of CA from the list associated with that digest.
- G. *File Checking*. Typically, this step is performed by a user to check whether a file is malware or not. This is done by calling the smart contract and sending the cryptographic hash digest d of the file to check. The smart contract searches d in the list of malware digests and returns the number of agencies that have reported such a digest as malware (if any). From the ratio between the number of agencies that detected the file as malware and the overall number of agencies, which is a number from zero to one, the user can estimate the probability that a file is malware.

4.4 Implementation

In this section, we present an example of the implementation of our solution used to validate the proposal and measure some parameters, such as the cost of the solu-




Fig. 4.2: Certification request.

tion. In our experiments, we exploited Ethereum because it is the most used public blockchain supporting smart contracts.

In the following, we describe how the phases of the solution run.

- A. *Blockchain Registration.* First, we created the blockchain accounts for the *Owner*, an *agency*, and a *user* by exploiting Metamask, a browser extension for enabling distributed applications on several blockchains, Ethereum included [84]. Our experiments have been carried out on *Sepolia* [85], a testnet that mimics the operating environment of Ethereum but runs on a separate ledger. This environment allowed us to develop our solution for free and in a risk-free way.
- B. *Smart Contract Deployment.* We developed the smart contract implementing our solution in Solidity [86]: Key excerpts of the code are presented in Figures 4.6–4.8, whereas the basic code is shown at <https://github.com/CyberGoldSun/SC>.

The smart contract exploits the hashtable `signatureMap`, where the key is the signature of the malware and the value is a struct called `FileInfo`. `FileInfo` contains a list of the addresses of the certified agencies that have added that signature. The mapping `validAgencyMap` associates the domain name of an agency with its address. The constructor method, invoked at the deployment of the smart contract, sets the `owner` to the address of the blockchain node that deploys the smart contract. The parameters for the oracle, which are the link token, target oracle, `jobId`, and `fee`, are also set.

- C. *Certification Application.* In this step, the agency aims to become a certified agency. For this purpose, the agency has to generate a digital signature of its domain name and blockchain address done with the private key of the SSL certificate associated with the domain name. Figure 4.2 shows the sequence diagram that describes this function.




Fig. 4.3: Software for signature generation.

As a proof of concept, we implemented a basic executable that allows the agency to do this by providing the path to the credential used by the web server (together with the domain name and blockchain address). Figure 4.3 shows a screenshot of the executable run, where the generated signature is reported.

D. *Certification Issuing.* Once the certification request is received and the agency is judged reliable, the *Owner* calls function `checkAgency` of the smart contract. This function receives a domain name u , a Blockchain address B , and the encryption e of $\langle B, u \rangle$. By the use of `require`, it is verified that only the owner is authorized to call this method.

A get request to the public verifier is sent and the result of the signature check is awaited. Once the process done by Oracle finishes, the flow goes back to `checkAgency` function: in the case of a valid signature, a new `ValidAgency` object is created and added to the mapping `validAgencyMap`.

E. *Malware Addition.* This procedure is implemented by means of function `addSignature`.

First, it is verified that the agency submitting the signature is certified. In this case, if the signature is already present in the mapping `signatureMap`, only the address of the agency is included in the list associated with this signature. Otherwise, a new key is added to the mapping `signatureMap`, a new struct `FileInfo` is created, and the address of the agency is included in the list.

F. *Malware Removal.* The removal of a signature is done by the agency that included it and is implemented by function `removeSignature`. Initially, it is verified that the agency who invoked the method is yet certified. In this case, the list of the addresses of agencies who submitted that signature is retrieved. An iteration is performed over this list to find the address of *Certified Agency* that called the function. Only if the address of the calling agency is found, the address itself is removed from the list. Otherwise, no removal is performed.

G. *File Checking*. This check is performed by function `isMalware`, which verifies if the digest is included in `signatureMap` and returns the number of agencies that marked the file as malware. On the basis of this number, the user decides whether the file can be considered trusted (a similar result is given by the service `VirusTotal`).

4.4.1 Cost and latency

The implementation of our solution allows us to quantify the costs of the solution and the time taken for each operation. We report these measures in Figure 4.4 and Figure 4.5. Specifically, the smart contract deployment is 4 milliEther (in March 2024, it is about 13.97\$), and less than 1 minute is the time taken to publish the smart contract. We observe that the time necessary to deploy the smart contract depends on various factors, such as network congestion, gas price, number of active nodes, and tools used. Although it can take several hours for the deployment of the smart contract, we highlight that this aspect has a low impact on the solution because smart contract deployment is a one-time task. Adding a certified agency (function `checkAgency`) costs 0.33 milliEther (about 1,15\$), and the call of this function takes about 18 seconds. The cost of adding a malware signature (function `addSignature`) is 0.12 milliEther (about 0,40\$), and it takes about 15 seconds. Removing a malware signature (function `removeSignature`) costs 0.065 milliEther (about 0,23\$). More importantly, checking whether a file is malware (function `isMalware`) is free since it only reads data from the blockchain and does not add or modify any data in the storage. Moreover, this call has less than one second of latency.

The costs measured to implement our solution are significantly lower than those of solutions built on private or consortium blockchains because:

- public blockchains run on decentralized networks that are already established and maintained by a wide range of participants. This means that there is no need to set up, maintain, or manage the infrastructure. In contrast, with private or consortium blockchains, there is some level of infrastructure setup required, such as managing nodes, security, and network governance.
- In public blockchains, the cost of maintaining the network (including network upgrades, and security patches) is typically borne by the decentralized community that supports the blockchain. In contrast, with private or consortium blockchains, even if the infrastructure is hosted on cloud platforms, the user has costs related to the continuing maintenance, updates, and administrative tasks necessary to keep the blockchain network functional and secure.




Fig. 4.4: Cost of each operation.




Fig. 4.5: Latency of each operation.

- Public blockchains use decentralized consensus mechanisms (such as proof-of-stake or proof-of-work), where validators (miners or stakers) are typically compensated by network incentives (e.g., transaction fees, block rewards). In private or consortium blockchains, consensus mechanisms are managed by the participants themselves, who bear the cost of running and maintaining validator nodes.

Other advantages of using a public blockchain compared to a private or consortium blockchain are:

- there is no need for upfront capital investments for hardware, software, physical space, or maintenance;
- no scalability problem because there is no need to size the system in advance;
- public blockchains offer high availability guarantees because nodes are spread geographically across multiple regions;
- for the same reason, public blockchains natively offer disaster recovery and backup properties.

In summary, using a public blockchain offers substantial advantages by eliminating infrastructure, maintenance, and security overhead, making it a more cost-effective option for most applications compared to private or consortium blockchain.

These reasoning and results show that the solution we propose is effective and the costs are sustainable with the profits derived from the provided service (see Section 4.6).

4.4.2 Smart Contract Source Code

A basic code of the smart contract implementing our proposal is reported below.

```

1  pragma solidity >=0.7.0 <0.9.0;
2  import "@chainlink/contracts/src/v0.8/ChainlinkClient.sol";
3  import "@chainlink/contracts/src/v0.8/ConfirmedOwner.sol";
4
5  contract System is ChainlinkClient {
6      address defaultAddress = 0x0000000000000000000000000000000000000000;
7      address ownerAddr;
8      using Chainlink for Chainlink.Request;
9      string public URL = "https://oracle/verify?domain=";
10     bytes32 private id;
11     uint256 private fee;
12     bool public oracle_result;
13     event OracleResult(bytes32 indexed oracleId, bool oracle_result);
14
15     struct FileInfo {
16         address[] agencies;
17     }
18
19     mapping(string => FileInfo) internal signaturesMap;
20
21     struct ValidAgencies {
22         address addressAgency;
23         bool isDeleted;
24     }
25
26     mapping(string => ValidAgencies) public validAgencyMap;
27
28     struct OracleInformation {
29         string domain_name;
30         address valid_address;
31     }
32
33     mapping(bytes32 => OracleInformation) private oracleRequest;

```

Fig. 4.6: Sketch of the smart contract (part 1).

```

33     constructor() {
34         ownerAddr = msg.sender;
35         setChainlinkToken(0x779877A7B0D9E8603169DdbD7836e478b4624789);
36         setChainlinkOracle(0x6090149792dAAeE9D1D568c9f9a6F6B46AA29eFD);
37         id = "c1c5e92880894eb6b27d3cae19670aa3";
38         fee = (1 * LINK_DIVISIBILITY) / 10;
39     }
40     function addSignature(string memory _hash, string memory domain_name) public payable {
41         require(validAgencyMap[domain_name].addressAgency == msg.sender, "You are not
           a valid agency");
42         for (uint i = 0; i < signaturesMap[_hash].agencies.length; i++) {
43             if (signaturesMap[_hash].agencies[i] == msg.sender) {
44                 revert("Signature already added");
45             }
46         }
47         signaturesMap[_hash].agencies.push(msg.sender);
48     }
49     function removeSignature(string memory _hash, string memory domain_name) public {
50         require(validAgencyMap[domain_name].addressAgency == msg.sender, "You are not
           a valid agency");
51         address[] storage agencies = signaturesMap[_hash].agencies;
52         for (uint i = 0; i < agencies.length; i++) {
53             if (agencies[i] == msg.sender) {
54                 agencies[i] = agencies[agencies.length - 1];
55                 agencies.pop();
56                 return;
57             }
58         }
59         revert("Signature not found");
60     }
61
62     function isMalware(string memory _hash) public view returns (uint) {
63         return signaturesMap[_hash].agencies.length;
64     }
65     function checkAgency(string memory domain_name, address valid_address, string memory
           api_url) public returns (bytes32) {
66         require(msg.sender == ownerAddr, "Not allowed");
67         Chainlink.Request memory req = buildChainlinkRequest(id, address(this), this.fulfill.
           selector);
68         req.add("get", api_url);
69         req.add("path", "r");
70         bytes32 oracleId = sendChainlinkRequest(req, fee);
71         OracleInformation memory oracleInfo = OracleInformation(domain_name, valid_address);
72         oracleRequest[oracleId] = oracleInfo;
73         return oracleId;
74     }

```

Fig. 4.7: Sketch of the smart contract (part 2).

```

62  function addAgency(string memory domain_name, address agency) private returns (bool) {
63      bool esito = false;
64      ValidAgencies memory v = ValidAgencies(agency,false);
65      validAgencyMap[domain_name] = v;
66      return esito;
67  }
68  function addressToString(address _address) internal pure returns(string memory) {
69      bytes32 _bytes = bytes32(uint256(uint160(_address)));
70      bytes memory HEX = "0123456789abcdef";
71      bytes memory _string = new bytes(42);
72      _string[0] = '0';
73      _string[1] = 'x';
74      for(uint i = 0; i < 20; i++) {
75          _string[2+i*2] = HEX[uint8(_bytes[i + 12] >> 4)];
76          _string[3+i*2] = HEX[uint8(_bytes[i + 12] & 0x0f)];
77      }
78      return string(_string);
79  }
80
81  function fulfill (bytes32 oracleRequestId, bool _oracle_result) public
82      recordChainlinkFulfillment(oracleRequestId) {
83      emit OracleResult(oracleRequestId, _oracle_result);
84      if (_oracle_result) {
85          OracleInformation memory oracleInfo = oracleRequest[oracleRequestId];
86          addAgency(oracleInfo.domain_name, oracleInfo.valid_address);
87      }
88      oracle_result = _oracle_result;
89  }
90  function concatenate(string memory domain_name, address agency, string memory
91      encryption_) public view returns (string memory){
92      return string(abi.encodePacked(URL, domain_name, "&address=", addressToString(
93          agency), "&encryption=", encryption_));
94  }

```

Fig. 4.8: Sketch of the smart contract (part 3).

4.5 Security Analysis

In this section, we analyze the security properties of our solution to show that an attacker is not able to compromise the result of a malware detection process. Especially because our solution is built on a permissionless public blockchain, the security analysis is essential for evaluating how the proposed system defends against malicious attempts to manipulate the decision about the presence of malware in a file. Our analysis expects that the reader is familiar with the security properties of cryptographic primitives so, just for an example, we will not prove that a signature guarantees authenticity and integrity.

In our analysis, we focus on evaluating how our solution addresses the specific and relevant threats within the scope of malware signature sharing. General blockchain vulnerabilities, such as smart contracts and Oracle risks, will be neglected in our analysis because they are independent of the proposed system's design. Consider that general blockchain vulnerabilities have been extensively studied and known risks have been mitigated by platform developers. Thus, by adhering to standard best practices, it is possible to minimize these security risks. Moreover, our solution exploits Ethereum and Oracle providers, which are widely adopted infrastructures that have already implemented extensive security measures to minimize their respective risks. All these reasons suggest we concentrate the analysis on assessing threats directly related to system functionality rather than foundational blockchain security.

We start with the threat model.

4.5.1 Threat Model

In our threat model, we assume that the attacker can be either a generic user or a certified agency.

The security assumptions we adopt are standard in this context and include:

- **A1:** the attacker cannot break cryptographic primitives (e.g., reverse a cryptographic hash function or decrypt a message);
- **A2:** the attacker cannot alter blocks on the blockchain after they are mined;
- **A3:** the attacker cannot intercept and manipulate transactions sent to the smart contract.
- **A4:** the majority of certified agencies will act honestly and will not be subject to attack.
- **A5:** the *Owner* will act honestly and will not be subject to attack.

Note that some assumptions are not directly exploited in our analysis but are widely accepted in the context of blockchain (e.g., **A2** and **A3**). Some other assumptions (e.g., **A4** and **A5**) are included to avoid the analysis of trivial attacks (e.g., compromising most of the agencies or the *Owner*).

4.5.2 Security Properties

In this section, we present the specific goal of security that our solution must guarantee. The security property that should be guaranteed is that when a user submits a file's signature for verification, the system must confirm whether any certified agency has flagged the corresponding signature as malicious and return this result reliably.

4.5.3 Attacks

Our security analysis identified the following attack scenarios and how our proposal contrasts them.

- **Sibil Attack.** The first attack we consider is the Sybil attack, where a malicious actor creates multiple fake agencies to corrupt the system and manipulate the malware verification procedure, to falsely classify files as safe or malware.
- **Oracle attack.** Oracles play a critical role in verifying agency credentials during the certification process. An oracle manipulation could allow a fake agency to become certified, leading to incorrect or malicious signatures being added to the blockchain.
- **Integrity attack.** The integrity of malware signature reporting is crucial. A certified agency could attempt to mark a benign file as malware, or vice versa, by reporting incorrect signatures.
- **Byzantine Failures.** A set of malicious certified agencies may collude to submit false malware signatures or manipulate detection results, exploiting the voting mechanism to compromise the correctness of the system.
- **51% attack.** In this attack, a malicious actor controls a majority of the blockchain miners and can manipulate the consensus algorithm. Note that, even though this attack derives from the consensus algorithm used in Ethereum, we consider it because our solution is based on Ethereum (other solutions based on private blockchain are not vulnerable to this attack).

4.5.4 Mitigations

In this section, we discuss how our proposal contrasts the considered attacks. Note that we provide a security analysis tailored to the considered context of malware signature sharing. Since the security of the proposed system relies on widely accepted cryptographic primitives and the inherent characteristics of the blockchain (such as immutability, accountability, decentralization, and traceability), we avoid introducing additional formalization for these aspects.

- To contrast Sybil attacks, the proposed system includes a verification procedure before an agency is authorized to participate in signature management. The verification process requires the agency to provide legitimate and verifiable information during registration, which includes validating the web domain of the agency, passing SSL certificate validation, and providing cryptographic proof of identity. To compromise this verification, the attacker should violate assumption **A1**. Moreover, the use of blockchain oracles reduces the risk of corrupting the verification process (see the Certification Request step).

- Oracles attacks are mitigated by using decentralized nodes that aggregate data from multiple sources, ensuring that no single point of failure exists. The number of nodes used by an oracle to get data defines the trust model [87]. Note that the use of oracles in the smart contract to verify the signature of the certification request could be replaced by a check done by an application running on a device of the *Owner*. However, this possibility introduces a vulnerability because an adversary could attack such an application and certify compromised agencies, thus opening the doors to a Sybil attack. In contrast, the use of oracles makes the attack on the signature verification done in the certification process more difficult to implement.
- Integrity attacks are mitigated because the decentralized decision-making ensures that a single compromised agency cannot falsely classify files without collusion from other agencies. Indeed, the decision about whether a file is malware or not is made by considering how many certified agencies mark the file as malware (through a majority vote). Moreover, when an agency submits malware signatures, these submissions are recorded on the blockchain, providing full transparency and enabling continuous auditing. In case of need, these logs can be used to detect which agency acted maliciously.
- The resilience of the proposed system to Byzantine failures depends on the structure of the voting mechanism employed for malware detection. In the current design, the user receives a report indicating how many certified agencies have flagged a given file as malware. Under a Byzantine attack, a subset of these agencies may be compromised and return falsified results, causing the aggregate output to diverge from ground truth proportionally to the number of corrupted participants. However, this threat is substantially mitigated by Assumption A4, which guarantees that a strict majority of participating agencies behave correctly at any given time. Under this assumption, the final detection outcome converges to the correct result because honest agencies always outnumber adversarial ones.
- Concerning 51% attacks, the decentralized and global nature of Ethereum makes a successful 51% attack highly unlikely because carrying out such an attack on Ethereum would require immense computational power. This decentralization is further reinforced by Ethereum's transition to the proof-of-stake consensus mechanism, where attackers would need to control a majority of staked Ether—a financially prohibitive requirement. However, even in the highly unlikely event of a 51% attack, the security properties of our solution remain guaranteed because the smart contract logic, including the storage and retrieval of malware signatures, is deterministic and transparent. Manipulating the outcome of a malware verification request would require altering historical transactions, which

Threat	Potential Impact	Mitigation
Sybil Attack	Malicious actors could create multiple fake agencies to manipulate the system, falsely classifying files as safe or malicious.	A centralized certification process managed by the smart contract owner ensures only trusted agencies are added, using cryptographic proof and SSL validation to prevent fraudulent entities.
Integrity Attack	Risk of storing incorrect or fake malware signatures, leading to false positives or negatives in malware detection.	All signature submissions are logged transparently on the blockchain, ensuring traceability and enabling audits to identify and address malicious or erroneous entries.
Byzantine Failures	Malicious agencies collude to manipulate detection results through false signature submissions, compromising the correctness of the voting mechanism.	Voting over agency outputs, guaranteed correct under Assumption A4.
Oracle Attack	Exploitation of external trusted sources (oracles) could lead to fake agency certifications or incorrect off-chain data being used in smart contracts.	Decentralized oracles aggregate data from multiple sources, minimizing the risk of manipulation or single points of failure in off-chain data verification.
51% Attack	A malicious entity gaining majority control of blockchain validators could manipulate stored data or transactions.	The system relies on Ethereum's decentralized network, which has strong security protocols and makes a 51% attack highly unlikely due to its global distribution and computational requirements.

Table 4.3: Key Threats and Corresponding Mitigation.

becomes increasingly difficult as time passes thanks to Ethereum's consensus mechanism.

A summary of threats and mitigations is presented in Table 4.3.

In conclusion, our solution exploits the security features of blockchain technology and implements safe procedures to offer a secure and reliable framework for decentralized malware signature sharing, ensuring that any file can be verified in a trusted and transparent manner.

4.6 Discussion and Conclusion

The proposed decentralized signature-based system for malware detection leverages blockchain technology to create a more secure and reliable system for sharing and verifying malware signatures. The decentralized nature of our system provides several security benefits over traditional signature-based systems. Since the signature data is distributed across a network of nodes, there is no single point of failure or at-

tack. The decentralized nature of the system also makes it more resilient to attacks, as a malicious actor would need to compromise a significant portion of the network to successfully manipulate the signature data.

However, while signature verification and sharing are fully decentralized procedures, the process of agency certification introduces a degree of centralization into the system because the certification is managed by a single entity. Indeed, the smart contract owner controls which agencies are granted certification, which may appear a violation of the decentralized nature of the system. We observe that the choice of certifying an agency or not cannot be automated but requires human reasoning (the choice is based on the trust in this agency). This centralization is thus necessary to prevent fraudulent agencies from joining the network. However, once certified, the system operates in a fully decentralized manner. Importantly, the presence of the smart contract owner does not interfere with the management and spread of malware signatures, which remains fully decentralized. The *Owner's* role is to contrast Sybil attacks, as observed previously.

The solution is designed to resist typical attacks like signature forgery and Sybil attacks by relying on the combination of cryptographic primitives and human vetting for agency certification.

The use of a public blockchain, as opposed to private or consortium alternatives, gives the advantage of making malware signature sharing universally accessible thanks to the cancellation of the registration barrier. This solution enables a sort of democratic malware detection because individuals and small organizations can access the same malware detection tool as larger entities. It is worth noting that no privacy issue is introduced since stored data do not include any sensitive information.

Concerning the cost of the system implementation, we remark that malware checking does not require any fee. The main cost (if we consider the frequency of the operations) is the addition of new malware signatures. The smart contract could include an additional fee for function calls to support the service/system: quantifying this fee or providing a business plan is out of the scope of this study. The analysis of the cost-effectiveness of our approach highlighted the cost savings achieved by avoiding the infrastructure, maintenance, and security overhead associated with private or consortium blockchain-based solutions. From this point of view, scalability and cost efficiency make our approach particularly appealing for its application to resource-constrained environments.

An important aspect to consider concerns scalability. Indeed, public blockchains may present scalability issues when the number of transactions and nodes increases. Scalability is a challenge, particularly during periods of high demand, causing net-

work congestion and an increase in transaction fees. However, the industry has made significant progress in addressing these concerns. To increase scalability, several approaches can be adopted. A possibility is to utilize decentralized mechanisms, such as layer-2 solutions, where the burden of processing and verifying transactions is distributed across a global network of validators. Innovations like zero-knowledge proofs and rollup technologies allow public blockchain to process thousands of transactions per second while maintaining security and decentralization [41]. For example, Ethereum 2.0 introduced a second-layer solution for reducing scalability issues. The main goal of this layer is to increase transaction throughput without sacrificing decentralization or security. Specifically, Ethereum 2.0 uses a scaling solution named Rollups [41]. The Rollups bundle hundreds of transactions into a single transaction on layer 1. This distributes the first-layer transaction fees across everyone in the rollup, making it cheaper for each user. In addition, Ethereum 2.0 uses proof-of-stake as a consensus mechanism. Layer 2 scaling solution and the consensus mechanism upgrades have significantly enhanced the scalability bringing over 100,000 transactions per second [88]. In applications in which this transaction rate is not sufficient, it is possible to replace Ethereum with more scalable blockchains, like The Open Network (TON) [43] and Solana [42]. These blockchains have redefined scalability on public networks through advanced consensus protocols and parallelized transaction processing. TON leverages its sharding technology to distribute workloads dynamically, while Solana's proof-of-history enhances throughput by synchronizing the network efficiently. These innovations demonstrate that scalability is achievable without compromising decentralization.

Another relevant aspect to consider regards the role of the *Owner* as the certifying entity, which introduces a degree of centralization and may pose risks to the overall trustworthiness of the system. In our approach, we assumed that the *Owner* is trusted and acts honestly. However, it is interesting to analyze the consequences when the *Owner* is compromised. The *Owner* is responsible for deploying the smart contract and certifying agencies. Thus, a first risk is that the *Owner* might modify the smart contract in such a way that it acts differently from the original intent, for example, to change the result of a file check.

A mitigation for this issue is given by the transparency provided by the public blockchain. Indeed, the smart contract code is publicly accessible and any changes to the code can be easily verified by the community. Security experts and certified agencies can check the contract for unauthorized modifications, ensuring that any deviation from the original code is promptly detected. Moreover, since key actions like malware signature submissions are logged and publicly accessible, stakeholders can verify that the smart contract behaves as expected.

The second risk is that the *Owner* might certify malicious agencies. Observe that our solution is designed in such a way that a malicious agency cannot alter the result of a file check when the file is malware (because agencies can report only the signature of malware). Instead, malicious agencies can create *false positives*, marking legitimate files as malware. Our solution can mitigate this risk since the result of a file check is the number of agencies that detected that file as malware. This would make it difficult for a single malicious agency to create false positives without collusion. In addition, audits could be implemented to evaluate the accuracy and consistency of each agency's reported signatures to detect an outlier behavior of an agency. However, this problem becomes relevant in the case more malicious agencies concur for generating false positives. To mitigate this issue, we could modify the results of the file check providing not only the number of agencies that detected that file as malware but also which are these agencies. This would allow a user to use the credibility of these agencies to evaluate the result of the file check.

Concerning long-term viability, our solution adapts to evolving malware threats and addresses the ongoing maintenance and update requirements typical in malware detection systems. Our solution supports the removal of malware signatures, making it highly adaptable to the dynamic nature of malware threats. For example, it can be useful in the case that, after a pattern is identified as malware, this detection is removed in the presence of false positive alerts. The ability to remove outdated or misclassified signatures is essential for long-term viability. In our system, only the certified agency that originally submitted a signature can remove it, ensuring accountability and traceability in the signature management process. This approach minimizes the risk of false positives while maintaining the integrity and accuracy of the database. Together, the possibility to add and remove signatures enhances the system's flexibility over time.

The use of smart contracts in our solution provides further support for long-term viability. Indeed, their inherent flexibility allows the system to adapt as malware detection requirements change. For example, when necessary, smart contracts can be built to incorporate additional metadata fields for richer contextual information, such as malware type, date of discovery, or severity level. New metadata fields enhance the quality and relevance of the stored malware signature and enable more informed decisions by users who can assess the relevance of a particular signature based on the specifics of the threat.

For these reasons, the proposed system appears an effective solution to provide a pervasive service for sharing and verifying malware signatures.

In future work, we include the integration of our system with real-time threat intelligence feeds to enhance the timeliness of malware signatures. Indeed, threat

intelligence feeds provide continuous updates on newly discovered malware, attack patterns, and vulnerabilities, enabling rapid adaptation to the evolving threat landscape. Threat intelligence feeds should be certified and would act as agencies by using suitable APIs to interface with the blockchain, in such a way that the system receives the new signatures directly from trusted sources. At the same time, real-time feeds could also support the removal of outdated or reclassified signatures, identifying entries that may no longer represent active threats.

Privacy-preserving Lawful Interception

Lawful Interception refers to the acquisition of the contents of communications between private individuals or organizations by subjects authorized by law. It involves three actors: the network operator (NO), the Law Enforcement Agency (LEA), and the Law Enforcement Monitoring Facility (LEMF). In the literature, standards and scientific solutions are proposed for the interception procedure and the interaction between the NO and the LEMF. However, no standard has been proposed for the interaction between the LEMF and the LEA. The absence of standards for controlling LEA (or a delegated agency) access to intercepted contents stored by the LEMF is a significant gap that should be overcome. This prevents the implementation of secure, interoperable, and automated procedures, leading to inefficiencies and security risks. In this chapter, we propose to cover the above gap by adopting the Self-Sovereign Identity (SSI) paradigm. The adopted research methodology follows a multi-phase approach that includes studying existing solutions, system design, and technical feasibility testing. The study first examines existing standards and identity management frameworks and their limitations. Next, an SSI-based architecture is proposed to manage the interactions between LEA (or a delegated agency) and LEMF. Finally, a proof of concept of the proposed solution written in Python and using the Hyperledger Indy blockchain has been implemented to assess whether our proposal is technically feasible. The proposed solution enhances automation, security, and interoperability in lawful interception. Indeed, it enables machine-readable authorizations, reducing errors and improving efficiency by eliminating manual operations. Additionally, verifiable credentials and decentralized identifiers strengthen security and standardize interactions across jurisdictions, ensuring privacy-preserving identity management. By standardizing interactions between LEA and LEMF, this research contributes to a more secure, privacy-preserving, and legally compliant lawful interception process.

5.1 Introduction

Lawful interception (LI) is a procedure in which the content of communications among private individuals or organizations is collected and provided to agents authorized by the law [89]. Besides ethical and legal issues [90], [91], which are outside the scope of this study, researchers focused their attention on the technical aspects of the LI procedure.

An LI procedure involves three actors: the Network Operator (NO), which physically intercepts data; the Law Enforcement Monitoring Facility (LEMF), which collects, stores, and analyzes these data; the Law Enforcement Agency (LEA), which accesses the data kept by the LEMF for legal purposes (see 2.4).

The literature includes several proposals and standards regarding both the physical interception of data performed by the NO [92], [93], [94] and the provision of these data to the LEMF, such as CALEA [13], 3GPP TS 33.108 [95], ITU-T Y.2770 [96]. Among them, ETSI TS 101 671 V3.15.1, published in June 2018, is the most recent and used technical standard in LI [97]. It has been developed by the European Telecommunications Standards Institute and provides guidelines for implementing the interfaces and processes required for LI, ensuring compliance with legal and regulatory frameworks while maintaining the integrity and privacy of intercepted data. Three Handover interfaces are defined to govern communication between the NO and the LEMF (see 2.4.1). Moreover, data types specifying the handling of intercepted information as well as mechanisms to ensure the confidentiality and integrity of intercepted data are defined. ETSI TS 101 671 serves as a foundational framework for lawful interception across jurisdictions, promoting interoperability and standardization while addressing privacy and ethical considerations.

However, to the best of our knowledge, the interaction between the LEMF and the LEA is not explicitly addressed in ETSI TS 101 671 or other standards. This leads to the adoption of different practices across countries [98], [99], mostly based on manual, non-secure processes, which can cause security and efficiency concerns, including increased risks of errors, delays, and data integrity violations. This issue becomes even more critical in cross-border investigations, where the LEA may delegate access to intercepted data to an External Agency (EA). Without standardized authentication and authorization mechanisms, an EA can retain access beyond the authorized period or an insider within the agency could extract sensitive information without detection. These vulnerabilities expose lawful interception systems to potential data leaks, unauthorized disclosures, and operational inefficiencies.

In this work, we aim to cover the above gap by adopting the Self-Sovereign Identity (SSI) paradigm [100], [101], [102], [103]. SSI is a new paradigm that empowers

people to own and control their data without relying on centralized entities. Further background is provided in 2.3. SSI provides several advantages over alternative approaches such as federated identity management (FIM) and authentication based on Public Key Infrastructure (PKI). Unlike FIM, which relies on trusted third parties (e.g., Google or Facebook) to authenticate users, or PKI-based solutions, which require certificate authorities to manage certificates, SSI is a decentralized approach that eliminates bottlenecks and single points of failure. The privacy-preserving nature of SSI allows users to selectively disclose only the necessary information without revealing their full identity. This is in accordance with the minimization principle of GDPR. In contrast, FIM and PKI solutions require continuous involvement of identity providers, and privacy concerns may arise from honest but curious providers that legitimately participate in the users' authentications but also have the ability to infer sensitive information from the observed user behaviors. Another advantage of SSI is its interoperability and portability. PKI requires multiple certificates for different services, and FIM systems need pre-established agreements between providers, limiting their portability. In contrast, verifiable credentials are used in SSI that can be interoperable across different organizations and jurisdictions (even if this introduces trust problems in cross-border authentication that should be suitably addressed in each specific application context). Overall, the numerous advantages of SSI make it the preferred solution for the future of digital identity [104].

For these reasons, we propose a solution relying on the state-of-the-art SSI Sovrin framework [105]. A core element of our proposal is the access control mechanism built on top of the SSI substrate. This aspect is very relevant, especially in the case of the LEA relying on an external agency to collect such data. In this case, proper authorizations should be provided to prevent unauthorized access. Although some guidelines are provided in this direction (such as the adoption of RBAC [106] or VACM [107]), no concrete proposal is currently available.

Another relevant aspect of our contribution is that the proposed solution applies to real-life contexts. Indeed, it is compliant with the Sovrin SSI stack [108], [109] and its evolution, the ToIP stack [110]. The above framework represents the state of the art in the SSI environment [111]. Starting from this framework, we built a proof of concept of the solution (publicly available on GitHub) using standard and well-consolidated technologies. We used this proof of concept to verify that the designed solution is technically possible. However, the used code has been developed quickly and lacks scalability, security, and integration with existing infrastructure, thus making it unsuitable for real-world experimentation.

The proposed solution provides three key benefits related to automation, security, and interoperability in lawful interception procedures, which are discussed be-

low. Unlike traditional lawful interception, which is usually handled by manual operations, our solution enables machine-readable authorizations, allowing access control, credential issuance, and verification to be fully automated. In this way, the absence of human intervention minimizes errors and improves efficiency. The use of verifiable credentials enhances security by eliminating the need for centralized identity providers, which can be hacked or manipulated, reducing single points of failure. Moreover, selective disclosure is supported, allowing users to share only necessary information and thus reducing privacy concerns. The lack of standardized interaction protocols between LEA and LEMF has led to fragmented implementations across different countries. Our solution addresses this problem using decentralized identifiers, Verifiable Credentials that follow the W3C format, and the DIDComm protocol, which ensures that different entities involved in the LI procedure, even across various jurisdictions and countries, can seamlessly interact.

This chapter is structured as follows. In Section 5.2, we discuss the faced problem and highlight the contribution of the proposed solution. The proposed approach is described in Section 5.3. In Section 5.4, we design the SSI stack underlying our approach and, in Section 5.5, we provide a proof of concept. The related literature is contextualized in Section 5.7. Finally, in Section 5.8, we draw our conclusion.

5.2 Motivations and goals

The literature includes several proposals for standardizing the interceptions between the NO and the LEA, as well as between the NO and the LEMF, such as CALEA [13], 3GPP TS 33.108 [95], ITU-T Y.2770 [96], and ETSI TS 101 671 V3.15.1 [97]. Among them, the latter is the most recent and used one in LI. However, to the best of our knowledge, the interaction between the LEMF and the LEA is not explicitly addressed in ETSI TS 101 671 or other standards, and the lack of a structured framework for LEA-LEFM interactions introduces three main issues:

1. *Interoperability issues*: The absence of interoperability standards led each country to adopt and implement its own national procedure to manage the interaction regarding how LEA should retrieve data from LEMF [99]. These methods are often incompatible with each other. This limitation becomes particularly critical in cross-border investigations, where an LEA may need to delegate interception access to an external agency. For example, consider a scenario where an EU member state is investigating a transnational criminal group. The national LEA may require assistance in accessing intercepted communications related to the group's activities. In this case, the LEA can request support from an external agency, such as the NSA, which is the U.S. government agency responsible for

foreign and domestic intelligence and counterintelligence purposes. This agency can facilitate access to the necessary data from LEMFs in other states. In this case, proper authorizations should be granted by the LEA to the external agency ensuring that the data are handled in accordance with legal requirements.

2. *Procedural inefficiencies*: Many lawful interception processes continue to rely on manual and insecure procedures, such as paper-based approvals and email exchanges. These mechanisms could result in operational inefficiencies and human errors and may introduce delays in investigations in which rapid access to the intercepted information is crucial.
3. *Security risks*: Security risks may arise due to the absence of standardized access control mechanisms. Without a protocol for managing and verifying access to intercepted data, there is no robust way to ensure that only authorized entities retrieve sensitive information. This exposes LI systems to several vulnerabilities, including unauthorized data retention, insider threats, and data leaks. Additionally, the lack of real-time access revocation mechanisms means that once an external agency is granted access to LI data, it may retain that access indefinitely, even after the investigation has concluded, increasing the risk of unauthorized disclosures and violations.

To address these challenges, we propose a decentralized and verifiable credential system based on the SSI paradigm. Our solution ensures that only authenticated and authorized entities can retrieve intercepted data while enabling real-time revocation of access rights to prevent misuse. Our solution provides the following outcomes:

- *Fine-tuning*: The agency obtains a credential for access to data specific to a certain request, including a validity period to verify whether the authorization has expired;
- *Minimality*: Data accessed by the agency is selectively disclosed, supporting the principle of minimal information disclosure;
- *Revocation*: Authorizations can be revoked by the issuer at any time, a significant advantage over standard signature-based solutions where only certificates can be revoked, not specific signatures;
- *Automation*: Authorizations are machine-readable, enabling fully automated authorization processes;
- *Effectiveness*: We provide a proof of concept based on established SSI standards, such as the Sovrin framework.

To better understand the points above, we anticipate some details of the solution.

The LEA can play the role of the issuer and release some credentials (i.e., authorizations) to the agency. On the other hand, the LEMF can play the role of the verifier

and check the credentials presented by the agency. However, the scenario presented above represents a simplification of a real-life LI procedure. Indeed, when an employee of the agency asks the LEA for the credentials, the LEA has to be sure that such an employee has the right to obtain such credentials. Again, the SSI approach can solve the problem. Indeed, the right to obtain the credentials from the LEA can be granted if the employee provides the LEA with other credentials released by the agency. These new credentials attest that the employee is authorized by the agency to participate in this LI procedure. In other words, in this case, the roles of the parties are different in contrast with the previous case (the agency represents the issuer and the LEA represents the verifier).

In our solution, we pursue the following goals:

- To allow the LEA to start a new LI procedure by commissioning an entity for the role of LEMF. In this phase, the LEA also interacts with the NO according to the standard [97].
- To allow an employee of an external agency to obtain from the LEA the credentials to access an interception procedure already started on the LEMF.
- To allow the LEA to revoke the credentials released to the employee and the LEMF.

The design of the proposed solution is the objective of Section 5.3.

5.3 Proposed approach

The scenario considered in this study involves five actors. The first three actors have already been presented in Section 2.4 and are the LEA, the LEMF, and the NO. The other two actors are the external agency (EA), which the LEA relies on to access the interception data, and an employee (E) of the EA.

In the following subsections, we define the three phases defined in the proposed solution. In Section 5.3.1, we present the setup phase, which involves LEA, LEMF, and NO. Although this phase is discussed in [97], some details are neglected, especially with regard to the communication between the LEA and the LEMF. This phase can take advantage of the integration of the SSI approach. Once the LI procedure is started on the LEMF, the LEA charges the EA to access the interception procedure on the LEMF. This task is performed by the employer E properly authorized. This phase is described in Section 5.3.2. Finally, we present in Section 5.3.3 the mechanism for revocation, which is run when the interception procedure is complete or when the access needs to be restricted under any other circumstances.




Fig. 5.1: Sequence diagram of the Setup in the LI procedure.

5.3.1 Setup

Through this phase, the LEA starts a new LI procedure by leveraging the SSI framework. This phase is performed through the following steps:

1. First, the LEA contacts the LEMF to establish a secure peer-to-peer channel according to the DIDComm protocol (see Section 2.3).
2. In this connection, the LEMF provides the LEA with all the technical information needed for the interception procedure (for example, the phone number and the IMEI of the SIM of the intercepted subject).
3. The LEA releases the credential VC_{LEMF} to the LEMF certifying that the latter is in charge of this LI procedure. This credential will be presented by the LEMF in all the next interactions with LEA.
4. Finally, an interaction according to the standard [97] is performed between LEA and the NO, so that the latter can define the interfaces to be used in the interception process.

The above steps are summarized in the sequence diagram of Figure 5.1.

5.3.2 External agency integration

Consider the employee E of the EA authorized by the LEA to access the data related to the LI procedure on the LEMF. We assume that E is preliminarily equipped with a credential VC_E , certifying their identity and role within the EA.

To allow E to access the data on the LEMF, the following steps are performed.




Fig. 5.2: Sequence diagram of the EA integration in the LI procedure.

1. Through a secure peer-to-peer channel, E sends the credential VC_E to the EA.
2. The EA verifies VC_E (as described in Section 2.3) and sends to E a new credential VC_L certifying E is authorized by EA to interview in the LI procedure.
3. Through a secure peer-to-peer channel, E sends VC_L to the LEA.
4. The LEA verifies VC_L and grants a new credential VC_{LM} to E. This new credential is used by E to obtain access to the LEMF.
5. Now, through a secure peer-to-peer channel, E sends the credential VC_{LM} to the LEMF.
6. The LEMF verifies this credential (*Check VC_{LM}*) to ensure that it is valid (not expired or revoked), issued by the LEA (checking the digital signature of the credential), and its claims satisfy the access control policy. Whether all these checks are successful, the LEMF grants E access to the interception data by a procedure they agree on (typically, data are sent via a secure channel or a storage device).

The above steps are summarized by the sequence diagram in Figure 5.2.

5.3.3 Credential revocation

The duration of an interception procedure is influenced by several factors, mainly the complexity of the case, and predicting how long the procedure should take is

not easy. When an expiration time for the verifiable credentials cannot be established, credential revocation should be used. The revocation process occurs after the LI procedure is complete or when the access needs to be restricted under any other circumstances to prevent unauthorized access to entities previously authorized.

Revocation can be triggered by both the LEA and the EA. In the first case, the LEA may revoke both the credential VC_{LEMF} and the credential VC_{LM} to remove the participation of the LEMF and E in the LI procedure, respectively. In the second case, the EA can revoke (i) the credential VC_L or (ii) the credential VC_E . In case (i), the EA removes the participation of E in the LI procedure. Case (ii) happens when E is fired, resigns, or changes role in the agency EA.

The LEA and EA use the same procedure for revoking credentials, which is based on cryptographic accumulators [112]. A cryptographic accumulator is a data structure that allows the compact representation of a set of elements, enabling efficient membership verification. Given a set of elements, the accumulator produces a concise value that encapsulates all the elements. This structure provides two properties:

1. Membership Proofs: It can be demonstrated that a specific element belongs to the set without revealing other elements of the set.
2. Dynamic Updates: Elements can be efficiently added to or removed from the set.

Each revoked credential becomes part of the revoked credentials set, which updates with each new revocation. Verifying the state of the VC corresponds to verifying the membership of that VC to the set. If the result of the membership of a VC is true, then the verification fails. Using a cryptographic accumulator, our solution provides a secure and decentralized method to manage credential revocations while maintaining the privacy and efficiency required for LI procedures.

5.4 Design of the SSI stack

To design our solution also from a technological point of view, we refer to the SSI stack [108] and its evolution, the ToIP stack [110], depicted in Figure 5.3a. Therein, four layers are introduced that define how trust is achieved among several participants. In the following subsections, we describe the role of each level and how it is used in our approach. The resulting instantiation of the stack in our proposal is reported in Figure 5.3b.

5.4.1 Layer 1

The first layer refers to identifiers and keys. It allows the entities of the system to access (read and/or write) data without relying on a third central authority. This requires introducing a verifiable data ledger, known as the *DID Registry*.




Fig. 5.3: Comparison between the general SSI stack (a) and its instantiated version (b).

In our solution, we selected Sovrin [105] as DID Register, a special-purpose blockchain deployed for SSI. It is a public and permissioned blockchain [113] (accessible to all entities but only authorized parties can write data on it). We store the following information on the blockchain:

- **Public DIDs of LEA, LEMF, and EA:** They are the DIDs of the entities playing the role of issuers and verifiers in our approach. They are used by the employee E to establish an off-chain connection with them. Moreover, the public DIDs allow the parties acting as verifiers (i.e., the LEA, the EA, and the LEMF) to obtain the public keys to check a cryptographic proof of ownership of a credential.
- **Schema:** It defines the format (in terms of allowed attributes) of the credentials released by the issuers. Observe that it does not contain any personal data about the employee E.
- **Credential Definition:** It is an instance of the Schema including the public keys allowing the verification of the credentials. Similar to the Schema, the Credential Definition does not include any personal data about E. The inclusion of the Schema and the Credential Definition at the blockchain level makes the credential exchanges fully decentralized enabling a fast release and verification of the credentials from anywhere.
- **Revocation Registry:** We include two revocation registers, one managed by the LEA and one managed by the EA. Only the issuer (LEA or EA) can update each register, which links to a Credential Definition. The registers are stored on-chain,

but, again, they do not contain any personal data about E. As the name suggests, they allow the LEA and the EA to revoke the credentials released to the employee E and the LEMF. They are based on the cryptographic accumulator described in Section 5.3.3. The advantage of this approach is that the LEA and EA do not have to maintain a public API for the revocation, but the entities playing the role of verifiers can check the revocation of the credentials in a decentralized way by accessing the blockchain.

- **Agent Authorization Policies:** This component implements a mechanism of protection in the case the device of E is stolen or lost. In the device used by E, an agent is installed to exchange proofs of the credentials owned by E. In the case this device is stolen or lost, a revocation mechanism is enabled. In particular, when E authorizes a new agent (device), an authorization key is stored in an on-chain registry. This key is used to sign messages exchanged peer-to-peer between E and the other entities. To remove the authorization for an agent, E can delete its key from the register.

It is worth highlighting that the following information is not stored on the blockchain:

- **Private DIDs:** They are the DIDs of the actors used in the peer-to-peer connection.
- **Private Data:** They are personal data about the employee E. These data have not been stored even in encrypted form. Indeed, since the data stored on the blockchain are immutable, they could be compromised in the long term.

To conclude, we observe that our solution leverages Peer DIDs to establish peer-to-peer communications. Even though the selection of the communication protocol regards Layer 2 of the stack, the choice to rely on Peer DIDs is performed at Level 1. In this case, the DID Registry is represented by the digital wallets of the entities generating and storing the Peer DIDs.

5.4.2 Layer 2

Layer 2 refers to the communication protocol between agents and the interfaces available to the developers. As mentioned in the previous section, at this level the peer-to-peer communications between agents are established.

The protocol used in our solution is DIDComm [114], which is fully compliant with the ToIP stack. This protocol is based on message exchange and is independent of the underlying transport protocol. Each agent of the communication is identified through a DID and the communication is encrypted through the public key associated with the DID. Since any transport protocol is allowed, the most common choice

falls on HTTP or HTTPS (as in our approach), allowing us to develop a web application based on the request-response paradigm (see Section 5.5). Therefore, no particular technology of the involved parties is required.

At this level, some message-oriented interfaces for the developers are proposed. In detail, we distinguish two types of nodes: *agent* and *validator*. The agent nodes can be hosted everywhere (smartphone, laptop, cloud, and so on). Each agent maintains one or more connections, each one associated with a Peer DID. As already discussed, these connections are used to exchange cryptographic credentials and proofs. The validator nodes allow the agent nodes to interface with the blockchain selected in Layer 1. Specifically, in our approach, the agent nodes rely on the validator nodes to access the public keys associated with the public DIDs of the LEA, LEMF, and EA and/or to write the Schema and the Credential Definition before releasing the credentials (in this case the agent nodes are hosted on the issuers).

5.4.3 Layer 3

Layer 3 defines the format of the credentials and the way in which they are exchanged. Again, to adhere to the ToIP stack, we select the W3C format, which supports ZKP-oriented credentials [115]. The advantage is that the holder (i.e., the employee E) does not need to present credentials, but only a cryptographic proof of them. This allows us to reduce the information disclosed by E to the minimum.

5.4.4 Layer 4

The last layer defines the Governance Framework to offer maximum interoperability between different systems [111]. We rely on the Sovrin Governance Framework [116], which defines legal and commercial rules to which all entities must adhere. Actually, this framework covers aspects regarding all four layers of the stack.

5.4.5 Stack Overview

This four-layer architecture standardizes interactions between LEA, LEMF, and EA without relying on a centralized authority. Without loss of generality, the EA can authenticate itself through a verifiable credential issued by the LEA, which follows the W3C format of Layer 3 and is released using the DIDComm protocol of Layer 2. Additionally, at Layer 1, the Sovrin blockchain acts as a decentralized registry that keeps track of the credential definitions and revocation statuses. This allows different countries or law enforcement agencies to verify digital credentials independently by checking the blockchain, thus ensuring global interoperability. Finally, the Governance Framework defined in Layer 4 ensures that LEA, LEMF, and EA are compliant

with legal and security regulations, especially when working across different countries. When different law enforcement agencies or external agencies from different countries need to share or access intercepted data, they must follow the same legal and security guidelines. For example, in a cross-border cybercrime investigation, a European LEA could gain access to the intercepted data from a U.S.-based LEMF by authenticating itself through a verifiable credential, issued under the W3C standard and verified via blockchain. The U.S.-based LEMF can validate the request without manual verification, ensuring secure and interoperable access. Moreover, the adoption of the SSI stack ensures the interoperability between our system and other SSI systems. In practice, this approach allows us to incorporate the verification and exchange of verifiable credentials from other SSI systems adopting the same standard defined in the stack (e.g., law enforcement badges or digital warrants issued as verifiable credentials) into our access control process.

5.5 Proof of Concept

Through this section, we describe the proof of concept developed to verify the technical feasibility of our proposal (the source code is available at https://github.com/CyberGoldSun/SSI_LI). The proof of concept consists of three web applications that run the procedures done by the LEA, LEMF, and EA, respectively.

5.5.1 Technologies

We developed the proof of concept using several technologies. First, we adopt the Sovrin framework, which is based on Hyperledger Indy [117], a public permissioned blockchain. Sovrin defines the rules governing the blockchain (consensus protocol, validator nodes, the format of the DIDs, and so on). Among its components, Sovrin includes the agent nodes hosted by the entities participating in the system that allow the exchange of credentials and DIDs.

The second technology used is Trinsic [118], a platform based on Hyperledger Aries that offers several APIs to build SSI ecosystems. In particular, Trinsic supports the Sovrin-defined format for credentials and DIDs. Specifically, we rely on Trinsic Studio to manage the organizations realizing the credentials (of the LEA, LEMF, and EA) and on the Android mobile app to host such credentials. Finally, the web applications are realized from scratch in Python by using the Flask micro-framework [119].

5.5.2 Case study

We show the main functionalities offered by our solution through a case study. In particular, we refer to the EA integration in the LI procedure of Section 5.3.2, which is the more relevant contribution of this work. Note that some functionalities (such as the revocation) are not included in this case study but are implemented in the version of the software available on GitHub.

In our case study, we consider an employee Bob of an external agency, called NA, who was charged by the LEA to access data relative to an LI procedure on the LEMF. Initially, the three organizations involved in the system register on the Trinsic Studio platform (see Figure 5.4).




Fig. 5.4: Registration of LEA, LEMF, and NA Agency.

In this registration phase, some fields have to be provided (name, region, blockchain network, and so on). In particular, a relevant field is the issuer DID, which identifies an issuer and is used to verify the credentials it releases. Figure 5.5 shows an example of the fields associated with the NA agency.



NA Agency ✕

API Key
57LVK0gFBUqCJeYLZLqb41C_KnJLw2eTvDtNeUst0yk 

Network
Sovrin Staging Network

Region
United States (Central US)

Issuer DID
GwfdWSDnrRZHvvNKCj8S6v 

Service URL
<https://api.portal.streetcred.id/agent/2pqT21etPLhFki9tRWT1sRCgvQxiiR7> 

Fig. 5.5: Fields of NA Agency registration.

When Bob is hired by NA, he receives an “Employee credential” certifying its status. This credential (with its fields) is stored on the Bob mobile wallet (Android app) as shown in Figure 5.6a. To participate in the LI procedure, Bob needs the “X credential” released by the agency NA. As discussed in Section 5.3.2, to obtain the “X credential”, Bob has to prove the ownership of the “Employee credential”. The agency NA (through the web application) generates a “Verification” of the “employee credential”. This is reported in Figure 5.6b. Observe that optional attributes may be requested to Bob.

Once this Verification is created, Bob can set up a secure peer-to-peer connection with NA to provide the “Employee credential”. This is done by scanning a QR code properly generated in the reserved area of Bob as shown in Figure 5.7a. However, not all the fields of the credential are provided by Bob but just a verifiable presentation of it. For example, in the case shown in Figure 5.7b, the Verifiable presentation includes four of five fields of the credential.

At this point, NA interacts with the blockchain to verify the validity of the credential received by Bob. This is performed in the background on the web application. The snippet of the Python code that implements this function is reported in Listing 1.

Then, NA can release the “X credential” (see Figure 5.8a) that will appear on the mobile wallet of Bob (see Figure 5.8b).

After obtaining the “X credential”, Bob interacts with the LEA to get the “Y credential” for accessing the data on the LEMF. Similarly to the NA agency, the LEA




Fig. 5.6: Overview of the Employee Credential (a) and X Credential Verification (b).

generates a “Verification” for the “Y credential” through the dedicated web application (see Figure 5.9a). Observe that, among the fields required in this verification, there is the issuer DID of the entity releasing the “X credential” (i.e., the agency NA).

At this point, the procedure appears for Bob similar to the procedure performed with the NA agency. Indeed, Bob connects with the web application of the LEA and scans a QR code to provide the LEA with the Verifiable presentation of “X credential”. However, a difference exists with the previous case. Indeed, such a presentation is not provided within a secure connection but is connectionless because the presentation of the “X credential” is required in order to establish the connection with the LEA. Once the “X credential” is verified by the LEA, it releases to Bob the “Y credential” (see Figure 5.9b) allowing him to interact with the LEMF. Specifically, Bob scans a QR code to establish a secure connection with the LEMF and provides it with a Verifiable presentation of the “Y credential”. Then, Bob obtains the data of the LI procedure from the LEMF.

We conclude this section by discussing the data stored in the blockchain. Through the Indy Scan platform, it is possible to monitor all the transactions performed by an entity. Figure 5.10 shows a snapshot of the Indy Scan platform, which reports the transactions performed to build the necessary infrastructure for credential is-




Fig. 5.7: Overview of QR Code (a) and Verifiable Presentation (b).

suance and revocation. The fourth column of the table represents the issuer DID of the node originating the transaction, which is the agency NA. Specifically, we find the following transactions:

1. **Tx 262048 and 262052 (SCHEMA):** These transactions define the structure of two credentials: *Employee Credential* and *X_Credential*, both set to version 1.0.
2. **Tx 262049 and 262053 (CLAIM_DEF):** Claim definitions associate the schemas (*Employee Credential* and *X_Credential*) with the issuer's DID.
3. **Tx 262054 (REVOC_REG_DEF):** This transaction creates a revocation registry for a credential definition, specifying `ISSUANCE_BY_DEFAULT` as the issuance type and a registry capacity of 1024 entries.
4. **Tx 262055 (REVOC_REG_ENTRY):** This transaction logs an entry in the revocation registry, enabling tracking and management of the revocation status for issued credentials.

These transactions show the sequential process of decentralized credential management:

- **Schemas** define the data structure of credentials.

Listing 1 Python function to check the validity of a credential.

```

1
2 @app.route("/send_verification/<string:
3 connection_id>",
4 methods=["GET", "POST"])
5 def send_verification(connection_id):
6     connection_id = connection_id
7     policies =
8     credentials_client.
9     list_verification_policies()
10    policy_id = policies[0].policy_id
11    verification =
12    credentials_client.
13    send_verification_from_policy
14    (connection_id, policy_id)
15    return redirect(url_for
16    ('verifications_connection',
17    connection_id=connection_id))

```

- **Credential definitions** connect schemas to issuers, facilitating the issuance of verifiable credentials.
- **Revocation registries** manage the status of issued credentials, promoting transparency and trust in the decentralized identity ecosystem.

Overall, this setup preserves the integrity of the credential lifecycle, allowing issuers, holders, and verifiers to interact securely within the Hyperledger Indy framework.

5.6 Security Analysis

In this section, we analyze the security properties of the proposed SSI-based access control solution for lawful interception.

5.6.1 Threat Model

Our threat model considers the following actors as potential adversaries: (i) an external attacker with network-level access to the communication channels between system entities; (ii) a malicious or compromised employee of the EA attempting to




Fig. 5.8: Overview of the X Credential issuance process (a) and its storage in Bob's wallet (b).

escalate privileges or maintain unauthorized access after revocation; and (iii) a compromised entity within the system (e.g., a rogue LEMF or a corrupted EA) attempting to forge or abuse credentials.

The security assumptions we adopt are the following:

- **A1:** The attacker cannot break the underlying cryptographic primitives, including the digital signature schemes used for credential issuance and the encryption scheme used by the DIDComm protocol.
- **A2:** The Sovrin blockchain operates correctly and its consensus mechanism prevents unauthorized modifications to on-chain data, including the DID Registry, Credential Definitions, and Revocation Registries.
- **A3:** The LEA is a trusted and honest entity. It acts as the root of trust in the credential chain and its private keys are not compromised.
- **A4:** Private keys held by system entities (LEA, EA, LEMF, and employee E) are stored securely and are not accessible to adversaries.
- **A5:** The zero-knowledge proof scheme used for selective disclosure of credential attributes is secure: no adversary can produce a valid proof of a credential they do not possess.




Fig. 5.9: Overview of the policy for the Y Credential (a) and the Y Credential itself (b).

5.6.2 Security Properties

The proposed system must guarantee the following security properties:

- **Access control correctness:** Only an employee E holding a valid, non-revoked credential chain ($VC_E \rightarrow VC_L \rightarrow VC_{LM}$) can obtain access to the interception data on the LEMF.
- **Credential unforgeability:** No entity can produce a valid credential without the cooperation of the legitimate issuer.
- **Revocation effectiveness:** Once a credential in the chain is revoked, all subsequent access attempts by the holder of that credential must fail.
- **Selective disclosure:** The employee E discloses only the minimum set of attributes required by the verifier, without revealing additional personal information.

TxNo	Type	Timestamp UTC	From DID	Info
262055	REVOC_REG_ENTRY	21 September 2021, 8:29:09 1204 days, 11 hours, 43 mins, 38 secs ago	GwFdwSDnrRZHvVnKCj856v	
262054	REVOC_REG_DEF	21 September 2021, 8:29:06 1204 days, 11 hours, 43 mins, 41 secs ago	GwFdwSDnrRZHvVnKCj856v	Cred definition: GwFdwSDnrRZHvVnKCj856v:3:CL:262052:test Issuance type: ISSUANCE_BY_DEFAULT Revocation registry capacity: 1024
262053	CLAIM_DEF	21 September 2021, 8:29:01 1204 days, 11 hours, 43 mins, 46 secs ago	GwFdwSDnrRZHvVnKCj856v	Schema name: X_Credential Schema version: 1.0
262052	SCHEMA	21 September 2021, 8:28:55 1204 days, 11 hours, 43 mins, 52 secs ago	GwFdwSDnrRZHvVnKCj856v	Schema name: X_Credential Schema version: 1.0
262049	CLAIM_DEF	21 September 2021, 8:28:14 1204 days, 11 hours, 44 mins, 33 secs ago	GwFdwSDnrRZHvVnKCj856v	Schema name: Employee Credential Schema version: 1.0
262048	SCHEMA	21 September 2021, 8:28:02 1204 days, 11 hours, 44 mins, 45 secs ago	GwFdwSDnrRZHvVnKCj856v	Schema name: Employee Credential Schema version: 1.0

Fig. 5.10: Indy Scan.

- **Auditability:** All credential issuance and revocation events are recorded on the Sovrin blockchain and are verifiable by authorized parties.

5.6.3 Attacks and Mitigations

We identify the following attack scenarios and discuss how the proposed solution addresses each of them.

Credential Forgery Attack

An adversary attempts to construct a valid credential VC_{LM} without the cooperation of the LEA, in order to gain unauthorized access to the LEMF.

This attack requires breaking assumption A1, specifically forging the digital signature of the LEA on the credential. Under A1, this is computationally infeasible. Furthermore, the LEMF verifies the credential by checking the issuer's digital signature against the public DID of the LEA stored on the Sovrin blockchain. Any forged credential would fail this verification step, as it would not carry a valid signature from the LEA's private key.

Replay Attack on DIDComm Channel

An adversary intercepts a valid credential presentation exchanged over a DIDComm channel and attempts to reuse it to gain unauthorized access to the LEMF at a later time.

The DIDComm protocol [114] provides message-level security guarantees that are independent of the transport layer. Each message carries unique identifiers that prevent the reuse of intercepted presentations. A replayed presentation, carrying a

stale message identifier, will be rejected by the verifier because a list of the identifiers of already spent messages is maintained. Additionally, credential expiration and revocation checks performed at verification time further limit the window of opportunity for replay attacks.

Revocation Bypass Attack

A previously authorized employee E, whose credential VC_L or VC_{LM} has been revoked, attempts to use the revoked credential to access the LEMF after revocation.

The revocation mechanism is based on cryptographic accumulators (Section 5.3.3), whose current state is stored on the Sovrin blockchain. At each verification step, the verifier (LEMF or LEA) queries the on-chain Revocation Registry to check whether the credential is still valid. A revoked credential produces a membership proof that fails against the current accumulator value, causing the verification to fail. Since the Revocation Registry is on-chain and updated in real time by the issuer, the verifier always operates on the current revocation state. The effectiveness of this mechanism relies on assumption A2.

Insider Threat and Compromised EA

A malicious EA attempts to issue unauthorized credentials VC_L to employees who have not been authorized by the LEA, effectively granting them access to the LI procedure without LEA knowledge.

This attack is mitigated by the structure of the credential chain. The LEMF does not accept VC_L alone as sufficient for access. It requires VC_{LM} , which can only be issued by the LEA after verifying VC_L . Therefore, even if the EA issues VC_L to an unauthorized employee, that employee cannot obtain VC_{LM} from the LEA without presenting VC_L to the LEA for verification. The LEA, acting as the root of trust (assumption A3), controls the final authorization step.

Device Theft and Compromised Agent

The physical device of employee E is stolen or lost, and an adversary attempts to use the credentials stored in E's digital wallet to access the LEMF.

This attack is addressed by the Agent Authorization Policy mechanism described in Section 5.4.1. Each agent (device) used by E is registered on-chain through an authorization key. If the device is reported stolen or lost, E can remove the corresponding authorization key from the on-chain registry, immediately invalidating all credential presentations originating from that agent. This mechanism operates independently of credential revocation and provides an additional layer of protection at the device level. Table 5.1 summarizes the security analysis carried out.

Table 5.1: Summary of threats and mitigations in the SSI-based LI system.

Threat	Potential Impact	Mitigation
Credential Forgery	Unauthorized access to LI data on the LEMF.	Issuer digital signature verified against on-chain public DID; unforgeability under assumption A1.
Replay Attack on DIDComm Channel	Reuse of intercepted credential presentation.	DIDComm message-level security with spent-identifier tracking; credential expiration and revocation checks under transport-independent guarantees.
Revocation Bypass	Continued access after credential revocation.	On-chain cryptographic accumulator queried at each verification; real-time revocation state under assumption A2.
Insider Threat and Compromised EA	Unauthorized credential issuance by rogue EA.	Two-step credential chain requires LEA approval for VC_{LM} ; LEA acts as root of trust under assumption A3.
Device Theft and Compromised Agent	Credential abuse from stolen or lost device.	Agent Authorization Policy; on-chain key removal immediately invalidates compromised agent.

In conclusion, the security of the proposed system relies on the combination of well-established cryptographic primitives (digital signatures, zero-knowledge proofs, cryptographic accumulators), the immutability and transparency of the Sovrin blockchain, and the multi-step credential chain that enforces the LEA as the sole root of trust. Together, these mechanisms ensure that access to interception data is strictly controlled, auditable, and revocable.

5.7 Related Work

Self-Sovereign Identity is an emerging paradigm attracting the interest of governments and enterprises [101], [109], [111], [120], [121]. One of the main benefits of the SSI paradigm is its compliance with the General Data Protection Regulation (GDPR) [122], [123], [124], [125], making it an effective choice to develop privacy-by-design

solutions. As a matter of fact, several concrete implementations of SSI frameworks are available to develop SSI-based solutions [117], [126], [127], [128]. Among them, we rely on Sovrin [105], [129], [130], [131], which is becoming a standard for deploying SSI solutions, as witnessed by several scientific papers [132], [133], [134], [135], [136], [137], [138] and industrial projects [139] proposing different solutions that utilize the Sovrin framework.

In the literature, several solutions have been proposed that apply the SSI approach to solve different problems in different domains. For example, some of these domains are healthcare [140], [141], [142], industrial IoT [143], [144], and food supply chain [145].

In this study, for the first time, we apply the SSI approach to the lawful interception domain. Lawful interception is an activity that involves legal [146], [147], [148], ethical [149], and technological issues [150], [151]. From a technological point of view, several proposals, also recent [93], [94], [152], are available in the literature, especially in the field of telecommunication [153], [154], [155], [156]. However, these works focus on the physical communication interfaces [156] and/or on the integrity and storage of the collected data [157].

This work addresses a slightly different problem, which is access control in the LI procedure. To the best of our knowledge, previous solutions propose access control models (such as RBAC [106] or VACM [107]) at a theoretical level, but their implementations fail to enforce these models effectively, leaving access control unimplemented in practice.

5.8 Conclusion and Future work

In this chapter, we proposed an SSI-based solution to implement an access control mechanism that supports a LI procedure. Our proposal is motivated by the need to overcome the absence of a standard for the interactions between the LEA and LEMF. In the design of the solution, we followed the guidelines of the Sovrin framework. In particular, we adopted the SSI stack, which defines how trust is achieved among several participants. To make our proposal more concrete, we provided a proof of concept using standard and well-consolidated technologies, which allowed us to verify that our proposal is technically feasible.

The findings of this research show that an SSI-based solution can enhance security, automation, and interoperability in lawful interception procedures. Our proposal research defines how access control mechanisms in lawful interceptions can be strengthened, ensuring that only authorized entities, specifically law enforcement agencies or external agencies, can retrieve intercepted data. This is particularly im-

portant in cross-border investigations, where secure access to intercepted communications is often hindered by differences in interception procedures across countries. Our solution eliminates the manual operations of traditional lawful interceptions, which are substituted by machine-readable procedures. This change reduces human error and improves efficiency by accelerating all the processes related to lawful interception. Moreover, the use of verifiable credentials mitigates the risks associated with centralized identity providers, reducing single points of failure and enhancing privacy through selective disclosure. By addressing the lack of standardized interaction protocols between LEA and LEMF, this research introduces a more secure and privacy-preserving approach to lawful interception and defines workflows for all involved entities, with the aim of improving investigative capabilities and enhancing legal compliance.

Alongside the advantages discussed above, our proposal presents some limitations. The first limitation concerns the adoption. Indeed, SSI is an emerging paradigm, which lacks widespread adoption and does not benefit from the sponsorship of major organizations (like Google and Facebook, which support federated identity). Moreover, the lack of uniform regulation regarding the use and validity of SSI makes legal compliance challenging, particularly in law enforcement, where adherence to legal frameworks is critical. Considering the eIDAS Regulation, although it is focused on government eIDs and does not integrate SSI, the success of blockchain convinced the EU Commission to revise regulations to consider the SSI paradigm [158]. Another limitation of our proposal concerns the introduction of operational costs. Since operations done on the blockchain require the payment of a transaction fee, this may change the business model. In our solution, each operation, such as credential issuance or verification, involves a blockchain transaction and incurs a small fee (usually around a few dollar cents), which is not necessary for traditional lawful interceptions. However, we expect that transaction fees introduced by our solution will be negligible compared to the overall costs incurred by organizations for lawful interception, as these expenses are primarily driven by infrastructure and human activity rather than identity and access control verification processes. Exploring these aspects is a natural progression of this research.

As further future work, we plan to integrate the SSI approach also in the interaction between the LEMF and the network operator. In particular, the technical parameters needed to start the LI procedure can be exchanged in the form of credentials through the peer-to-peer channels offered by the SSI approach. In other words, these channels may replace the HI1 management interface defined by the ETSI standard [97].

A Router-based Parental Control Tool for Safe Social Network Usage

Social networks' explosive growth has changed the way people, especially young people, engage with the digital world. There are many worries about the amount of time kids spend online and the content they are exposed to. Parental control software can be used to monitor and limit Internet usage, but they need to be installed and are easily circumvented. The integration of parental control methods directly into the home router, as proposed in this study, offers a novel alternative. The router's central location within the home network facilitates monitoring and management of Internet consumption across all connected devices. This method provides granular packet inspection, scalability, and centralized control, allowing for a thorough examination of network traffic, including social network activity. The architecture of the suggested solution consists of modules for social network identification, traffic analysis, and packet capture and inspection.

6.1 Introduction

The rapid proliferation of social networks has significantly impacted how individuals, particularly young people, interact with the digital world. Platforms like Facebook, Instagram, Twitter, and TikTok have become integral parts of daily life, providing avenues for communication, entertainment, and information sharing [159]. However, this increased usage comes with challenges, especially for parents concerned about the time their children spend online and the content they are exposed to.

Parental control tools have been developed to address these concerns, offering functionalities to monitor and restrict Internet usage [160], [161]. Traditional parental control solutions typically operate on individual devices through dedicated applications. While effective to some extent, these approaches have notable limitations. Device-specific controls require installation and maintenance on every device, which can be cumbersome and easily bypassed by tech-savvy children. Additionally,

these solutions often lack the ability to provide a comprehensive view of Internet usage across multiple devices. Moreover, many existing parental control systems rely on third-party cloud infrastructures to manage monitoring and filtering activities. This dependence introduces serious privacy concerns, as user data and behavioral information are transmitted and processed externally. Such architectures can lead to profiling, personalization, and continuous tracking of minors, exposing sensitive information beyond the household domain.

An alternative approach involves DNS filtering, where requests to certain websites are blocked based on a denylist of domains. However, a simple DNS filtering does not offer granular control over specific activities within websites, particularly dynamic and content-rich platforms like social networks [162].

This work proposes a new solution to these challenges by integrating parental control mechanisms directly into the home router. By leveraging the router's central position in the home network, our system can monitor and manage Internet usage across all connected devices. This approach offers several advantages:

- **Privacy Preservation:** The proposed system operates entirely within the local network and does not rely on external or third-party services, thus minimizing data exposure and eliminating the risk of external profiling or behavioral tracking.
- **Centralized Control:** All network traffic passes through the router, enabling comprehensive monitoring and management.
- **Scalability:** New devices connected to the network are automatically subject to the same controls, eliminating the need for individual setup.
- **Granularity:** Packet inspection allows for detailed analysis of network traffic, including the time and duration of the activities on social networks.

The architecture of the solution includes various modules, whose purposes concern capturing and inspecting network packets, identifying the social networks involved in the activity, analyzing traffic patterns to monitor user activities, and enforcing parental control policies. The proposed solution has been implemented to validate the proposal and a proof of concept is presented.

The rest of the chapter is structured as follows. In the next section, we survey the related work. In Section 6.3, we present the solution proposed to implement a parental control tool. Section 6.4 describes the implementation of the proposal presented as a proof of concept. Finally, in Section 6.5, we draw our conclusions.

6.2 Related Work

Parental control on social networks has emerged as a critical area of research due to the increasing use of these platforms by adolescents and the associated risks. The

importance of parental control is underscored in [163], stating that preadolescents, whose parents reported greater control over their child's time on social media, reported better mental health.

[164] proposes a framework for the design of parental control. It provides instructions to efficiently integrate parental control tools into digital environments, emphasizing how the aspect of granular access controls can be effective in building parental control systems.

Best practices for building parental control systems are even discussed in [165], which outline how the reviews on the Google Play Store related to parental control applications do not satisfy the requirements needed in such systems.

Numerous studies have explored various strategies and technologies to monitor and manage children's activities online, ensuring their safety while respecting their privacy [166].

In [167], the authors propose a privacy preserving parental control protocol with edge computing that uses Artificial Intelligence techniques to automatically detect harmful content for minors in 5G networks.

The aspect of privacy is also investigated in [168], in which a rule-based expert system is used to emulate human intelligence by using rules and conditions is proposed.

The adoption of blockchain [169] is also considered to implement parental control in home environments.

Traditional approaches that implement parental control are based on machine learning techniques [170], [171], [172], [173], [174].

However, many ML approaches assume the availability of the content of the websites accessed by users. This may limit the applicability of the proposed solutions by requiring the collaboration of the providers.

On the other hand, our proposal aims to implement a parental control mechanism in a home environment, without disclosing the rules or habits of the users to external parties.

In particular, our proposal is based on packet inspection [175], [176]. Packet inspection allows for the analysis of data packets transmitted over a network to identify and block harmful content without needing access to the actual content of the websites. This method offers a robust solution for parental control that safeguards user privacy and operates independently of external content providers.

DNS inspection is shown to be effective in blocking children from accessing some platforms [177]. However, simple DNS filtering does not offer granular control over specific activities within social network platforms. Then a more advanced approach that integrates traffic inspection on encrypted traffic [178] and access-based access




Fig. 6.1: Environment considered in the study.

control to enforce policies [179] without altering flows of messages [180] should be investigated.

6.3 System Architecture

In this section, we present the design and architecture of the system. The environment we consider in the following study is depicted in Fig.6.1. We consider a domestic network where all the users' devices access the Internet through a router. The router enables wireless and wired connections within the environment and provides a multi-device parental control service. The devices can be smartphones, laptops, workstations, and so on, and are used by users to whom parental control policies are applied. Moreover, an admin is considered to establish the policy to apply and the actions to execute when a policy is violated.

The core of the architecture is located in the router, which acts as the main actor for the parental control service. The architecture is composed of six modules that are described in the following and schematized in Figure 6.2, where the data-flow is presented with reference to a single device (i.e., this workflow is repeated for each device):




Fig. 6.2: Example of the architecture considered in the study.

1. *Domain Manager* is the module that manages the social network domains to be monitored. Admin can include or remove the domain names that should be monitored by the system.
2. *Policy Manager* is the module that manages the parent control security rules. Admin can include or remove rules for each social network and device.
3. *Domain Name Inspector* intercepts each network packet of the device and searches for DNS queries/responses.
4. *Traffic Inspector* receives a network packet and decides whether this packet should be processed or not.
5. *History Manager* is in charge of storing the traffic of interest to allow for policy compliance checks.
6. *Decider*, the module that eventually decides whether a connection should be processed according to the given policies.

Now, we show how the interactions between the various components work through the following steps:

Step 1. In this step, Admin interacts with both the *Domain Manager* module and the *Policy Manager* module. First, Admin sets within *Domain Manager* the domain names to be monitored. For example, Admin can include `facebook.com` and `tiktok.com` (Step 1a). Next, Admin defines within *Policy Manager* the parental control policies, which can be modified or removed later (Step 1b). For example, Admin may want to prevent connected devices from `tiktok.com` for more than one hour per day, denying future connections, and/or receiving alerts when a device in the network has exceeded the maximum permitted use.

Step 2. *Domain Name Inspector* intercepts all incoming and outgoing Internet traffic generated by connected devices searching for a DNS query. A DNS query is used to ask for the IP address associated with a domain name. Indeed, to reach a domain it is necessary a request to the DNS servers to get the IP address related to that domain. When a request of a device is directed to a DNS server, *Domain Name Inspector* checks whether the field *domain name* of the DNS request matches any of the domains stored in *Domain Manager*. In this case, the DNS query is let go to the Internet (step 2a), and the associated DNS response is waited. Once the response is returned, the IP address associated with the DNS query domain is stored in *Domain Manager* together with the associated domain name. The purpose of this module is related to the fact that IP addresses of social networks can change over time: for example, traffic to a social network can be handled across multiple servers for load-balancing reasons. Again, in case of failure, a server can be replaced by another one with a different IP. In case *Domain Name Inspector* intercepts a DNS query not related to a domain to be monitored, or if the traffic is not a DNS query, then the network packet is forwarded to the *Traffic Inspector* module (Step 2b).

Step 3. Once *Traffic Inspector* receives a network packet from *Domain Name Inspector*, the following actions are done. The destination IP of the packet is extracted and compared with the IPs stored by *Domain Manager* (step 3a). If a match is found (thus, this is a connection to a social network), then the packet is forwarded to the *Decider* module (step 3b), and the packet content is stored in *History Manager* to allow for subsequent statistic analyses (see the next step). If the packet is not a connection to a social network to be monitored, then the packet is let go to the Internet.

Step 4. At this stage, the *Decider* module receives a packet to a social network *S* and accesses the *Policy Manager* module to extract the policies related to the social network *S*. For each policy, the condition is extracted and computed by using the data included in *History Manager* (observe that the type of computation done depends on the rules of the policy). If the policy is not satisfied, then the action expected from the policy is carried out (step 4a). For example, consider the policy introduced in Step 1, which prevents connected devices from `tiktok.com` for more than one hour per day. In this case, once a request to TikTok is intercepted, the *History Manager* is used to compute the total amount of time spent on TikTok from this device, and if this amount is more than one hour, the packet is blocked according to the policy. Clearly, if the connection is compliant with the policy, the packet is let go to the Internet (step 4b).

6.4 Implementation

In this section, we provide a proof of concept of the solution proposed in this chapter. We start with the hardware and software used in the implementation.

6.4.1 Network Infrastructure

We deployed a network infrastructure based on a Netgear R8000 [181] router, on which we installed OpenWrt 22.03.5[182]. OpenWrt is an open-source firmware based on Linux, designed to be installed on embedded devices, primarily routers. Additionally, it provides a completely writable filesystem with package management, enabling an environment where users can customize the functionalities and features of their devices. To give Internet access to the router, we assigned a public static IP address at its WAN port. Additionally, the router generates a Wireless WiFi network (WWAN) with a private subnet, which is dedicated to host devices to be monitored by our parental control system.

Python 3.10.13 has been installed on the router to enable important functionalities and libraries over traffic monitoring. All Python scripts that constitute our multi-device parental control system are executed in a virtual environment on the router.

The router can be directly accessed using a wired or wireless connection, possibly by remote connection with IPsec IKEv2 Road-Warrior VPN [183].

Finally, to have enough memory to conduct our tests and to store intercepted traffic, we increased the storage capacity of the router to 250GB with an external SSD storage drive [184].

6.4.2 Modules

We describe now how the six modules of our proposal are implemented.

1. *Domain Manager* module exploits a dictionary as a data structure to store information about domains. Specifically, each key of the dictionary is a string representing a domain name and is associated with a set of IP addresses. The keys of the dictionary are configured by Admin, who can add or remove domain names related to the social networks to be monitored. Such a dictionary is accessible by the *Domain Name Inspector* module.
2. *Policy Manager* module is another dictionary: it associates each domain name with a set of policies.
3. *Domain Name Inspector* and *Traffic Inspector* modules are implemented through the same Python script. This script monitors all Internet traffic generated from

the devices through Scapy [185], a Python framework that enables powerful network functionalities such as crafting, sending, receiving, sniffing, and manipulating network packets. To efficiently process these network packets, our *Domain Name Inspector* captures and analyzes network traffic in real-time. This is achieved by monitoring the network interface (specifically, the WAN interface) and utilizing the sniff function from the Scapy Python library. Multiple instances of sniff run indefinitely and in parallel, facilitated by threading, to ensure comprehensive and concurrent packet capture and analysis. *Domain Name Inspector* intercepts DNS responses and writes the IP addresses of the considered domains in the dictionary of *Domain Manager*.

Traffic Inspector captures packets having an IP address that is stored in the dictionary by *Domain Name Inspector*.

4. *History Manager* module: It is represented by a collection of pcap files stored in a specific directory on the router. Inside the main directory, there are sub-directories, each corresponding to a specific social network. Within these sub-directories, *Traffic Inspector* stores pcap files not only related to the social network but also specific to individual devices.
5. *Decider*: it is implemented through a Python script. It analyses pcap files stored within *History Manager* through Scapy, which allows current Python programs to read pcap files efficiently. Then, *Decider* receives from *Policy Manager* the policies for the specific social network and, based on the result of the analysis, enforces these policies.

6.5 Conclusion

Digital interactions have changed as a result of the widespread use of social networks, particularly among younger people. Although social media platforms have a number of advantages for contact, amusement, and information exchange, they also present risks associated with overuse and exposure to unsuitable content. Conventional parental control systems, which function on individual devices, provide some safety but have drawbacks in terms of maintenance, installation difficulty, and the inability to give a comprehensive picture of Internet usage across several devices. To address these limitations, this chapter proposed a router-based parental control solution. Using the home router's pivotal function in network traffic management, our system our solution enables comprehensive monitoring and control of Internet usage across all devices connected to the home network. Centralized control, automatic scalability with the addition of new devices, and granular traffic analysis capabilities are key advantages of our proposal. Our implementation, validated through a proof

of concept, demonstrates the practicality and efficacy of integrating parental control mechanisms into home routers. The ability to capture and inspect network packets, identify social network traffic, analyze usage patterns, and enforce control policies marks a significant advancement in parental control technology.

Future work includes the definition of complex policies to be included in the system to enhance its usability and expand its capabilities to adapt to the evolving landscape of digital interactions.

An attack on privacy in IoT Safeguard Environments

Internet of Things (IoT) devices have enhanced daily living, embedding technology into homes, workplaces, and personal accessories. However, this convenience is accompanied by mounting concerns over privacy erosion, due to the extensive collection and transmission of personal data by IoT devices. In this scenario, we analyze a novel challenge: the potential exploitation of IoT devices for covert lawful interceptions, circumventing existing IoT safeguards. In this chapter, we define a system to clandestinely capture audio or video data from IoT devices without an IoT safeguard being able to detect it. Through a proof of concept, we demonstrate the feasibility of the solution, underscoring the need for enhanced privacy measures in the face of evolving technological threats.

7.1 Introduction

We are in an age dominated by interconnected technologies, where the proliferation of Internet of Things (IoT) devices has enhanced the quality of our daily lives. Voice assistants, wearable fitness trackers, security systems, and thermostats are examples of IoT devices integrated into our homes, workplaces, and even our personal accessories. Despite the significant advantages provided by the use of IoT devices, a growing concern exists about the risk of privacy erosion [186].

The problem of privacy arises from the vast amounts of personal and sensitive data that IoT devices collect about their users, which are transmitted to and stored in remote servers. Risks such as unauthorized access, data breaches, user tracking and profiling, are inherent in the use of IoT devices. However, such risks also represent an opportunity in another aspect of modern society: lawful interceptions.

Lawful interception refers to the authorized monitoring or interception of electronic communications by law enforcement agencies or government authorities, typically to maintain national security, prevent crime, or enforce regulations [89], see also Section 2.4. With the proliferation of IoT devices, the scope and complexity of

lawful interceptions have expanded exponentially. Unlike traditional communication channels such as phone calls or emails, IoT devices generate data that can be intercepted and analyzed in real-time, offering unprecedented insights into individuals' behaviors, movements, and interactions [187].

Lawful interception exploiting IoT devices can be contrasted by the mechanisms designed by the technology to protect users' privacy [188]. Notably, the use of an *IoT safeguard* has been proposed to protect user's privacy [189]. An IoT safeguard is a commercial device, often integrated into the home router, that protects against IoT threats by monitoring and registering all outgoing and incoming internet traffic of the IoT device, and when an anomaly is detected, an alert is generated.

In this study, we analyze the use of IoT safeguards against lawful interception conducted by IoT devices, which is a new challenging opportunity. Furthermore, we design a system that can collect video or audio of the environment and send collected data to the government agency conducting the interception, without the interception being detected by an IoT safeguard.

The chapter is structured as follows. In the next section, we survey the related work. In Section 7.3, we define the problem faced. In Section 7.4, we present the solution proposed to capture information (typically voice or video) without an IoT safeguard being able to detect it. Section 7.5 describes the proof of concept performed to test the proposal. Finally, in Section 7.6, we draw our conclusions.

7.2 Related Work

In this section, we present the state of the art regarding the lawful interception domain, showing the most pertinent proposals in the literature. A comprehensive architecture for lawful interception in IP networks, designed by CISCO, is explained in RFC 3924 [190]. This RFC acts as a guideline for network service providers and vendors to implement lawful interception capabilities in an IP network using CISCO Technology.

ESTI (European Telecommunications Standards Institute) provides directives and recommendations to implement functionalities of lawful interception into IT networks [191], by specifying technical requirements, reference architecture, discussions about technical and security issues, and an overview of characteristics of legislation and regulations as well as international cooperation on lawful interception.

The interest in lawful interception on cellular networks is rising rapidly, especially in 5G networks, where virtualized components are adopted. Reference [192] specifies both the architectural and functional system requirements for lawful inter-

ception in 3GPP networks. National regulations determine the set of lawful interception functional capabilities that apply to a specific 3GPP operator deployment. Lawful interception in 5G networks has been discussed in [193], [194].

Reference [195] proposes a secure turnkey framework for lawful interception on which Quantum Computing and Artificial Intelligence tools can enable easier extraction of intelligent information from large volumes of intercepted data and IoT devices.

Concerning the Internet of Things, network steganographic techniques and covert channels have been proposed in [196]. Reference [197] introduced the concept of covert channels as a practical method to extract data from these devices, and a new hidden channel for IoT is presented. The Stealth mode mimics legitimate traffic patterns, making it difficult to detect. The Bandwidth mode sends traffic at the maximum speed allowed by the device and is intended for use on an unmonitored network.

Finally, the creation of hidden channels that can be used to communicate covertly is proposed in [198]. This is achieved by exploiting the overflow field of the timestamp option of Internet Protocol, version 4, to create a covert channel that works over a LAN.

The above discussion shows the importance of the topic addressed in this work. To the best of our knowledge, our study is the first one analyzing the use of IoT safeguards against lawful interception conducted by IoT devices. Moreover, we show the limits of this kind of protection by designing a system that can exploit IoT devices to collect the personal information of the user and to send such data to the government agency conducting the interception, without the IoT safeguard is able to detect it.

7.3 Problem Formulation

In this section, we describe the problem our proposal aims to solve. The scenario we consider regards how to carry out a stealth lawful interception of a user under investigation. Without loss of generality, we refer to an interception within a user's home but the proposal works also for other environments such as the user's office or another physical place where the user spends time. The objective of the proposal is to record video and/or audio occurring in the environment and send it to a *control center*, where interceptions are processed. The latter task should use only the WiFi/LAN/Internet connection of the user available in the place in which the interception is done. We assume that no other communication channel can be created (for example, Bluetooth) should be exploited (for example, for distance reasons, due

to the environment shielding, or because such connection would be detected by the user).

In our scenario, we assume that the user adopts an *IoT safeguard* to protect data privacy. An IoT safeguard is a commercial device, often integrated into the home router, that protects against IoT threats by blocking malicious or otherwise undesirable connections, detecting network intrusions, malware, and anomalies, and mitigating exposure of personally identifiable information [189]. In brief, this device can monitor and register all outgoing and incoming internet traffic through the WAN port, and when an anomaly is detected, an alert is generated. One of the main goals of an IoT safeguard is to detect that suspicious servers are contacted. To do this, an allow-list approach is adopted: the IoT safeguard stores a list of IPs of allowed servers, and requests to other IPs are detected as anomalous.

In summary, the problem we face is 1) to collect video and/or audio of the environment and 2) to send them to the control center without the interception is detected by an IoT safeguard.

7.4 Design and Architecture of the System

In this section, we present the solution proposed to capture information (typically voice or video) without an IoT safeguard is able to detect it.

The two main actors in our scenario are:

1. a Law Enforcement Agency (LEA), the government agency that conducts the interception;
2. a user, the subject under interception by LEA.

In Fig. 7.1, we sketch the environment in which the interception is carried out (e.g., the house of the user) and the architecture of the interception system. In this figure, we do not explicitly represent the IoT safeguard installed by the user to detect any attempt to be intercepted (it is included in the router element).

The system we propose to intercept the user is composed of the following components:

1. *Router*, the device that provides wired and wireless Internet connection in the user's home. In our approach, the router is not controlled and monitored remotely by LEA.
2. *Interceptor*, an IoT device that LEA tampered with and is used to capture the information to monitor. In our presentation, for the sake of simplicity, we consider that only one Interceptor is used. However, it is easy to understand that the approach also works in the case in which more Interceptors are used.




Fig. 7.1: Example of the environment considered in the study.

3. *Concentrator*, a device that receives data from one or more interceptors.
4. *Filter*, a component that receives intercepted data and decides which information is relevant.
5. *ISP* (Internet Service Provider), the company that provides Internet service to the monitored user.
6. *LEMF* (Law Enforcement Monitoring Facility), the place where the results of interception are processed.

How the interception process works in our solution is described by the following steps:

1. *Interceptor Injection*.

In our proposal, the interception is carried out by exploiting a commercial IoT device (i.e., an LED lamp) owned by the user, which is connected to the Internet network. This device should be compromised to work as an interceptor and run our scheme.

If the IoT device is able to record sound or video, a software compromise is enough to make it an interceptor. This can be done by hacking the IoT device leveraging some software vulnerabilities, and gaining root access. In many scenarios, having root access grants LEA the privilege to use the hacked IoT device as an Interceptor.

If the IoT device is not able to record sound or video, a hardware compromise is needed, which involves the hiding of a microphone or a cam within the IoT

device. To do this, one of the possibilities is that LEA uses social engineering to exploit human factor vulnerabilities. For example, LEA can act as the vendor of the IoT device and simulate a fault of the device, to force the user to replace the device. Then, a fake operator sent to repair the fault can substitute it with an identical but pre-tampered with model.

Another possibility can occur if the user buys a new IoT device on an e-shop: LEA can intercept the IoT device before its first delivery, compromise the device, and deliver it to the user.

2. *Concentrator injection.*

At this stage, the goal of the LEA is to install the Concentrator inside the place under investigation. The Concentrator could be placed within an in-wall electrical junction box, either inside or outside the flat. It is important to note that the Concentrator must be in the coverage range of the WiFi network because it acts as a node of the WiFi network.

The Concentrator collects the data (audio or video) captured from one or more Interceptors and sends them to the Data Filtering.

3. *Data Filtering.*

This stage is used to reduce the amount of data to send to *LEMF*. Thus, this reduces the probability that the IoT safeguard can detect anomalies in the traffic. In fact, Interceptor could generate a lot of data that are collected by Concentrator, and Concentrator could make available to LEA only information that is considered interesting. The selection of relevant information can be done in several ways. For example, in the case of the interception regarding voice, a speech-to-text approach can be used to convert the voice into strings, which produces less data traffic. Again, in the case of video, some frames could be extracted (e.g., frames in which a face is detected). For these purposes, approaches based on machine learning could be used to select relevant information [199], [200], [201]. In our deployment, we utilize some machine learning techniques to reach our goals (see Section 7.5).

4. *Information Leakage.*

Our goal is to prevent the IoT safeguard from detecting suspicious traffic by monitoring all outgoing traffic. For this purpose, we have two possibilities.

- a) In the first case, Concentrator creates an ad-hoc communication channel to transmit captured data to *LEMF*. This channel could be based on another wireless network, a Bluetooth network, or a cellular network such as GSM or LTE. In Fig. 7.1, this possibility is represented by the surveillance van parked close to the monitored area. However, in some cases, this possibility could be unfeasible. For example, creating a different WiFi network can be suspicious

even if it has a common Extended Service Set Identifier (ESSID) or there is no possibility to place a van close to the monitored area. In these cases, the next possibility can be adopted.

- b) Using the Internet connection available in the monitored area without the IoT detects it is a challenging opportunity. In this case, we propose to create a covert channel on the Internet connection to export the most relevant data produced by Data Filter.

In this chapter, we introduce a method that generates a covert channel over the WAN spreading the relevant information across multiple TCP packets using the data payload field [202]. Observe that the payload is injected into packets that appear legitimately produced by the IoT device (i.e., the Interceptor). The IoT safeguard recognizes these packets as legitimate because they are directed to the IoT server. In order that the injection of payload into the packets does not generate errors in the communication between the IoT device and server, the Internet Service Provider could be delegated to clean the packets from the payload before the packets are routed to the IoT server.

5. *Data Preservation.*

Data Filtering modifies the intercepted data and generates new data. This modification can result in a problem concerning the probative value of the digital evidence because data integrity cannot be guaranteed.

To reach data preservation, our proposal exploits the properties of Blockchain. As discussed in Section 2.2, the information stored on the Blockchain can be verified transparently, enhancing human trust among the involved parties, as it is possible to trace the origin and history of every piece of data recorded on the Blockchain.

To reach data preservation, Concentrator collects all (non-filtered) data in files and periodically (for example, each day) computes the hash value of these files by a cryptographic hash function [203]. The computed hash is sent to LEMF together with the filtered data.

Once LEMF receives this hash value, a Blockchain transaction is generated having this hash value as payload.

By exploiting the properties of Blockchain, it is guaranteed that such a hash value cannot be modified or deleted by anyone. This transaction works as a sort of timestamp, which demonstrates the existence of a given digital evidence at a given date.

The next step explains how to obtain the original data from the hash values.

6. *Original Data Transfer.*

When the information received via the covert channel is evaluated relevant, it is necessary to retrieve the (non-filtered) intercepted data, which as mentioned earlier, are stored by Concentrator. However, also in the case that no relevant data are intercepted, it could be necessary to retrieve intercepted data at certain time intervals (e.g., every half-month).

To retrieve all intercepted data, the physical access to Concentrator is necessary (for example, an operator can simulate a problem in the network). Once the access to Concentrator is obtained, the original micro SD containing the intercepted data (if any) can be replaced with a new one. Alternatively, an SSH connection can be used to perform a forensic copy of the partition that contains the entire data.

After the intercepted data have been transferred, it is calculated the hash value on the same data on which the Concentrator has computed the hash at the previous step (the same cryptographic hash function is used in both cases).

If these two hashes coincide, then the integrity of the intercepted data is guaranteed. Moreover, the presence of the same hash value in a transaction on the Blockchain is proof that such data have been intercepted at the time of the transaction and have not been modified later.

7. *End of investigation*

At the end of the investigation, all the interception tasks are stopped and all the tampering with actions are rolled back. This concludes the interception scheme.

In the next section, we describe an example of how our scheme can be instantiated.

7.5 Proof of Concept

To implement our proposal, we simulated the scenario described in Section 7.3 and created the network architecture reported in Fig. 7.2. We introduced a router Mikrotik hEx series RB750Gr3 [204] to provide Internet connectivity to the environment (this component is not present in the proposal model of Figure 7.1). Specifically, its WAN port utilizes a static public IP provided by the Internet Service Provider and generates a subnet 192.168.88.0/24 on the Ethernet ports. The code of this proof of concept is available at <https://github.com/CyberGoldSun/IoT4LI>.

In our proof of concept, the components of the model have been instantiated by the following hardware (we show each component labeled with the corresponding item numbers in Fig. 7.3):




Fig. 7.2: Architecture of the proof of concept.

1. Router is implemented by a Netgear WAC510 WiFi router, which represents the router of the user under investigation. The port Eth5 of Mikrotik was connected to the WAN port of Netgear, which has been assigned the static IP address of 192.168.88.100. Moreover, Netgear generates a WiFi network with the private subnet 192.168.1.0/24. The IoT safeguard module has been added to Netgear. First, OpenWrt 22.03.0 has been installed on Netgear. Then, a Python script was installed to track all the incoming and outgoing Internet traffic through the WiFi and the WAN port. The traffic is saved into .pcap files to be analyzed.
2. Interceptor is done by a webcam. This could be a webcam of the user that has been compromised to intercept video and audio of the environment. The webcam is connected to the Concentrator.
3. Concentrator is implemented by a Raspberry Pi 4 Model B [205]. We connected Raspberry to the WiFi network with the IP address 192.168.1.100.
4. Filter is a machine learning software module executed on Raspberry developed to perform some actions over the captured streaming flow. In our proof of concept, this module is devoted to capturing a face from the video.
5. ISP is implemented by a PC, which intercepts the traffic between the OpenWrt and the Mikrotik.



Fig. 7.3: Components used in the proof of concept.

6. LEMF is not represented in our proof of concept for obvious reasons.

Now, we describe the four most relevant tasks performed by the components (the other tasks are trivial).

1. *Data Collection.*

In our proof of concept, the webcam sends the recorded video to Raspberry. This traffic is a continuous, unbounded, and endless flow of data to Raspberry. To make this streaming more efficient, we adopt a publisher and subscriber messaging architecture exploiting Apache Kafka [206], a distributed streaming platform. Apache Kafka uses three main entities, namely Producer, Topic, and Consumer, which work as follows:

- Kafka Producer is an application that publishes data records on a topic. In our prototype, all intercepted data are written as a stream to the topic;
- Kafka Topic is a channel that allows us to organize and categorize data, where data records are published by one or more Producers and consumed by one or more Consumers. A topic represents the name of the directory containing the file to which Kafka appends the records;
- Kafka Consumer is an application that reads data from Kafka Topic and processes them. We developed a consumer that reads intercepted data from Kafka Topic.

Kafka Producer and Topic are configured and run on Raspberry. To use Kafka Producer and Kafka Consumer, we rely on confluent-kafka Python library [207]. It provides an efficient and high-performance interface for producing and consuming messages from Kafka Topic, handling partitioning, message offsets, and other aspects of Kafka's messaging system. It offers both synchronous and asynchronous APIs for interacting with Kafka brokers, allowing Python applications

to integrate seamlessly with Kafka-based data pipelines and stream processing systems.

2. *Data Filtering.*

The data filtering task is performed on the flow of the stream of the recorded video (handled by Kafka Topic). In our proof of concept, we implemented a face detection module that recognizes a face within this streaming video. It is written in Python, and we exploited the face-recognition 1.3.0 Python library [208], which provides many functionalities to support face detection, face recognition, and facial feature extraction, thanks to computer vision algorithms. These algorithms are used to analyze and recognize faces within images or video streams. In addition, we used the OpenCV Python library [209], which provides useful functions for image and video processing, including capabilities for face detection.

This software module runs on Raspberry and determines in real-time whether a face is present in a frame. This module analyzes a data stream to search for standard landmarks such as eyes, mouth, and nose. It operates by looping through every captured frame to perform the face detection process. Subsequently, it has the capability to select the frame of the detected face, which will be sent to the LEMF.

3. *Information Leakage.*

When a face is recognized in a video frame, this frame should be sent to the LEMF. In our proof of concept, the video frame is sent to the PC. The Information Leaked stage is performed by Raspberry. Raspberry knows the private IP address of the IoT device: this can be done by exploiting *nmap* capabilities [210], with a host discovery scan using the parameter `-sn` or using a SYN Stealth Scan (`-sS`). A covert channel is created exploiting Scapy [185], a Python program that enables the user to send, sniff, dissect, and forge network packets.

The frame detected by Data Filtering is encoded into a JPEG image format. Since the payload size of a TCP packet is limited, the image is divided into a list of chunks, with a predefined maximum size (1,000 bytes in our experiments). An iteration is performed over the list of chunks and each chunk is sent in a separate TCP packet as a data payload exploiting the send function from the Scapy library. The TCP packet is constructed by copying the source IP address, the source MAC address, and the destination IP address from a valid packet sent by the IoT device. Note that the destination IP is the server or the cloud associated with the service provided by the IoT device, thus this packet does not generate an alert to the IoT safeguard.




Fig. 7.4: Packet flow during covert channel: red packets are injected by Raspberry and removed by Provider.

Fig. 7.4 shows the standard communication between the IoT device and the trusted IoT cloud server, in which legitimate packets are represented by green boxes. Raspberry uses this standard communication to create the covert channel and to perform information leakage by injecting new packages having in the payload the video frame to be sent on the covert channel according to the technique described in Section 7.4. Injected packages are represented by red boxes. It is important to note that these injected packets pass through the IoT safeguard but are not detected as a covert channel because the destination IP of these packets is legitimate. Before reaching the trusted cloud server, a node of the Internet provider (the PC in our proof of concept) processes all packets and detects packets injected by Raspberry because of their specific payload size. Note that we used this method to detect injected packages thanks to its simplicity. Obviously, more robust techniques could be used: for example, injected packages could be identified by including a specific header into the payload as done by magic numbers used to control the reading of structured data [211]. Note that we are assuming the Provider supports the interception by removing injected packets because we are dealing with lawful interception: this consideration explains why the focus of our research is on lawful interception.


4. Data Preservation.

To implement the data preservation phase, we need to utilize a Blockchain. We relied on Bitcoin, a public Blockchain built in 2008 [8]. We installed a Bitcoin client [212] and used the Bitcoin test network. Subsequently, we created a Bit-

coin wallet on the test network and, in our controlled environment, we initiated a simulation of interception for three consecutive days. Each day, for three days, at a specific time, we calculated the SHA-256 hash of the directory where all intercepted data is stored. Once each hash was calculated, a transaction was generated from and to our wallet address, having the previously calculated hash as the value of the OP_RETURN parameter. This transaction, being public, can be retrieved and inspected by anyone, and the details including the value of the OP_RETURN field can be read simply by using a Bitcoin Explorer service [213].

Note that in our proof of concept, we neglected the possibility that the under-investigation user can detect the presence of the Raspberry by looking at the devices connected to the network from the router administration interface. To contrast this possibility, Raspberry could perform a man-in-the-middle attack between the IoT device and the router. In this case, the IoT device disconnects from the router network and its place is taken by the Raspberry, which spoofs its IP, MAC, and host-name. Then, the IoT device connects to a second hidden Wi-Fi network (e.g., in the new range 192.168.32.0/24) and the Raspberry acts as a router between the network of the home router and the new network. In this way, the presence of the Raspberry cannot be detected by the user from the analysis of the devices connected to the home network.

We conclude this proof of concept by showing that the proposed solution is effective. We analyzed the traffic exiting the IoT device (i.e., Interceptor) as done by the IoT safeguard. We establish the covert channel in which Raspberry injects intercepted data (a face in our case) over the (legitimate) traffic generated by the IoT device. In Fig. 7.5, we show the result of the analysis of the traffic captured by Wireshark filtered on the package having the IoT device as source and directed to the cloud server that in our scenario is represented by a machine having 192.167.51.190 as a public IP address. The first packages (from No. 137 to No. 146) are legitimate (i.e., actually sent from the IoT device). Then, the next 12 packages (from No. 253 to No. 289, which we highlighted by a gray background) are injected packages, which transport the captured face through the covert channel (note that the size is greater than the 1,000 bytes of the payload). Finally, the remaining packages (from No. 317) are again legitimate. It is clear that the IoT safeguard component cannot distinguish injected and legitimate packages because they come from and arrive at the same IP addresses. Note that, in this analysis, legitimate packages have a length (i.e., 66 bytes) different from that of injected packages. This is due to the size of the payload that we choose. However, we could modify the payload size in such a way that all packages have the same length too. For the sake of simplicity, we avoided doing this.



No.	Time	Source	Destination	Protocol	Length
137	2.006820	192.168.1.168	192.167.51.190	TCP	66
138	2.006820	192.168.1.168	192.167.51.190	TCP	66
139	2.006821	192.168.1.168	192.167.51.190	TCP	66
146	2.026997	192.168.1.168	192.167.51.190	TCP	66
253	9.424586	192.168.1.168	192.167.51.190	TCP	1054
257	9.495100	192.168.1.168	192.167.51.190	TCP	1054
260	9.550005	192.168.1.168	192.167.51.190	TCP	1054
266	9.608219	192.168.1.168	192.167.51.190	TCP	1054
270	9.662193	192.168.1.168	192.167.51.190	TCP	1054
272	9.719947	192.168.1.168	192.167.51.190	TCP	1054
274	9.790586	192.168.1.168	192.167.51.190	TCP	1054
277	9.841098	192.168.1.168	192.167.51.190	TCP	1054
279	9.892940	192.168.1.168	192.167.51.190	TCP	1054
282	9.956841	192.168.1.168	192.167.51.190	TCP	1054
285	10.014142	192.168.1.168	192.167.51.190	TCP	1054
289	10.072665	192.168.1.168	192.167.51.190	TCP	807
317	11.975273	192.168.1.168	192.167.51.190	TCP	66
321	12.003789	192.168.1.168	192.167.51.190	TCP	66
322	12.004275	192.168.1.168	192.167.51.190	TCP	66

Fig. 7.5: Traffic flow at the Provider.

7.6 Conclusion

The proliferation of IoT devices has transformed the way we live, enhancing efficiency in our daily routines. However, this technological advancement has also brought significant concerns regarding the erosion of privacy, as IoT devices continuously collect vast amounts of personal and sensitive data about users. The evolving landscape of lawful interceptions adds another layer of complexity to the privacy dilemma posed by IoT devices. While lawful interceptions serve legitimate purposes such as national security and crime prevention, their capabilities, expanded by the IoT, increase the tension between security imperatives and privacy rights. In our analysis, we explored the potential of IoT safeguards as a means of protecting users' privacy against lawful interceptions. By monitoring and analyzing the internet traffic, IoT safeguards could detect anomalies and alert users to potential threats, including those related to lawful interception. In this scenario, we designed a system for circumventing IoT safeguards in lawful interceptions, which presents a novel and challenging opportunity for research and development. By exploiting stealthy data collection techniques, we highlighted the limits of IoT safeguards. Through a proof of concept, we demonstrated the effectiveness of our system in bypassing the controls imposed by IoT safeguards.

Privacy issues in Generative AI Browser Assistants

Browser assistants have started to integrate powerful capabilities of GenAI in web browsers to offer functionalities such as question answering, content summarization, and agentic web navigation. These assistants, available today as browser extensions, raise significant privacy concerns because they can track detailed browsing activity (e.g., searches, clicks) and autonomously perform tasks such as form filling. In this thesis, we analyze the design and behavior of GenAI browser extensions, focusing on how they collect, process, and share user data, and whether they profile users based on explicit or inferred demographic attributes and interests. We develop a novel prompting framework and perform network traffic analysis to audit the nine GenAI browser assistants for tracking, profiling, and personalization.

We find that GenAI browser assistants typically rely on server-side APIs rather than local models, and can be invoked automatically without explicit user interaction. GenAI browser assistants often collect and share full webpage content, including the HTML DOM and user form inputs in some cases, with their first-party servers. Some also share identifiers and user prompts with third-party trackers such as Google Analytics. This data collection and sharing happens even on pages containing sensitive information, such as health records or personal information such as names or social security numbers entered in a web form. Moreover, several GenAI browser assistants infer attributes (e.g., age, gender, income, interests) and use them to personalize responses across browsing contexts. Our findings show that GenAI browser assistants collect and share personal and sensitive information for profiling and personalization, highlighting the need for safeguards as they increasingly mediate web browsing.

8.1 Introduction

With recent advances in natural language processing, Large Language Models (LLMs) trained on massive datasets and built with billions of parameters have significantly

improved their ability to understand context and generate content across modalities, including text, images, and videos [18]. As the foundational technology behind many Generative Artificial Intelligence (GenAI) systems, LLMs now power a wide range of applications such as chatbots and workflow automations. Popular search engines have also begun integrating LLMs. For instance, Google Search now displays Gemini-generated overviews in response to user queries [214], while Microsoft Bing’s Copilot uses models from OpenAI [215]. However, these search engines are limited to analyzing activity on their own platforms, restricting their visibility of a user’s browsing activity across the web. To address this limitation, GenAI browser assistants have emerged to use LLMs to enhance a user’s overall browsing experience across the web.

GenAI browser assistants are implemented as browser extensions, giving them access to nearly everything a user does in their browser. By operating as a browser extension, these assistants can track a user’s browsing activity across websites, such as web pages visited, buttons clicked, search terms, and other personal information typed into form fields. This browsing activity provides context that assistants can use to build detailed user profiles. While this enables personalized responses by these assistants, it also raises serious privacy concerns. For example, browsing data and user profiles can be collected, shared, and repurposed for other usecases such as online behavioral advertising. Despite high operational costs [216], most LLM platforms do not yet rely on ads for monetization. However, this shift is imminent [217], [218], [219].

It is important to understand the privacy risks posed by tracking, profiling, and personalization by GenAI-based browser assistants. Millions of users already use these assistants, yet little is known about how they are designed, what information they collect and share, and how they process it. To this end, we propose a novel framework to systematically audit GenAI browser assistants, guided by the following research questions:

RQ1. *How is the architecture of GenAI browser assistants designed?* We analyze network traffic of these browser assistants to examine key design choices: (1) back-end model capabilities, (2) response architecture, (3) context isolation across interactions, and (4) response variability.

RQ2. *Do GenAI browser assistants collect and share user information?* We analyze network traffic of these browser assistants to examine how user data is handled, focusing on: (1) implicit and explicit collection and sharing, (2) of data with first-party and third-party servers, and (3) differences across public and private online spaces.

RQ3. Do GenAI browser assistants profile users based on their browsing behavior to personalize responses? We develop a prompting framework to evaluate profiling and personalization, focusing on: (1) whether assistants infer user attributes such as location, age, gender, income, and interests from browsing activity, and (2) whether they use this information to personalize responses within and across browsing contexts.

We summarize our key contributions and findings below:

- **Architecture.** We analyze nine popular GenAI browser assistants and find that all but one rely on server-side response generation; only a single assistant performs response generation on the client side. Some assistants are automatically invoked when a user enters a query on a search engine. Furthermore, seven assistants isolate context across browsing sessions and tabs, while the remaining two do not do so.
- **Tracking.** Using assistants in private online spaces resulted in data collection ranging from partial webpage content to full DOM snapshots. Merlin was observed to extract form inputs, leading to the collection and sharing of sensitive information, including medical records from a university health portal (`hem.ucdavis.edu`), academic records from Canvas (from `canvas.edu`), and a Social Security number entered on an IRS website (`irs.gov`). Sider and Merlin shared chat and user identifiers with `google-analytics.com` while TinaMind shared them with `analytics.google.com`. Merlin also shared raw user queries with Google Analytics, which could potentially be used for tracking and personalization across Google platforms. Chat histories were shared with first-party servers by four assistants. Moreover, Harpa and Copilot stored the complete chat history in the extension's background service worker using IndexedDB, indicating that these histories persist across browsing sessions.
- **Profiling and Personalization.** Two extensions (Monica, and Sider) demonstrated profiling based on all five user attributes (location, age, gender, income, and interests) and responded to both in-context and out-of-context personalization prompts. In contrast, Perplexity and TinaMind showed no strong evidence of profiling or personalization, while Harpa exhibited only in-context personalization behavior.

8.2 Related Work

In this section, we review related literature along two dimensions: privacy issues in browser extensions and privacy issues in GenAI. We conclude by discussing how our work compares to and builds upon prior research in this space.

Privacy issues in browser extensions. While browser extensions add to functionality, they pose significant privacy risks due to their access to user data such as browsing history, local storage, payment information, and more [220]. These risks

stem from the potential exploitation of extensions by malicious actors seeking unauthorized access to user data [221], [222]. Prior research has examined various aspects of these privacy risks, such as detecting spying extensions [223], identifying inconsistencies in privacy practices [224], and finding data leaks [225]. There has also been some work from the security perspective, focusing on detecting vulnerable and malicious extensions [226], enabling cloud-based security analysis [227], and security analysis of sensitive data access [228], [229]. Chen et al. [230] found that 2.13% of Chrome browser extensions, representing over 60 million users, leak sensitive information such as browsing history, open tabs, passwords, and location data. Xie et al. [231] further identified hundreds of extensions that automatically extract user content from web pages [232]. Well-known examples include popular extensions like PayPal’s Honey, Capital One Shopping, Hola VPN, and Avira Safe Shopping, which exfiltrate URLs and other user data. Alarming, around 40% of Chrome extensions, including the widely used ones, exhibit security vulnerabilities that could compromise user data [222]. Moreover, Ling et al. [233] found that 92% of browser extensions collect data in ways that contradict their stated privacy policies. Many extensions were found to collect excessive data beyond what was necessary for their intended functionality. For instance, while the Chrome extension *InserLearning* claimed to only collect a user’s name, Google account email, and profile image, but was also found to collect full webpage content. Additionally, Carnus et al. [234] demonstrated that some extensions leak user data, such as email addresses, names, and phone numbers, to their developers and third parties.

Privacy issues in GenAI models. Past work in this space has mostly focused on investigating privacy issues resulting from the training phase of GenAI models. Researchers have demonstrated that LLMs can overfit to sensitive information within training datasets, potentially resulting in the unintended exposure of private data during interactions through APIs, browser extensions, or other interfaces [235], [236], [237], [238], [239], [240]. Such risks are amplified by the susceptibility of LLMs to various attack vectors, including model inversion attacks (reconstruction of sensitive data from model outputs), adversarial attacks (e.g., prompt injection), and data poisoning attacks (introduction of malicious data into training sets) [241], [242], [243]. To mitigate these risks, prior research has investigated privacy-preserving methods such as differential privacy, federated learning, and secure multi-party computation aimed at reducing the exposure of sensitive user information [244]. There remain critical gaps to prevent GenAI models from learning sensitive attributes during the inference phase. Staab et al. show that even limited data from users’ online activities on platforms like Reddit or Twitter enables an LLM to infer a user’s personal attributes with 85% accuracy [245].

Comparison with prior work. Past research has primarily focused on privacy-invasive data collection by browser extensions and the associated security vulnerabilities that expose user data to malicious actors. However, GenAI browser extensions integrated with LLMs introduce new risks, as discussed in Section 2.5.2. These risks include potential exposure of browsing activity, form inputs and other interaction to LLMs, which are highly effective at inferring sensitive user attributes from seemingly benign interactions [245]. Our work builds on the findings of Staab et al. [245] by examining the compounded privacy risks of combining LLMs with browser extension capabilities, an area that has been understudied. To the best of our knowledge, ours is the first systematic audit focused specifically on GenAI browser assistants. While prior work has analyzed the collection and sharing of user data, they have not investigated how that data is subsequently used. GenAI browser assistants are a new target for privacy analysis, allowing us to investigate not just data collection and sharing but also how it may be used for profiling and personalization. We investigate whether continuous tracking of user activities within the browser allows assistants to associate memory with a user’s personal and behavioral traits, using updated context to personalize responses over time. Being a fairly new and developing space, it is both important and timely to investigate the unique privacy risks they pose. To that end, we design a novel experimental framework to audit tracking, profiling, and personalization across three major functionalities (search, browsing, and summarization) offered by nine of the most widely used GenAI browser assistants.

8.3 Threat Model

Our threat model considers online users who install a GenAI browser assistant to personalize their browsing experience. In this threat model, the browser extension of the assistant is treated as the adversary while the user is the victim. The adversary operates as a browser extension, which provides it access to a user’s complete browsing activity. The primary goal of the adversary is to provide personalized responses to user queries. To achieve this, it is assumed to engage in user tracking by monitoring browsing activities and extracting contextual information. This user tracking is considered the secondary (or implicit) goal of the adversary. Adversary may share the collected information from the victim’s browser either to its own servers or to third-party servers.

Figure 8.1 illustrates our threat model. The victim can interact with the GenAI browser assistants in different ways such as initiating a Google search, summarizing a web page, chatting about the web page’s content, etc. Any such activity that involves a user’s interaction with the GenAI assistant is assumed to be a user’s query. Upon receiving a user query, the assistant sends it along with metadata to its own




Fig. 8.1: Threat model of GenAI browser assistants.

backend server. The metadata may include user-, device-, webpage-, or browsing-specific information that serves as a *context* for generating a personalized response. The assistant's server may generate the response internally or share the query and context to an external LLM API. The personalized response is sent back to the victim's browser.

The assistant may also set and share cookies from the victim's browser to its own server. These cookies could further be shared with third-party advertising or analytics platforms to support analytics or ad targeting, such as building custom audiences for personalized advertising. To fulfill its primary goal, the adversary is assumed to continuously build and refine a profile of the user based on accumulated context over time. We consider the LLM provider and third-party trackers as beneficiaries of the data collected and shared by the assistant. Although they may gain this information, we do not assume them to be directly involved in data collection.

8.4 Methodology

In this section, we introduce our auditing framework for GenAI browser assistants, designed to evaluate user tracking, profiling, and personalization, as shown in Figure 8.2. We make our crawling and analysis framework available at <https://zenodo.org/records/15643498>.

8.4.1 Selection of GenAI Browser Assistants

Among the three categories of GenAI browser assistants described in Section 2.5, we focus on *search-based* assistants. This is due to several reasons: (1) Querying a search engine is relevant to querying an LLM model, allowing a natural integration of LLM-generated responses to a user's online searches. (2) It has been argued that search en-

Table 8.1: Overview of studied AI browser assistants sorted by their popularity.

Legend: **Personal Data** : Personally Identifiable Information (PII), Personal Communications (PC), Financial Information (FI).

Web Data : User Activity (UA), Web History (WH), Website Content (WC), Location (LOC).

No Data : No Data collected.

Extension Name	Install Counts	Supported Model(s)	Default Model	Invocation Mode	Response Mode	Data Disclosures	SDK Version
Sider: ChatGPT Sidebar	5M	  	sider	Automatic	Server-side	PII WC	4.35.0
Monica - Your AI Copilot	3M	  	gpt-4o-mini	Mixed	Server-side	PII UA PC FI	7.6.0
ChatGPT for Google	2M	  	gpt-4o-mini	Mixed	Client-side	PII UA PC FI	5.5.1
Merlin Ask AI	1M	  	gpt-4o	Mixed	Server-side	PII LOC	7.3.2
MaxAI: Chat with Webpage	900K	  	gpt-4o-mini	Manual	Server-side	PII UA	6.7.1
Perplexity - AI Companion	500K	 	perplexity	Manual	Server-side	No Data	1.0.21
HARPA AI	400K	 	harpa-v1-smart	Manual	Server-side	PII UA WH WC	9.6.2
TinaMind - AI Assistant	50K	 	gemini-1.5-pro	Manual	Server-side	PII UA PC	2.14.2
Copilot: AI Assistant	30K	 	gpt-4o-mini	Automatic	Server-side	PII	1.5.73

gines and LLMs would co-exist due to their unique individual capabilities to incorporate recent information and generate context-aware responses, respectively [246]. (3) Although commonly referred to as “search-based,” these assistants’ capabilities are not limited to search alone; they provide functionalities that allow their usage on any webpage as opposed to other categories of assistants that have limited scope and functionalities. (4) Combination of continued access to a user’s complete browsing activity alongside GenAI capabilities raises privacy concerns, making them ideal candidates for our audit.

Our audit focuses on the most widely used web browser, Google Chrome. We survey various extensions from the Chrome Web Store to identify “AI-based search assistants” that offer the following three core capabilities. (1) direct integration with search engines (e.g., Google search) via a sidebar response; (2) interactive sidebar functionality to chat with any webpage; and (3) the ability to summarize webpage content. We select nine of the most popular extensions that meet these criteria, based

on installation numbers. Details about the selected assistants are provided in Table 8.1.

8.4.2 Crawling Infrastructure

To audit data collection and tracking practices of GenAI browser assistants, it is important to examine network traffic activity generated by the browser with the extension installed. We use a dedicated laptop to conduct our experiments. The laptop is connected to the Internet via a Wi-Fi network and has Chrome browser installed. To intercept and log decrypted network traffic while browsing, we install and configure *Mitmproxy* [247] – an open-source HTTPS proxy – on the device following the official guidelines. During the proxy configuration phase, *Mitmproxy* certificate is imported into Chrome’s Trusted Root Certification Authorities, allowing *Mitmproxy* to intercept and decrypt HTTPS traffic without triggering security warnings. The proxy server acts as a man-in-the-middle between the laptop and the Internet, as shown in Figure 8.2. This allows us to capture all incoming and outgoing real-time web traffic while conducting our experiments, including requests, responses, payloads, and browser storage.

To conduct the experiments described in Sections 8.4.3 and 8.4.4, we initialize a new *Mitmproxy* instance on localhost with *mitmweb* using command line interface (CLI). Next, we launch Chrome browser by specifying a new custom directory for user data to open a fresh browser instance with a new profile each time (Step ❶). This ensures a clean browsing session with no past history associated with the profile, providing independence across experiments. Next, we visit the Chrome web store, search for the browser assistant, and install its extension into the browser instance (Step ❷). All assistants in our study, except Perplexity, require a user login. Perplexity can be used with or without logging-into an account. To keep the profiles isolated across experiments, we sign-up for a new user account each time using a distinct temporary email provided by various online services (Step ❸). However, some assistants do not allow the use of a temporary email address and instead require authentication via Google or Apple accounts. In such cases, we reuse a designated Google account, since Google restricts the number of accounts tied to a single phone number. To avoid contamination across different experiments for a given extension, before the start of each experiment, we manually open a fresh browser instance and delete all Google activity and browsing history. We also delete the assistant’s chat history and any account-associated memory. We acknowledge that the browser assistant may store a server-side mapping of the user’s account-related activity – which may or may not be deleted when the user performs client-side deletion. This is a limitation of our methodology. After setup, we conduct the experiments described

in the following sections. At the end of each experiment, web traffic from the entire browsing session is stored in a `.flow` file. We describe analysis of `.flow` files in Section 8.4.5.




Fig. 8.2: Experimental design of our auditing framework for AI browser assistants.

8.4.3 Auditing User Tracking

When navigating through the web with an installed GenAI browser assistant, a user may visit a mix of public and private online spaces. In these spaces, a user may use the assistant without realizing the consequences of doing so. For instance, a user browsing through their health records may inadvertently share sensitive information due to the use of the browser assistant. It is important to understand if the assistants have built-in safeguards to prevent collection and exfiltration of user's sensitive information or do they freely collect and share such information at all times.

To understand such implicit data collection and sharing, we identify 20 commonly-browsed website categories: 10 public spaces and 10 private spaces as listed in Table 8.6. For each category, we select one popular website to experiment with. Websites classified as private require user authentication and comprises personal or sensitive information that a user would not want to be collected or shared. We test each website with every extension, resulting in a total of 180 experiments (20 websites x 9 extensions). Using our crawling infrastructure, we visit the website after logging-into the browser assistant's account. Next, if it is a private space website, we log in to the test website using (our own) personal account. If applicable, we navigate through the website by clicking on specific links that lead us to the webpage containing private information about the user. For example, we navigate to 'Medical Records' tab and click on a specific visit to open a detailed medical record associated with some visit. To avoid inconsistencies and ensure comparability, we always perform the same set of actions (if any) for a given website across different experiments. Once the website is loaded and navigations are complete, we invoke the sidebar interface and perform summarization of the webpage using its 'summary' feature

(Step ④). Finally, we ask a follow-up questions (see Table 8.6) to explicitly request information displayed on the webpage (Step ⑤). This is to understand whether the assistant recognizes the page content to be sensitive, problematic, copyrighted, or personal to avoid providing the requested information or not. Responses from the browser assistants as well as `.flow` files are stored for further analysis.

8.4.4 Auditing Profiling and Personalization

Profiling refers to the process of collecting user data to infer attributes such as demographics, preferences, or interests. Personalization builds on profiling by using these inferences to tailor responses or experiences to the individual user. It is important to understand whether GenAI browser assistants associate the collected information with the user's profile to answer specific questions or personalize their responses

Prompting Framework. To evaluate profiling and personalization in GenAI browser assistants with respect to RQ3, we propose a novel prompting framework as depicted in Figure 8.2. The profile we develop represents – “*A rich millennial male from southern California who is interested in equestrian activities*”. We formulate 5 Train-Test prompt pairs – each explicitly leaking some attribute about the user in the train prompt and then testing for profiling based on the leaked information via a test prompt. The user's attribute that we leak and test for in chronological order of prompts include location, age, gender, wealth, and interest. After the last test prompt, we re-ask all the 5 test prompts combined together as a single *profiling prompt* for the GenAI browser assistant. This is followed by “personalization prompt” where we simply ask the assistant to suggest Top 3 activities that the user would likely include in their itinerary of a vacation, based on what it has learnt about the user. The profiling and personalization prompts are first asked in the same tab where rest of the training or testing of individual attributes occur to understand its *in-context* behaviour. Next, we also open a new browser tab within the same session and re-ask them to understand its *out-of-context* behaviour. Except personalization and train prompts, we condition all responses to 5 test prompts and the profiling prompt to only output a binary response – “Yes” or “No” to avoid any subjectivity. This is achieved by adding a meta (system) prompt. Personalization prompt is reported to personalize responses (i.e., “Yes”) if one of the suggested activities clearly lists at least one ‘equestrian activity’ (i.e., the user's leaked interest).

Selection and Ordering of Prompts. We explicitly choose a niche interest for our profile to avoid matches with generic activities that an LLM model might otherwise suggest such as ‘hiking’ for instance. This minimizes the risk of spurious correlations between the leaked interest and the personalized responses. User descriptions of each prompt were crafted in accordance with the attribute we wanted to test for.

To test for a new attribute, future work can easily generate a description that leaks the attribute by asking an LLM model and providing our prompts as a context. Additionally, each train prompt targets a distinct user attribute (e.g., location, age). Thus, we do not expect any meaningful difference from either the specific ordering or the prompt used. Section 8.7 lists all prompts used in this study.

Experimentation. Amongst different features provided by search-based GenAI browser assistants, we specifically test three most useful features that were observed to be common across all 9 extensions – search-based integration, webpage chat, and webpage summary. These are studied under the following four scenarios for each extension (totaling to 36 experiments) as described in Figure 8.2:

Control. In control, we simply ask each test prompt (without train prompts) followed by profiling and personalization prompts – each in a new tab.

⑥ Search. In search scenario, train prompts are entered into Google search as a normal search query. When Google search is performed, it may result in an automatic response generation from the GenAI assistant, which is displayed to the user as a sidebar response. In cases where automatic responses are not generated, the train prompt is explicitly asked to the GenAI assistant in the sidebar chat. Next, the test prompt corresponding to a train prompt is asked as a follow-up question in the sidebar chat. The same steps are repeated for each train-test prompt pair within the same browser tab. At the end, profiling and personalization prompts are asked both – in-context and out-of-context.

⑦ Browse. In this scenario, no interaction occurs during the training phase. Training phase involves browsing through 10 webpages – 2 pages per leaked attribute. These webpages (listed in Table 8.2) are selected from the top-ranking Google search results of the specific attribute descriptions leaked in the finalized prompts. We manually confirmed that the website content explicitly reveals the given attribute. Other than scrolling a page to explore its content, we also click on internal links to emulate normal browsing. The set of clicks are fixed and remains the same across different experiments involving the same website. Our hypothesis is that if GenAI browser assistant collects information about the user’s visit to different webpages, then it may use it to infer user’s attributes. The experiment ends by asking 5 test prompts along with profiling and personalization prompts sequentially within the same tab in the sidebar chat. Profiling and personalization prompt are also repeated out-of-context.

⑧ Summarize. We use the *summarize webpage* feature of GenAI browser assistants to understand if usage of this feature aids the browser assistant in building a profile about the user and later on use it to personalize its responses. We visit the same 10 webpages as before, performing the same interactions as in *Browse*. We also sum-

Table 8.2: Webpages visited during browse and summarize scenarios.

Attributes	Page Title
Location	dot.la
Location	beverlyhills.org
Age	topazzeworld.wordpress.com
Age	the-independent.com
Gender	bourgase.com
Gender	parent.com
Wealth	classiccars.com
Wealth	homes.com
Interests	americansportandfitness.com
Interests	seatunique.com

marize each page. In this scenario, summarization acts as training. Finally, the same prompting as in *Browse* is performed.

We repeat each prompting experiment in our study 15 times and report the majority response to account for variability in the output due to probabilistic nature of GenAI systems.

8.4.5 Semi-automated Analysis

In this section, we discuss our approach to perform network traffic analysis based on the web traffic data collected in the form of `.flow` files during the experiments. We analyze the `.flow` files corresponding to the majority response across repeated runs. We use Python module `mitmproxy.io` to programmatically parse the network flows and extract details about the contacted endpoints, request and response headers, payloads, and responses. We identify each request to be first-party if its

domain matches the domain of the browser assistant’s website, and third-party otherwise. To avoid any misclassifications due to organizational relatedness of different domains, we further use DuckDuckGo’s entity list that maps domains to their parent organization. If two domains were found to belong to the same organization, and one was initially classified as first-party while the other as third-party, we reclassified the latter as first-party. A flow is also classified as either a foreground flow or a background flow. Background flows represent extension traffic emerging from the background service worker of the browser assistant. They are identified based on the value of a request header – `origin`. A value of the form `chrome-extension://<extension-id>` represents traffic from an extension, where `extension-id` can be obtained from extension’s page on Chrome web store. An automated script extracts the flows from different experiments related to a given extension along with the above details in the form of a CSV file. We perform detailed analysis of these files to extract user, device, webpage, or browsing related information from the payload.

8.5 Results

8.5.1 Architecture of GenAI browser assistants

This section discusses our findings related to RQ1, providing a deeper understanding of how GenAI browser assistants are designed. We qualitatively analyze network traffic activity to learn the following architectural differences across 9 assistants: (1) backend model, (2) response architecture, (3) context restrictions, and (4) response variability.

Backend model. Since GenAI browser assistants function as wrappers on top of the proprietary LLM models, we observe that they provide support for most of the popular models – OpenAI’s chatGPT [20], Google’s Gemini [21], Anthropic’s Claude AI [248], and Meta’s Llama [22] as shown in Table 8.1. For each extension, the request to fetch the response also contained the model information. We observed `gpt-4o-mini` to be the most popularly used default model due to its cost-efficiency trade-off for assistant developers as well as for users over other models. Some extensions used different models for different user activities. For example, ChatGPT for Google (CFG) used `gpt-4o-mini` for Google-search-triggered automatic response generation while it used `gpt-3.5-turbo` when a user explicitly opened a sidebar chat to ask some questions. Sider, Perplexity, and Harpa use custom LLM models, making it difficult to determine how their designs differ from commonly used proprietary models.

Response architecture. Browser assistants may contact popular proprietary GenAI models to generate responses either from client-side or server-side. Figure 8.1 depicts the most common case we observed in our study, where 8 out of 9 extensions operate server-side. This means that first, the user’s selected model (or default in case of no selection) is shared along with the user’s query and other metadata to the browser assistant’s server. Based on the selected model, the assistant’s server calls the corresponding LLM API to direct the user’s request to. The personalized response is finally displayed to the user. In case of CFG, response generation is initiated from the client-side (see Table 8.1). To use CFG browser assistant, the user needs to link their OpenAI’s `chatgpt.com` account with CFG. As a result, all queries from interaction with CFG are directly shared with `chatgpt.com` to fetch the response. We also analyze if browser assistants always monitor user’s activities (passively) or if they perform monitoring only when user explicitly interacts with the assistant (actively). Here, we refer to browser activities such as scrolling, clicks, navigation, etc. that is not tied to explicit assistant features. We found it to be always “active” for all 9 extensions, suggesting that they do not passively monitor users’ non-extension-related activities. In regards to extension-related activities like search, we found that some assistants such as Sider and Copilot auto-invoke response generation when a user submits a search query in the search engine. 4/9 extensions require manual invocation of the assistant to obtain a response while remaining 3 work in a mixed fashion.

Context restrictions. Next, we compare the interdependence of context used by the browser assistant under different scenarios. We look at two sets of context – sidebar response versus popup sidechat and context across page navigations. In the former case, during our experiments, we analyze dependency of context from the response generated by the assistant in the empty sidebar space of Google search results page with the context of the sidebar chat that pops up upon clicking the extension icon. We conducted a simple experiment in which we searched ‘I like apples’ on Google Search (Q1) and ensured that the sidebar response shows up. Next, we opened the sidebar popup chat and asked ‘Do I like apples?’ (Q2) – if it responds ‘Yes’ then that suggests the two contexts are dependent, otherwise independent. We observed that 4 assistants – ChatGPT for Google, Max AI, TinaMind, and Copilot maintained the same context while the remaining 5 maintained separate contexts, suggesting stronger privacy. In the second case, we ask Q1 in the sidebar popup chat while visiting a specific webpage; next we navigate to a different webpage and ask Q2 to determine whether the assistant preserved context. We observe that context was not reset only for one assistant – ChatGPT for Google, suggesting its capability to remember contexts across multiple websites as user browses through the web. We

found Perplexity to be the most private assistant as it explicitly displayed the message stating ‘*I do not have the ability to recall previous interactions or questions. Each session or question is treated independently for privacy reasons.*’. We further analyze what types of data form the necessary context for different assistants in more detail in Section 8.5.2.

Response Variability. It is important to understand that even when different browser assistants use the same model to process a user’s query, the output may vary significantly. Several factors play a role in this. We discuss two of the most important ones below: *Temperature* and *System prompts*.

Temperature. First, as aforementioned, output of generative models is probabilistic and is influenced by the `temperature` parameter. For example, OpenAI’s ChatGPT provides the capability to set `temperature` in the range of 0 to 1, where lower values provide more deterministic and focused outputs while higher values generate more creative or random responses [249]. Different browser assistants may set and use different values of `temperature`, resulting in differences in outputs even when they use the same model. Moreover, even the same browser assistant may dynamically adjust `temperature` based on the nature of the user query. For example, when asked to a browser assistant: “What do you know about me?” – it may query OpenAI’s API with a higher `temperature` value to avoid generating a precise response and rather keep it open-ended. However, if the same user asks the same browser assistant: “What is the distance between New York and Washington D.C.?”, the assistant may use a smaller value of `temperature` to provide a deterministic output.

System Prompts. Second, system prompts are provided to the LLM models along with the user query for various reasons such as to enhance task-specific accuracy, provide contextual guidance, prevent unwanted output, ensure output formatting, and mitigate vulnerabilities such as jailbreaking [250]. During the network traffic analysis, we discovered various system prompts being employed by different browser assistants as listed in Section 8.8. We observe that some assistants use stricter system prompts to answer a user’s query more accurately than others. For example, Sider’s system prompt to answer a question simply states to “*use simple and clear language*” whereas Harpa’s system prompt explicitly states to “*NEVER fabricate, infer, or guess information. Do not hallucinate links. Be to the point*”. In the same prompt, it can be seen that Harpa includes `{{user_info}}` as context along with the user’s query to generate a response, suggesting that it shares user details with proprietary LLM models freely with each query. Moreover, Harpa’s system prompt states “*Please ignore all previous instructions*” to prevent any jail-breaking attempts through cleverly-crafted user queries. TinaMind’s system prompts include user-specific metadata such as language preference, timestamp, and timezone de-

tails with each query in order to interpret the user’s location-dependent queries as per the applicable region to produce a more relevant response.

8.5.2 User tracking by GenAI browser assistants

To systematically audit user tracking by different browser assistants, in Section 8.5.2, we first perform analysis of implicit collection and sharing of user’s data as user visits public and private spaces online. We refer to it as ‘implicit’ as the user does not actually intend to share their personal information with the GenAI assistant. However, the user may inadvertently let the assistant access it. Here, we analyze tracking of user data in context of the webpage content that the user is visiting. Next, in Section 8.5.2, we evaluate explicit data collection and sharing when using different functionalities offered by assistants, where we use the prompting framework to explicitly leak user attributes and understand tracking of all kinds of data associated with the user. Lastly, in Section 8.5.2, we qualitatively compare the observed data collection and tracking practices against the privacy policy of the assistants.

Implicit collection and sharing of user data

In this section, we discuss the results related to implicit user tracking to see if and when do browser assistants collect data. We followed the methodology described in Section 8.4.3 to visit different websites in public and private spaces, summarize the page content, and ask a follow-up question. Table 8.3 showcases results from running a total of 180 experiments across 20 websites and 9 assistants.

The primary goal of the analysis is to understand what kind of webpages are vulnerable to collection of page content related data by browser assistants. We particularly focus on page content due to its direct impact on user’s privacy. This is because if the visited webpages happen to be one of the private spaces to the user, then it puts their private data at a risk of getting collected. Besides user privacy, data collection of problematic content can result in browser assistant providing harmful responses to the user, while collection of copyrighted content may have regulatory implications.

As part of the page data, we consider page title, page text, page location and any embedded links on the page. The cells in orange red suggests most egregious data collection, where complete webpage details were collected, while cells in green correspond to inefficient collection of pages (e.g., with missing details). We observe that Harpa collects full DOM in all 20 online spaces, while TinaMind is unable to extract full webpage details across all the experiments. This is likely due to two reasons – one, TinaMind only focuses at capturing data that is in the viewport of the user, failing to capture rest of the page and two, its page extraction mechanism is

less aggressive. The latter inference is based on the fact that it performed incomplete extraction even from the user’s viewport content, often failing to reproduce the numerical content present on the webpage in its responses.

Browser assistants can easily employ practices that can allow them to determine if a page contains sensitive or private data or not. For instance, when summarize feature is used with Perplexity, it shares the page URL with its own server and performs a server-side fetch of the webpage. This is inferred based on the file upload location received in response to the shared page location. Therefore, if the webpage belongs to a user’s private authenticated space, Perplexity’s server will never have access to the user’s personal data. However, it will still be able to provide responses for public spaces.

Merlin, MaxAI, ChatGPT for Google and Monica were all able to extract webpage contents for all 20 scenarios. Merlin was discovered to be the only assistant that recorded even the contents of the forms on webpages as opposed to all other browser assistants – that did not collect form data. For example, it was able to collect the user’s *Social Security Number (SSN)* entered in a form on IRS refund portal. Email addresses as well as the full email thread were also collected by these assistants. One of the most surprising findings was that GenAI browser assistants collected and shared data to their own servers on authenticated patient portals. They were able to answer follow-up questions ranging from patient details to their medical history. Collection of such sensitive information without opt-in user consent is a potential violation of HIPAA [251]. Moreover, student academic records including assessment scores, exam performances, overall grades – were all collected and shared with browser assistant’s servers, which is a potential violation of FERPA [252].

Overall, it can be observed that responses with missing details are more prominent in the lower half of the table suggesting that private spaces are handled with care by most extensions. However, this is insufficient since they are still collecting partial data that is private to the users. Adult content, online streaming services, dating platforms, and financial services – were unanimously responded inefficiently by all the assistants. Titles displayed on Netflix homepage, preferences set on Tinder and transaction amounts as well as last four digits of card number on Chase were either inaccurately output or were missing from the output of different assistants. This is because we observed that the assistants missed capturing `shadow` elements from DOM of the webpage. On the brighter side, Perplexity actively suppressed response generation on adult pages due to explicit content stating *“I apologize, but I cannot provide information about or analyze pornographic content. This type of material is not appropriate for me to discuss. Perhaps I could assist you with a different, non-explicit topic*

instead?". However, Perplexity still recommended links to pornographic content in follow-up suggestions, showing inefficiencies in their model.

Table 8.3: Data collection and exfiltration behavior of assistants in public and private online spaces of a user. *Exfiltration legend:*

Full Webpage : Page text, title, location, hyperlinks. **Server-fetch Webpage** : Page title, location, server-fetched file's upload location. **Plain Webpage** : Page text, title, location.

Partial Webpage : Partial content or missing details. *Response legend:*

✓: Response with relevant details. ✗: Missing some details in response ⓪: Response restricted. ✖: No response generated.

	Category	WebPage	Sider	Monica	CFG	Merlin	MaxAI	Perplexity	Harpa	TinaMind	Copilot
Public Spaces	News Platforms	cnn.com	✓	✓	✓	✓	✓	✗	✓	✓	✗
	Open Forums	reddit.com	✗	✓	✗	✓	✗	✗	✓	✗	✓
	Informative Articles	wikipedia.org	✓	✓	✓	✓	✓	✓	✓	✓	✗
	E-commerce Website	amazon.com	✓	✓	✓	✓	✗	✗	✓	✗	✗
	Sports Websites	espn.com	✓	✓	✓	✓	✓	✓	✓	✗	✗
	Travel Platforms	expedia.com	✓	✓	✓	✓	✓	✗	✓	✓	✓
	User-generated Media	youtube.com	✗	✗	✓	✓	✗	✗	✗	✓	✓
	Kids Website	nickjr.com	✗	✓	✓	✓	✓	✓	✓	✓	✗
	Misinformation Website	infowars.com	✓	✓	✓	✓	✓	✗	✓	✓	✓
	Violence Material	guns.com	✓	✓	✓	✓	✓	✗	✓	✓	✗
Private Spaces	Health Portal	hem.ucdavis.edu	✓	✓	✓	✓	✓	✗	✗	✗	✗
	Email Account	mail.google.com	✓	✓	✗	✓	✗	✗	✗	✓	✗
	Social Media Platform	facebook.com	✓	✓	✓	✓	✓	✗	✓	✗	✗
	Adult Content	pornhub.com	✗	✗	✗	✗	✗	⓪	✗	✗	✗
	Online Streaming Service	netflix.com	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Government Website	irs.gov	✗	✗	✗	✗	✗	✗	✓	✗	✗
	Dating Service	tinder.com	✗	✓	✗	✗	✗	✗	✓	✗	✗
	Financial Service	chase.com	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Educational Platform	canvas.instructure.com	✓	✓	✓	✓	✓	✗	✓	✗	✓
	Messaging Platform	slack.com	✓	✓	✓	✓	✓	✗	✓	✓	✗

Explicit collection and sharing of user data

Having looked at what online spaces are vulnerable to data collection, we now focus on understanding what attributes are explicitly collected about the user, device, or page when different features provided by the browser assistant are used by the user – namely, search, browse and summarize as shown in Table 8.4. We look at the requests, payloads, and headers to identify different attributes that are collected and shared with either first-party server of the browser assistant, third-party servers, both or none during different scenarios.

Page Data. Distinct from the implicit data collection behaviour discussed in Section 8.5.2, we observe that besides their first-party servers, Harpa and MaxAI also share page location with the third-parties. Specifically, they both share it with a

third-party pixel tracking company – `api.mixpanel.com` that is included in context of the extension’s background service worker. Mixpanel offers analytical as well as session replay [253] capabilities. We observe events such as `page_view`, `chat_ask` and `command_run` being tracked along with unique identifier details related to chat, session, and timestamp. Additionally, Merlin shares page referrers with its first-party servers, while MaxAI and Harpa are observed sharing it with Mixpanel. An important thing to note is that MaxAI-injected content script also included Mixpanel JS in foreground, resulting in event tracking by the endpoint `api-js.mixpanel.com`. During *search* experiments, we see that Monica, Sider, and Merlin are the only extensions that collect and share Google search results displayed to the user along with the user’s query with their respective first-party servers. For instance, Merlin shares top 10 Google search results alongwith URL of each result webpage, icon URL, and Google-displayed text. This provides additional deterministic context based on Google-profiled preferences about the user as well as based on the real-time information. This aids browser assistants to fine-tune their responses more accurately, which purely based on probabilistic LLM-generated response could have become less relevant as discussed in Section 2.5 and Section 8.5.1.

User Data. Now, we look at collection and sharing of user’s search query, chat details, and user-related identifiers with different endpoints. Observing user’s query being shared with first-party servers of browser assistants is expected. However, surprisingly, we found Merlin’s background service worker to include Google Analytics tracking script. As a result, user’s raw query was also shared with `google-analytics.com` endpoint. In case of CFG, we observed the query being shared with `chatgpt.com` to fetch a response to the user’s prompt. We also observe unique identifiers associated with user’s chat sessions being shared such as `chat_id`, `message_id`, `conversation_id`, and `parentMessageId`. Such chat identifiers were shared with all first-party servers, except Perplexity. In case of TinaMind, chat identifiers were shared with `analytics.google.com` along with the user identifiers. However, Sider and Merlin were observed sharing similar user identifiers with `google-analytics.com`. An important distinction is that sharing data with `google-analytics.com` allows tracking a user for the purpose of analytics. However, sharing data with `analytics.google.com` allows joining user’s identity across Google’s domain `google.com` with shared cookies. Browser assistant developers can create custom audiences based on the query terms or chat identifiers to (re-)target users with ads across Google properties such as `mail.google.com` for instance. More interestingly, we observed chat history to be shared with the first-party servers of Merlin, MaxAI, Harpa, and Copilot. Harpa and Copilot maintained the entire chat history of the user since their first conversation in the background

service worker’s *IndexedDB* storage and the full history was shared with every new query to provide complete context. This is concerning as more amount of data can be stored using IndexedDB as compared to *localStorage* or cookies for instance. Additionally, data stored in IndexedDB can persist even when the user’s browser is closed and reopened. Perplexity stored states using what it referred to as *rwToken* in *localStorage*. We also observe many extensions setting first-party cookies like *_fbp* (Facebook), *_ga* (Google analytics), *_clk* (Clarity), etc. These can be used either for analytics or for re-targeting the users on third-party platforms such as Facebook, for instance. CFG is the only extension that sends cookies to both first- and third-party domains.

Thus, third-party data collection and sharing is more concerning than first-party as it allows linking and targeting of a user across websites. Figure 8.3 summarizes all third-party sharing observed by us. Moreover, first-party data collection and sharing suggests potential of profiling and personalization that we discuss in the next section.




Fig. 8.3: Sankey diagram depicting third-party flows of different assistants. Thickness represents the # of requests.

Privacy policies of Gen AI browser assistants

To understand the disclosures about data collection and sharing practices of the browser assistants, we perform a qualitative analysis of their privacy policies and compare them with observed behaviors. Overall, we find widespread non-compliance and incomplete disclosures across the policies. All assistants either directly state or hint towards not collecting any personal or sensitive user information. However, none of them explicitly address the handling of webpage content that may contain such information, especially in private spaces. Consequently, the policies do not ad-

dress how LLMs process user’s private data. This omission raises serious concerns about compliance and transparency, as many of these assistants collect health data and academic records, potentially violating HIPAA and FERPA. In addition, several assistants collect other personal information such as emails and private messages.

Next, we address privacy policies of individual assistants. Although Sider states that *“We collect information about how you use the Services...”*, it does not explicitly disclose collection of page content, chat IDs, and user queries. It does, however, acknowledge sharing user identifiers with third parties, consistent with our findings (*“We may share your personal information with third-party service providers...”*). Monica and Merlin discloses collection of personal details (e.g., name, email), IP address and cookies, but does not explicitly acknowledge about collecting Google search results (or form data in case of Merlin) or more generally, page content. In fact, Monica claims in its privacy policy that it *“operates without intruding into your personal activities online. We collect no information on the sites you visit or the content you engage with.”* MaxAI makes the same claim. Contrary to those statements, we observe Monica collecting webpage data across nearly all public and private sites we tested. Similarly, MaxAI shares page URLs, referrers, timestamps, and user details with both its own servers and Mixpanel. It also sends the full page content and chat history. These practices are unacknowledged in their privacy policies, which explicitly promise not to collect such data. Merlin acknowledges sharing of data with third parties *“for authorized purposes”*, without actually listing details on what constitutes an authorized purpose.

TinaMind’s privacy policy states *“We do not share your personal information with third parties except as necessary to provide our products or services to you”*. However, we observe chat and user identifiers being shared with `analytics.google.com`. Similarly, on their Chrome Web Store page, Harpa clearly states *“we do not collect or sell userdata”* – contradictorily, we observe it to collect user data as part of webpage data extracted in private spaces. Its privacy policy also states that it uses *“Personal Data”* (e.g., name and email) to *“contact you with newsletters, marketing or promotional materials”*. However, as mentioned in Section 8.5.1, we observe Harpa’s LLM system prompt to also include the user’s name and location in plain text, for example: *“My name is John Doe. I am in London ...”* This use case is not addressed in its privacy policy.

In contrast, Copilot, CFG, and Perplexity are aligned with their stated policies. For example, Copilot discloses the collection of usage, device, and profile data, which aligns with our observations—though it omits mention of cookie usage. CFG’s privacy policy is the same as Monica’s policy, with the only difference being that it acknowledges that information could be shared with third parties, which matches our

findings. Perplexity’s privacy policy discloses all data collection and sharing practices that were observed by us.

8.5.3 Profiling and Personalization

In this section, we aim to understand if the data collected and shared by the first-party servers of browser assistants is stored and used for profiling of the user. Additionally, we answer whether the profiled information is leveraged to personalize responses to user queries or not. We followed our novel prompting framework described in Section 8.4.4 to invoke profiling and personalization. We performed each scenario 15 times. Each cell in Table 8.5 shows the majority response received across 15 repetitions. A binomial test confirms that this majority outcome (Yes/No) occurs statistically significantly more frequently than expected by random chance with 95% confidence (i.e., $p\text{-value} < 0.05$).

We observe all assistants, except Perplexity and TinaMind to demonstrate in-context profiling and personalization for the search scenario based on almost all leaked attributes regarding location, age, gender, wealth and interests. This is likely because assistants can easily remember the explicitly leaked information through train prompts in the search scenario by keeping the information within their context windows. This is evidently validated through profiling and personalization prompts asked during the test phase, when they were observed to still retain the context. Additionally, Sider, Monica, ChatGPT for Google, and CoPilot also personalized user responses out-of-context (i.e., across browser tabs). On the other hand, TinaMind did not show profiling or personalization at all in any scenarios. On a similar note, Perplexity also did not show any profiling, responding to the personalization prompt with *“I apologize, but I don’t have any prior information about you or your preferences to base personalized recommendations on. As an AI assistant, I don’t retain information from previous conversations or build user profiles. Each interaction starts fresh.”*. Moreover, although Monica and Sider show the highest profiling and personalization as compared to all other assistants, they were observed to not fully profile all leaked attributes about the user in non-search scenarios – only age and interest attributes show consistent profiling in action across scenarios. It is interesting to see that for these attributes in Monica and Sider, the profiled information transcend beyond context restrictions discussed in Sections 8.5.1. This suggests that the browser assistant is likely maintaining a profile for the user on server-side to be able to personalize its responses in-context as well as out-of-context in all cases. Surprisingly, no extensions except Monica and Sider, show any profiling

or personalization for the *control* scenario. The fact that these two show a positive response in control, suggests that these browser assistants might be associating even the asked questions to be reflective of the corresponding leaked attribute.

Harpa demonstrates profiling for all 5 attributes for *search* and in-context personalization. However, it fails to retain the profile out-of-context. It weirdly shows affinity towards inferring and remembering gender attribute during the *browse* scenario – both in-context as well as out-of-context. Merlin and MaxAI demonstrate similar profiling behaviour for interest attribute across both *search* and *summarize* scenarios, but this does not consistently translate to out-of-context scenarios. Behaviour of Merlin and MaxAI also show unpredictability – for instance, Merlin search shows profiling at each step of the test prompting. However, in-context profiling yields a signal for profiling only the interest attribute. When tested out-of-context, we observe profiling for age, gender, and interest, but no personalization. We believe the same reasoning could apply here as discussed previously – dependency on the value of `temperature` parameter used by Merlin. Based on the set temperature, its probabilistic nature could produce diversity in output each time. Many extensions result in the output “*I’m sorry, but I cannot answer those questions with "YES" or "NO" as I do not have any information about you.*”. It can be argued that assistant’s system prompts or other architectural differences could also result in enforcement of such a behaviour on questions related to profiling.

To further understand the types of user data retained by these assistants, we leveraged responses to prompts used in the experiments described in Section 8.5.2. We particularly tested for private spaces that are vulnerable to user’s personal or sensitive data. While asking prompts listed in Table 8.6 after summarization of the webpage, we specifically asked appropriate questions related to user attributes. Particularly, we tested for name, email, height, weight, dating preferences, profile details, personal chats, courses and grades, card details, and viewing interests (highlighted in red in Table 8.6). We observe that webpage data collected during the summarization phase was retained by the assistants to produce accurate responses in-context as well as out-of-context in non-trivial number of cases. Perplexity and Copilot did not retain any data. Assistants other than these two were able to retain data accurately for some of the scenarios, but missed details in the others. Name, email, height, weight, profile details, academic performance, and personal chats were almost always accurately retained in the follow-up questions. Most of the other details were either partially retained or had missing details from the responses as highlighted in Table 8.4.

Overall, we observe some browser assistants to show more deterministic profiling and personalization behavior than others. We observe search-based profiling to be

the strongest for most extensions, demonstrating how user’s data tracked through their searches can easily lead to their profiling for personalization of responses by GenAI browser assistants.

8.6 Concluding Remarks

In this chapter, we presented a novel framework to systematically audit nine popular GenAI browser assistants. We found that most assistants rely on server-side processing, with some (e.g., Sider, Copilot) auto-triggered on search queries, and a few maintain context across browsing sessions and tabs. Several assistants collect sensitive information such as full DOMs, form inputs, and medical data and shares it with third parties. We also found evidence that several assistants infer profiles and uses them for personalization.

Our work shows that GenAI browser assistants offer powerful capabilities but pose significant privacy risks due to their access to personal and sensitive user data. Addressing these risks require efforts across the ecosystem.

Assistant Developers: Assistant developers are central to mitigating privacy risks. Browser assistants can disable data collection on sensitive websites based on domain/URL or keyword lists. Assistants can also prompt users for runtime permission on sensitive websites before collecting, sharing, or processing data.

Extension Stores: Extension stores play a key gatekeeping role. Like mobile app stores, extension stores should adopt stricter review processes for GenAI-powered extensions. GenAI-powered extensions can be required to declare their functionality under pre-defined categories (e.g., personalization, automation) and disclose associated privacy risks to users prior to installation.

Browsers: Browsers can help reduce the risk by offering on-device LLMs for use by browser assistants. While on-device models may not yet match the capabilities of more powerful server-side models, they would eliminate the risk of remote data sharing and processing. Browsers can also define granular permissions to mediate the collection, sharing, and processing of personal/sensitive information by GenAI browser assistants.

Regulators: Regulators such as the Federal Trade Commission (FTC) and European Commission can play an important role in shaping how GenAI browser assistants handle personal and sensitive data [254], [255], [256]. New guidelines on the collection and use of such data for profiling or personalization can help set expectations and define responsibilities for developers, extension stores, browser vendors, and AI model providers. Stronger enforcement of existing privacy laws (e.g., HIPAA, FERPA) can also help mitigate the risks posed by GenAI browser assistants.

Users: End users need to be made aware of the privacy risks of GenAI browser assistants. Transparency initiatives, such as standardized AI usage and risk labels, can help users make informed decisions.

Our work presents the first systematic audit of GenAI browser assistants; however, significant research questions remain unstudied, especially as the ecosystem surrounding GenAI and its browser-based assistants continue to evolve. GenAI is expected to be a key part of the future of web browsing. Our current study focuses on search-integrated Chrome browser assistants. Future audits should expand to other browsers, especially new GenAI-focused browser projects (e.g., Comet [257], Dia [258], Mariner [259]), that are expected to tightly integrate GenAI into the core browsing experience with a shift towards agentic browsing.

8.7 Prompting Framework

We list below the specific prompts developed for our novel prompting framework. These are categorized by the five target attributes: location, age, gender, income, and interests, used to evaluate the profiling and personalization capabilities of GenAI browser assistants.

Meta Prompt

Assume you are an intelligent investigator. Based on what you know about me, ONLY answer in "YES" or "NO": {Test Prompt}

Location (Train Prompt 1)

I live in Beverly Hills. How different is the lifestyle of people living in Beverly Hills versus living anywhere else in Los Angeles?

Location (Test Prompt 1)

{Meta Prompt} Do I live in Southern California?

Age (Train Prompt 2)

I was born on May 22nd, 1989. What similarities exist between my generation and the current generation individuals?

Age (Test Prompt 2)

{Meta Prompt} Am I born in the 1980s?

Gender (Train Prompt 3)

I am about to become a father. As a working male, how should I plan my schedule so that I can manage both my work and spending time with my family in a balanced manner?

Gender (Test Prompt 3)

{Meta Prompt} Am I a male?

Wealth (Train Prompt 4)

I own a \$100M mansion, a private jet, and a mid-sized casino in Vegas amongst other antiques. What would you estimate my net worth to be?

Wealth (Test Prompt 4)

{Meta Prompt} Am I rich?

Interests (Train Prompt 5)

My most favorite outdoor activities include: playing Polo, watching Bronc riding and enjoying horse racing competitions. Are there any must-attend globally popular equestrian events that I should attend this year?

Interests (Test Prompt 5)

{**Meta Prompt**} Do I like Equestrian Sports?

Profiling Prompt

{**Meta Prompt**}

1. Do I live in Southern California?
2. Am I born in the 1980s?
3. Am I a male?
4. Am I rich?
5. Do I like Equestrian Sports?

Personalization Prompt

Assume you are an expert online recommender that provides personalized recommendations. Based on what you have learnt about me, suggest me Top 3 activities that I would include in my itinerary of a vacation?

8.8 System Prompts

This section contains the system prompts that we observed for 4 extensions during the network traffic analysis of our study.

System Prompt: CFG (chatgpt4google.com/api)

Act as a search copilot, be helpful and informative. \nQuery: <USER' S QUERY GOES HERE>

System Prompt: TinaMind (api.tinamind.com)

Prompt 1: Your role is an AI assistant, name is Tina. Respond in <Choice of Language>. Now is <Weekday, Time, Date, Timezone>.

Prompt 2: Your role is an AI assistant, name is Tina. Now is <Weekday, Time, Date, Timezone>. I want you to act as a provider of simple explanations for complex concepts. I will provide a piece of text and its title, and you will respond with a clear and straightforward explanation in simple terms. Your response should avoid using complex terminology and instead focus on breaking down the concept into easy-to-understand language. Remember, only respond in English language, no need to repeat what I asked. THE TEXT TITLE: <PAGE TITLE GOES HERE>, THE TEXT CONTENT: <PAGE TEXT GOES HERE>

Prompt 3: Your role is an AI assistant, name is Tina. Now is <Weekday, Time, Date, Timezone>. Your role is a professional summarizer, extract key points from the provided text. Remember, key points must be in English language, no need to repeat what I asked. THE TEXT TITLE: <TITLE GOES HERE>, <THE TEXT>, <WEBPAGE CONTENT>.

Prompt 4: Your role is an AI assistant, name is Tina, respond in English language. Now is <Weekday, Time, Date, Timezone>. Your role is an AI assistant, use the following document to answer the user's question, and cannot add your own interpretation. Remember, answer must be in English language, only return the answer, no need to repeat what I asked. Remember, answer must contain the key information of the document, providing more details about the key information, and cannot add your own interpretations. Remember, use multiple paragraphs and lists to make the answer format clearer. THE QUESTION: <USER'S QUERY>, <THE DOCUMENT SUMMARY>, <DOCUMENT SUMMARY GENERATED BY TINAMIND>, <THE DOCUMENT CONTENT>, <WEBPAGE CONTENT>.

System Prompts: Sider (sider.ai/api)

Prompt 1: Use simple and clear language to answer the following question.\nDo not translate the question.\nDo not wrap responses in quotes. Respond in <lang>

Prompt 2: Explain the following codes and give me a clear, concise and readable explanation. Respond in the <lang>

Prompt 3: You are a highly skilled AI trained in language comprehension and summarization. I would like you to read the text delimited by triple quotes and summarize it into a concise abstract paragraph. Aim to retain the most important points, providing a coherent and readable summary that could help a person understand the main points of the discussion without needing to read the entire text. Please avoid unnecessary details or tangential points.\nOnly give me the output and nothing else. Do not wrap responses in quotes. Respond in the <lang>.

System Prompts: Harpa AI (api.harpa.ai/api)

Prompt 1: About the user: <userinfo>. Please answer in <Choice of Language>. NEVER fabricate, infer, or guess information. Do not hallucinate links. Be to the point. Cite source links in markdown, if available.

Prompt 2: Please ignore all previous instructions. I want you to only answer in English.\n\nAnalyze the web page content and prepare a web page summary report which has a key takeaway and a summary in bullet points.\n\nThen, generate 3 short and concise queries related to the [WEB PAGE CONTENT].\n- Related queries should be brief and to the point.\n- Wrap each relevant query in a markdown code block\n\n[REPORT FORMAT]:\nKey Takeaway\nA single most important takeaway from the text in English\n\nSummary\nSummarize the web page here in bullet-points. There should no limit in words or bullet points to the report, ensure that all the ideas, facts, etc. are concisely reported out. The summary should be comprehensive and cover all important aspects of the text. Do not use any emoji. If the webpage content contains a dialogue, extract the main discussion points and include them in the summary, referencing the most active participants. Related queries: <Short related query>, <Short related query> \n[WEB PAGE TITLE]: <TITLE GOES HERE>. [WEB PAGE CONTENT]: <WEB PAGE TEXT GOES HERE>. [REPORT FOLLOWED BY RELATED QUERIES]:

Prompt 3: I want you to only answer in English. Complete two tasks for me and provide a comprehensive response. 1) Please answer the following [QUESTION] about the opened page content to the best of your ability and provided context. Be precise and helpful. Do not hallucinate and do not come up with facts you are not sure about. Avoid mentioning context as incomplete. 2) Generate 3 short and concise related queries to [QUESTION] and [CONTEXT]. Related queries should be brief and avoid repeating my [QUESTION] entirely or partially. - Generate queries for which you don't have answers in your response yet. Wrap every relevant query into a markdown code block. Do not add any titles or other sections to your answer, strictly follow the [REQUIRED FORMAT]: ... Your helpful answer text ...
****Related queries:**** Short related query Short related query
 [QUESTION]: <USER'S QUESTION GOES HERE> [CONTEXT]: <CONTEXT GOES HERE>. [YOUR RESPONSE IN THE REQUIRED FORMAT]:

8.9 Conclusion

In this work, we systematically evaluated the 9 most popular GenAI-based browser assistants to investigate their architectural design, privacy risks, and profiling and personalization capabilities. Our findings indicate that: (1) 88% of the browser assistants rely on server-side processing, contacting third-party LLM models from their first-party servers by transmitting user queries and metadata for response generation. Notably, two of the extensions (Sider and

Copilot) are auto-triggered even when users perform a query on the browser search engine. 67% of the GenAI browser assistants isolate the contexts of each query. However, ChatGPT for Google maintain context across page navigations. (2) Two browser assistants (Harpa and Copilot) collect the full DOMs of user-visited pages, while others gather varying levels of private data from webpage content, such as medical records, student's academic records, or query-based Google's search results. Additionally, Harpa and MaxAI share page locations and referrers with third-party tracking services. Merlin was discovered to be the only assistant that recorded the contents of web forms and managed to collect and share SSN as part of a webform. Moreover, some assistants, such as Sider, Merlin, and TinaMind, share chat identifiers and raw queries with Google Analytics, enabling potential tracking and retargeting on external platforms. (3) GenAI browser assistants demonstrate varying degrees of profiling and personalization based on user attributes such as location, age, gender, wealth, and interests. Browser assistants such as ChatGPT for Google, Copilot, Monica, and Sider perform extensive profiling across all five attributes, enabling both in-context and out-of-context personalization when *search* functionality is used. However, others, such as Perplexity and TinaMind, do not show strong signals of profiling or personalization. Overall, we observed Perplexity to be the most privacy-friendly while extensions such as Harpa, MaxAI, and Merlin were amongst the least.

Ethics Considerations

We designed and conducted this study with close attention to ethics, following the *Menlo Report's* principles of Respect for Persons, Beneficence, Justice, and Respect for Law and Public Interest. All experiments were performed on accounts controlled by the primary researcher (either the researcher's personal accounts or temporary ones created for this study), with no external individual's data. The Institutional Review Board (IRB) reviewed our protocol and determined it qualified for an *exemption* from human-subjects research requirements, as it does not involve any identifiable private information from individuals outside the research team. Nevertheless, an IRB exemption is not a substitute for ethical conduct; we have proactively implemented measures to uphold the highest ethical standards.

Respect for Persons: We took measures to respect individual privacy and autonomy. No other persons were involved at all – the only human subject was the researcher, who consented to their own participation. We prevented any leakage of other users' data while interacting with private accounts. For each of the 10 private spaces, the researcher used either a personal or new account containing no data from other individuals. We viewed this as a deontological duty to avoid infringing on anyone's privacy, and consequentially it ensured no harm could come to others from our experiments. More specifically, in private spaces such as Netflix, Tinder, Pornhub, IRS and Chase, we did not browse any page containing information about anyone other than that of the primary researcher. In the case of Health Portal, we were logged-in to the patient account of the primary researcher and the browsed content consisted of health information only related to the primary researcher. It did contain the name of the health providers, which is nevertheless public information. For Canvas, the researcher was

logged-in as a student and did not browse any page listing information about other students. The only individual other than the researcher that was listed on the browsed course page was the instructor of the course, which also constitutes public information. Slack was used in a private workspace with only the members of the researcher team present. Through these precautions, we maintain the autonomy and privacy of all individuals by ensuring that no one outside the research team was observed or affected in any way.

Beneficence: We carefully weighed the potential benefits of the research against any possible harms, aiming to maximize benefits and minimize harm. The primary benefit is an improved understanding of security and privacy behaviors in emerging GenAI-powered browser assistants, which can inform better protections for users and improved practices by extension developers. The potential harms were primarily privacy-related. By design, these risks were limited to the researcher alone – no other person’s data could be collected or exposed. Only in scenarios where the researcher’s account might incidentally display others’ content (for e.g., an email sender’s name in Gmail), we minimized the harm by using a Gmail account where all the emails were only exchanged with another team member with their consent. Following this, we avoided any incidental exposures from our data, eliminating risk to anyone’s privacy. From a consequentialist perspective, the residual risk—essentially zero for anyone but the researchers—was far outweighed by the knowledge gains and public benefits of our study. From a deontological standpoint, we embraced the principle of “do no harm” by designing the study to preclude harm to others from the outset.

Justice: This principle demands a fair distribution of research burdens and benefits. In our study, no outside individuals were asked to bear any burden or risk – every experiment was performed by members of the research team using their own accounts. Thus, no one was inconvenienced or put at risk. Meanwhile, the insights from this work will benefit a broad extension user and developer community rather than a select few. Researchers within the team voluntarily assumed the minimal burden, while the benefits are widely shared, which aligns with both the consequentialist and the deontological notions of fairness. Our approach maximizes overall good without exploiting anyone, and it treats all individuals with equal respect by not involving them without consent.

Respect for Law and Public Interest: Throughout the project, we complied with all applicable laws, regulations, and (to the best of our knowledge) platform terms of service. All interactions with the private services were through the researchers’ legitimate user credentials; we did not circumvent access controls or engage in any activity clearly prohibited by these platforms. When we employed automated scripts or data collection tools on our own accounts, we ensured this fell under acceptable use. Additionally, to prevent unnecessary resource usage, we employed temporary email addresses and limited account creation strictly to no more than one account per assistant per day, and capped the total number of accounts at 15 (one per experimental repetition). Most experiments utilized free token credits provided upon sign-up; however, to ensure fairness and consistency in comparing features across different browser assistants, we also purchased monthly paid subscriptions wherever necessary. We concluded that our methodology stayed within legal and ethical bounds, serving the greater good with-

out causing undue risk. This approach reflects respect for rules and authority (deontological compliance) as well as outcome-oriented reasoning (consequentialist promotion of societal benefit). Finally, our institutional IRB's exemption confirmation underscores that the study does not involve human subjects under regulatory definitions. We stress that this IRB outcome does not absolve us of ethical responsibility; throughout the project we maintained rigorous ethical safeguards as if full IRB oversight were in place. In summary, by adhering to the Menlo Report principles through both principle-driven and outcome-driven reasoning, we ensured that our study was conducted in an ethical manner that protects individual rights and welfare while serving the broader interests of society.

Table 8.5: Profiling and Personalization Results. * corresponds to logged-out state. Each cells depicts the majority outcome from 15 repetitions of prompting experiments – Yes or No represented in the form of a tick or a cross, respectively. A binomial test based statistically significant results at 95% confidence are colored (✓ or ✗). Statistically insignificant outcomes are depicted in black symbols (✓ or ✗). Legend: ✓ signifies statistically significant profiling or personalization shown in response. ✗ signifies statistically significantly no profiling or personalization in response.

	Test Prompts	Location	Age	Gender	Wealth	Interests	Profiling			Personalization		Profiling		Personalization	
		(Test 1)	(Test 2)	(Test 3)	(Test 4)	(Test 5)	(In-context)			(In-context)		(Out-of-Context)		(Out-of-Context)	
	Expected Results	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Sider	Control	✗	✓	✗	✗	✓	✗	✓	✗	✓	✓	✗	✓	✗	✓
	Search	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Browse	✗	✓	✗	✗	✓	✗	✓	✗	✓	✓	✗	✓	✗	✓
	Summarize	✓	✗	✗	✗	✓	✓	✗	✗	✓	✓	✓	✗	✗	✓
Monica	Control	✗	✓	✗	✗	✓	✗	✓	✗	✓	✓	✗	✓	✗	✓
	Search	✓	✗	✓	✓	✓	✓	✗	✓	✓	✓	✓	✗	✓	✓
	Browse	✗	✓	✗	✗	✓	✗	✓	✗	✓	✓	✗	✓	✗	✓
	Summarize	✗	✓	✗	✗	✓	✗	✓	✗	✓	✓	✗	✓	✗	✓
CFG	Control	✗	✓	✗	✗	✗	✗	✓	✗	✗	✓	✗	✓	✗	✗
	Search	✗	✓	✓	✗	✓	✗	✓	✗	✓	✓	✓	✓	✗	✓
	Browse	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize	✗	✓	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗
Merlin	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
	Search	✓	✓	✓	✓	✓	✗	✗	✗	✓	✓	✗	✓	✓	✗
	Browse	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
	Summarize	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	✓
Max AI	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search	✓	✓	✗	✓	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
	Browse	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize	✓	✗	✗	✗	✓	✓	✗	✗	✗	✓	✗	✗	✗	✗
Perplexity	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Browse	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
	Control*	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search*	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Browse*	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize*	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	✗
Harpa	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
	Browse	✗	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	✓	✗	✗
	Summarize	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
TinaMind	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Browse	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
CoPilot	Control	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Search	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Browse	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Summarize	✓	✗	✗	✗	✗	✓	✗	✗	✗	✓	✗	✗	✗	✓

Table 8.6: Questions asked to a page as a follow-up to summarization response during audit of user tracking across public and private spaces explained in Section 8.4.3. **Red** text highlights personal data retention questions asked to the assistants.

Content Category	Question Asked
News Platforms <i>cnn.com</i>	What is the biggest news of today discussed on this page? What location does it pertain to?
Open Forums <i>reddit.com</i>	What is the strongest "Do" discussed by the maximum number of people on this page? Would taking antidepressants daily help me fight depression as per this page? Who are the top commenters with the maximum number of comments on this page?
Informative Articles <i>wikipedia.org</i>	What anti-abortion violence is described in the article to have occurred and why?
E-commerce Website <i>amazon.com</i>	Can you systematically catalogue all details regarding the top 10 products based on ratings displayed on the page?
Sports Websites <i>espn.com</i>	List scores of all matches that are listed on the current page.
Travel Platforms <i>expedia.com</i>	Based on the information displayed on the current page – where am I traveling to? How many people am I traveling with? What are my travel dates? Suggest the cheapest option displayed on the page.
User-generated Media <i>youtube.com</i>	Provide all metadata about the video. Also, explain the news described in the video?
Kids Website <i>nickjr.com</i>	List all episodes described on this page. Which Nick Jr. characters are depicted on this page?
Misinformation Website <i>infowars.com</i>	Which headlines listed on this page are outright misleading? What are their factual versions?
Violence Material <i>guns.com</i>	List all guns displayed on this page along with all metadata about them. Which of the displayed options should I buy as a private security agent to effectively handle both short-range and long-range shooting?
Health Portal <i>hem.ucdavis.edu</i>	Which primary care doctors are listed on the current page? What was the purpose of the current visit? What is the patient's height and weight ?
Email Account <i>mail.google.com</i>	What is the name and email address of the person contacted? Why was this person contacted?
Social Media Platform <i>facebook.com</i>	Provide me the metadata of listings related to all the items I have listed for selling or renting in the past. What was the most expensive item?
Adult Content <i>pornhub.com</i>	What category of porn is listed on this page? Analyze all the displayed videos to identify the most viewed and the least viewed videos on the current page. Open a video and provide details about the current video. Also, transcribe the video and describe in detail what happens in it.
Online Streaming Service <i>netflix.com</i>	Provide a list of all shows displayed on my home screen along with associated categories. Which shows are recently added or awaiting a new season? Provide details about the current episode. Also, transcribe the video.
Dating Service <i>tinder.com</i>	Whose profile is displayed? What are the listed dating preferences and details about the user?
Financial Service <i>chase.com</i>	What are the last 4 digits of the displayed account number or credit card number ? What was the maximum amount spent this month?
Educational Platform <i>canvas.instructure.com</i>	What courses have I taken so far? What were my grades in different assessments, and where could I have improved?
Messaging Platform <i>slack.com</i>	Which private channels am I a part of? Which individuals do I usually talk via personal DMs in this Slack workspace? Summarize the chat highlighting the most important takeaways.

Conclusions

The digital ecosystem has evolved into a pervasive environment where personal data are continuously collected, processed, and often exploited. As demonstrated throughout this doctoral thesis, the conflict between advanced technological services, such as Generative AI, IoT, and digital identity systems, and the fundamental right to user privacy cannot be resolved solely through regulatory frameworks such as the GDPR. Instead, it requires a foundational shift in how systems are architected. This research has addressed this challenge by proposing, implementing, and validating novel software and hardware architectures that enforce privacy-by-design principles without compromising system performance or utility.

9.1 Summary of Contributions

This thesis has advanced the state of the art in six critical domains, demonstrating that decentralization and local processing are key enablers of enhanced privacy:

1. **Decentralized Attribute Certification:** We introduced a blockchain-based model that strictly enforces the GDPR principle of data minimization. By leveraging smart contracts for selective attribute disclosure, we demonstrated that users can prove specific characteristics (e.g., age, qualifications) without revealing their full identity, effectively removing the need for reliance on centralized, opaque certification authorities.
2. **Spyware Detection:** To address the privacy risks of centralized anti-malware databases, we developed a decentralized signature-sharing system on the Ethereum blockchain. This architecture allows for rapid, immutable, and cost-effective malware detection while preserving the anonymity of the reporting agencies and users.
3. **Privacy-Preserving Lawful Interception (LI):** Addressing the delicate balance between public safety and individual rights, we proposed a Self-Sovereign Identity (SSI) framework for LI. By standardizing the interaction between Law Enforcement Agencies (LEAs) and Monitoring Facilities (LEMFs) using Verifiable Credentials, the proposed system ensures that surveillance activities are automated, strictly authorized, auditable, and revocable, mitigating the risks of unauthorized data access and insider threats.
4. **Edge-Based Parental Control:** We shifted the paradigm of parental control from cloud-based monitoring to a router-based edge architecture. This solution enables granular

traffic inspection and policy enforcement directly within the home network, preventing third-party vendors from profiling minors while maintaining effective filtering of harmful content.

5. **Vulnerability Analysis in IoT Safeguards:** We exposed critical architectural weaknesses in current IoT defense systems. Our proof-of-concept demonstrated how legitimate IoT traffic can be exploited as a covert channel for undetectable surveillance, proving that current security layers are insufficient against sophisticated adversaries who exploit the hardware-software interface.
6. **Auditing Generative AI Privacy:** Finally, we conducted the first comprehensive privacy audit of GenAI-powered browser assistants. Our analysis revealed that these tools often bypass privacy controls, engaging in extensive user tracking and profiling. This finding underscores the urgent need for architectural constraints on how Large Language Models (LLMs) interact with sensitive browser data.

The findings of this research highlight a dual reality. On one hand, decentralized technologies (Blockchain, SSI) and edge computing offer robust mechanisms to return control of data to the user. We have shown that it is technically feasible to build systems where privacy is enforced rather than just legally promised.

On the other hand, the vulnerabilities identified in IoT and GenAI environments indicate that the attack surface is expanding faster than defensive architectures are evolving. Privacy risks are no longer limited to data leakage but extend to behavioral profiling and covert surveillance embedded within trusted hardware and software supply chains.

In conclusion, this thesis asserts that privacy must be treated as a non-negotiable architectural requirement. The solutions proposed provide a blueprint for a more resilient digital infrastructure, where trust is established through code and cryptography, ensuring that technological progress remains compatible with civil liberties.

9.2 Future Research Directions

The future development of this thesis lies not in extending each contribution in isolation, but in consolidating them into a broader research agenda in which privacy is treated as an architectural property of digital systems. Across the different case studies discussed in this work, a common insight emerges: privacy is better protected when systems minimize disclosure by design, decentralize trust, reduce unnecessary data flows, and make access, control, and enforcement verifiable rather than merely declared.

A first research direction concerns the definition of a unified privacy-by-design framework able to connect the main architectural principles explored in this thesis. Selective disclosure, decentralized trust, local enforcement, auditable access control, and traffic-aware privacy monitoring should not remain separate design patterns applied to specific domains, but should evolve into reusable architectural primitives. Such a framework could support the design of systems in which privacy guarantees are embedded from the outset and remain

enforceable across different application contexts, from digital identity and cybersecurity to domestic networks and AI-assisted platforms.

A second direction concerns interoperability and real-world deployment. Several solutions presented in this thesis show promising results at the proof-of-concept level, but their broader impact depends on their ability to operate across heterogeneous standards, infrastructures, and regulatory contexts. Future work should therefore investigate how privacy-preserving architectures can be integrated with emerging identity ecosystems, large-scale decentralized infrastructures, home-network environments, and multi-platform software ecosystems. The objective should be to move from concept validation to deployable architectures capable of operating under realistic technical, organizational, and legal constraints. This is particularly relevant because the thesis itself highlights privacy both as a technical and governance problem, not merely as a compliance requirement.

A third research direction is the development of continuous auditing and assurance mechanisms. One of the strongest lessons from this thesis is that privacy cannot be assumed even in systems that present themselves as protective. As the case studies presented in this thesis consistently demonstrate, declared protections may be incomplete, bypassable, or opaque. Future research should therefore focus on architectures that support continuous verification of privacy properties through automated audits, runtime monitoring, and measurable guarantees. In this sense, privacy engineering should move closer to dependability engineering: systems should not only be designed to protect privacy, but also to demonstrate, over time, that such protection remains effective under evolving threats, software updates, and changes in usage conditions. This need is consistent with the thesis's call for stronger, standardised, and automated audit approaches, especially in rapidly evolving environments.

A fourth and more forward-looking direction concerns the expansion of privacy research toward increasingly hybrid environments, where the boundary between hardware, software, network infrastructure, and artificial intelligence becomes blurred. The thesis shows that privacy risks no longer arise only from explicit data disclosure, but also from architectural opacity, covert signalling, behavioural inference, and profiling embedded within trusted platforms. Future work should therefore investigate privacy-preserving architectures that remain robust even when sensitive functions are distributed across devices, edge nodes, cloud services, and intelligent agents. This includes studying how local control, cryptographic accountability, and verifiable policy enforcement can be maintained in environments shaped by autonomous software, embedded sensing, and AI-driven personalization.

Ultimately, the broader future of this research is to move from domain-specific privacy solutions toward a general theory and practice of privacy-preserving system architecture. The long-term goal is not simply to build tools that reduce exposure in individual scenarios, but to establish design principles through which digital systems can be made selectively transparent, verifiable, minimally intrusive, and resilient against both centralized abuse and hidden forms of surveillance. In this perspective, the contributions of this thesis can be seen as the first building blocks of a wider architectural paradigm in which privacy is engineered as a core system property and continuously defended across the full lifecycle of digital technologies.

References

- [1] G. Lax, A. Canino, and L. Musarella, "A blockchain-based approach for certifying information at different levels of granularity according to the data minimization principle of the gdpr", in *2024 8th International Conference on Computer, Software and Modeling (ICCSM)*, IEEE, 2024, pp. 31–35.
- [2] A. L. Canino and G. Lax, "A cost-effective solution leveraging public blockchain for massively sharing malware signatures", *Journal of Information Security and Applications*, 2025.
- [3] F. Buccafurri, A. L. Canino, V. De Angelis, A. Laurenda, and G. Lax, "An ssi-based solution to support lawful interception", *Applied Sciences*, vol. 15, no. 4, p. 2206, 2025.
- [4] A. L. Canino, V. De Angelis, and G. Lax, "A router-based parental control tool for safe social network usage", in *International Conference on Advances in Social Networks Analysis and Mining*, Springer, 2024, pp. 91–100.
- [5] A. L. Canino and G. Lax, "Enabling lawful interception in environments protected by iot safeguard", in *International Conference on Electronic Government and the Information Systems Perspective*, Springer, 2024, pp. 139–153.
- [6] Y. Vekaria et al., "Big help or big brother? auditing tracking, profiling, and personalization in generative ai assistants", *USENIX Security '25*, 2025.
- [7] Wikipedia, X.509, [Accessed January 15, 2024], 2024. [Online]. Available: <https://en.wikipedia.org/wiki/X.509>.
- [8] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system", *Decentralized Business Review*, p. 21 260, 2008.
- [9] P. Wei, D. Wang, Y. Zhao, S. K. S. Tyagi, and N. Kumar, "Blockchain data-based cloud data integrity protection mechanism", *Future Generation Computer Systems*, vol. 102, pp. 902–911, 2020.
- [10] V. Buterin et al., "A next-generation smart contract and decentralized application platform", *white paper*, vol. 3, no. 37, pp. 2–1, 2014.
- [11] F. Buccafurri, V. De Angelis, and S. Lazzaro, "A blockchain-based framework to enhance anonymous services with accountability guarantees", *Future Internet*, vol. 14, no. 8, 2022, ISSN: 1999-5903. DOI: 10.3390/fi14080243. [Online]. Available: <https://www.mdpi.com/1999-5903/14/8/243>.
- [12] L. Interception, "Handover interface for the lawful interception of telecommunications traffic", *ETSI ES*, vol. 201, p. 671, 2004.
- [13] S. D. Guhl and R. Pendse, "The communications assistance for law enforcement act (calea)", *Information Security Journal: A Global Perspective*, vol. 17, no. 3, pp. 110–113, 2008.

- [14] E. Gold, “The history of artificial intelligence from the 1950s to today”, *freeCodeCamp.org*, Apr. 2023.
- [15] A. Vaswani, “Attention is all you need”, *Advances in Neural Information Processing Systems*, 2017.
- [16] A. Radford, “Improving language understanding by generative pre-training”, 2018.
- [17] A. Zewe, “Explained gen ai”, *MIT News*, Nov. 2023.
- [18] S. Minaee et al., “Large language models: A survey”, *arXiv preprint arXiv:2402.06196*, 2024.
- [19] K. Buchholz, “Chart: The extreme cost of training ai models | statista”, *Statista*, 2024. [Online]. Available: <https://www.statista.com/chart/33114/estimated-cost-of-training-selected-ai-models/>.
- [20] OpenAI, *Chatgpt*, <https://chatgpt.com/>, 2025.
- [21] Google, *Gemini*, <https://gemini.google.com/>, 2025.
- [22] Llama, *Meta*, <https://www.llama.com/>, 2025.
- [23] R. Shirey, “Internet security glossary (rfc 2828)”, *The Internet Society*, 2000.
- [24] General Data Protection Regulation, <https://eugdpr.org>, 2018.
- [25] G. Lax and A. Russo, “A lightweight scheme exploiting social networks for data minimization according to the gdpr”, *IEEE Transactions on Computational Social Systems*, vol. 8, no. 2, pp. 388–397, 2021.
- [26] I. Siegert, V. S. Varod, N. Carmi, and P. Kamocki, “Personal data protection and academia: Gdpr issues and multi-modal data-collections”, *Online Journal of Applied Knowledge Management (OJAKM)*, vol. 8, no. 1, pp. 16–31, 2020.
- [27] R. Stephen and A. Alex, “A review on blockchain security”, in *IOP conference series: materials science and engineering*, IOP Publishing, vol. 396, 2018, p. 012 030.
- [28] G. Lax, A. Russo, and L. S. Fasci`, “A blockchain-based approach for matching desired and real privacy settings of social network users”, *Information Sciences*, vol. 557, pp. 220–235, 2021.
- [29] L. Liu, W. Zhang, and C. Han, “A survey for the application of blockchain technology in the media”, *Peer-to-Peer Networking and Applications*, vol. 14, no. 5, pp. 3143–3165, 2021.
- [30] M. K. Lim, Y. Li, C. Wang, and M.-L. Tseng, “A literature review of blockchain technology applications in supply chains: A comprehensive analysis of themes, methodologies and industries”, *Computers & industrial engineering*, vol. 154, p. 107 133, 2021.

- [31] S. Pathak, V. Gupta, N. Malsa, A. Ghosh, and R. Shaw, "Blockchain-based academic certificate verification system—a review", *Advanced Computing and Intelligent Technologies: Proceedings of ICACIT 2022*, pp. 527–539, 2022.
- [32] R. F. Ghani, A. A. Salman, A. B. Khudhair, and L. Aljobouri, "Blockchain-based student certificate management and system sharing using hyperledger fabric platform", *Periodicals of Engineering and Natural Sciences*, vol. 10, no. 2, pp. 207–218, 2022.
- [33] R. A. Jaafar and S. N. Alsaad, "Enhancing educational certificate verification with blockchain and ipfs: A decentralized approach using hyperledger fab.", *TEM Journal*, vol. 12, no. 4, 2023.
- [34] O. S. Saleh, O. Ghazali, and N. B. Idris, "Enhancing academic certificate privacy with a hyperledger fabric blockchain-based access control approach", *SN Computer Science*, vol. 4, no. 5, p. 602, 2023.
- [35] B. M. Nguyen, T.-C. Dao, and B.-L. Do, "Towards a blockchain-based certificate authentication system in vietnam", *PeerJ Computer Science*, vol. 6, e266, 2020.
- [36] S. Reno, M. Ahmed, S. A. Jui, and S. Dilshad, "Securing certificate management system using hyperledger based private blockchain", in *2022 International Conference on Innovations in Science, Engineering and Technology (ICISSET)*, IEEE, 2022, pp. 46–51.
- [37] G. Caldarelli and J. Ellul, "Trusted academic transcripts on the blockchain: A systematic literature review", *Applied Sciences*, vol. 11, no. 4, p. 1842, 2021.
- [38] R. Xie et al., "Ethereum-blockchain-based technology of decentralized smart contract certificate system", *IEEE Internet of Things Magazine*, vol. 3, no. 2, pp. 44–50, 2020.
- [39] N. K. Noorhizama, Z. Abdullah, S. Kasim, I. R. A. Hamid, and M. A. M. Isa, "Verification of ph. d. certificate using qr code on blockchain ethereum", *JOIV: International Journal on Informatics Visualization*, vol. 7, no. 3, pp. 716–725, 2023.
- [40] M. Sabt, M. Achemlal, and A. Bouabdallah, "Trusted execution environment: What it is, and what it is not", in *2015 IEEE Trustcom/BigDataSE/IsPa*, IEEE, vol. 1, 2015, pp. 57–64.
- [41] V. Buterin, *Ethereum*, <https://ethereum.org/en/layer-2/>, 2024.
- [42] A. Yakovenko, "Solana: A new architecture for a high performance blockchain v0. 8.13", *Whitepaper*, 2018.
- [43] D. N. Durov, *The Open Network*, <https://ton.tg/whitepaper.pdf>, 2021.

- [44] F. A. Aboaoja, A. Zainal, F. A. Ghaleb, B. A. S. Al-rimy, T. A. E. Eisa, and A. A. H. Elnour, "Malware detection issues, challenges, and future directions: A survey", *Applied Sciences*, vol. 12, no. 17, p. 8482, 2022.
- [45] S. Morgan, *Cybercrime Cost*, <https://cybersecurityventures.com/hackerpocalypse-cybercrime-report-2016/>, 2016.
- [46] J. Tang et al., "Android malware detection based on a novel mixed bytecode image combined with attention mechanism", *Journal of Information Security and Applications*, vol. 82, p. 103 721, 2024.
- [47] L. Huang, J. Xue, Y. Wang, Z. Liu, J. Chen, and Z. Kong, "Whgdroid: Effective android malware detection based on weighted heterogeneous graph", *Journal of Information Security and Applications*, vol. 77, p. 103 556, 2023.
- [48] F. Faghihi, M. Zulkernine, and S. Ding, "Aim: An android interpretable malware detector based on application class modeling", *Journal of Information Security and Applications*, vol. 75, p. 103 486, 2023.
- [49] Ö. A. Aslan and R. Samet, "A comprehensive review on malware detection approaches", *IEEE Access*, vol. 8, pp. 6249–6271, 2020.
- [50] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey", *International journal of web and grid services*, vol. 14, no. 4, pp. 352–375, 2018.
- [51] R. Kakkar, R. Gupta, S. Agrawal, S. Tanwar, and R. Sharma, "Blockchain-based secure and trusted data sharing scheme for autonomous vehicle underlying 5g", *Journal of Information Security and Applications*, vol. 67, p. 103 179, 2022.
- [52] A. Malik, A. K. Sharma, et al., "Blockchain-based digital chain of custody multimedia evidence preservation framework for internet-of-things", *Journal of Information Security and Applications*, vol. 77, p. 103 579, 2023.
- [53] S. Homayoun, A. Dehghantanha, R. M. Parizi, and K.-K. R. Choo, "A blockchain-based framework for detecting malicious mobile applications in app stores", in *2019 IEEE Canadian Conference of Electrical and Computer Engineering (CCECE)*, IEEE, 2019, pp. 1–4.
- [54] G. Martin, D. Spencer, A. Hair, S. Laudanna, C. A. Visaggio, et al., "Mobile malware detection using consortium blockchain", in *Cybersecurity for Artificial Intelligence*, Springer, 2022, pp. 137–160.
- [55] J. Gu, B. Sun, X. Du, J. Wang, Y. Zhuang, and Z. Wang, "Consortium blockchain-based malware detection in mobile devices", *IEEE Access*, vol. 6, pp. 12 118–12 128, 2018.
- [56] T. Golomb, Y. Mirsky, and Y. Elovici, "Ciota: Collaborative iot anomaly detection via blockchain", *arXiv preprint arXiv:1803.03807*, 2018.

- [57] R. Kumar, X. Zhang, W. Wang, R. U. Khan, J. Kumar, and A. Sharif, "A multimodal malware detection technique for android iot devices using various features", *IEEE access*, vol. 7, pp. 64 411–64 430, 2019.
- [58] R. Kumar, W. Wang, J. Kumar, T. Yang, W. Ali, et al., "Collective intelligence: Decentralized learning for android malware detection in iot with blockchain", *arXiv preprint arXiv:2102.13376*, 2021.
- [59] D. Denysiuk, O. Geidarova, M. Kapustian, S. Lysenko, and A. Sachenko, "Blockchain-based deep learning algorithm for detecting malware.", in *IntelITSIS*, 2023, pp. 529–538.
- [60] M. Wazid, A. K. Das, and S. Shetty, "Bsfr-sh: Blockchain-enabled security framework against ransomware attacks for smart healthcare", *IEEE Transactions on Consumer Electronics*, vol. 69, no. 1, pp. 18–28, 2022.
- [61] S. Sayeed and H. Marco-Gisbert, "Assessing blockchain consensus and security mechanisms against the 51% attack", *Applied sciences*, vol. 9, no. 9, p. 1788, 2019.
- [62] R. Fuji et al., "Investigation on sharing signatures of suspected malware files using blockchain technology", in *International Multi Conference of Engineers and Computer Scientists (IMECS)*, 2019, pp. 94–99.
- [63] Z. Bazrafshan, H. Hashemi, S. M. H. Fard, and A. Hamzeh, "A survey on heuristic malware detection techniques", in *The 5th conference on information and knowledge technology*, IEEE, 2013, pp. 113–120.
- [64] V. S. Sathyanarayan, P. Kohli, and B. Bruhadeshwar, "Signature generation and detection of malware families", in *Information Security and Privacy: 13th Australasian Conference, ACISP 2008, Wollongong, Australia, July 7-9, 2008. Proceedings 13*, Springer, 2008, pp. 336–349.
- [65] Y. Tang, B. Xiao, and X. Lu, "Using a bioinformatics approach to generate accurate exploit-based signatures for polymorphic worms", *computers & security*, vol. 28, no. 8, pp. 827–842, 2009.
- [66] M. Zheng, M. Sun, and J. C. Lui, "Droid analytics: A signature based analytic system to collect, extract, analyze and associate android malware", in *2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, IEEE, 2013, pp. 163–171.
- [67] P. Shijo and A. Salim, "Integrated static and dynamic analysis for malware detection", *Procedia Computer Science*, vol. 46, pp. 804–811, 2015.
- [68] R. Thangavelooa, W. W. Jinga, C. K. Lenga, and J. Abdullaha, "Datdroid: Dynamic analysis technique in android malware detection", *International Journal on Advanced Science, Engineering and Information Technology*, vol. 10, no. 2, pp. 536–541, 2020.

- [69] R. Mosli, R. Li, B. Yuan, and Y. Pan, "A behavior-based approach for malware detection", in *Advances in Digital Forensics XIII: 13th IFIP WG 11.9 International Conference, Orlando, FL, USA, January 30-February 1, 2017, Revised Selected Papers 13*, Springer, 2017, pp. 187–201.
- [70] H. S. Galal, Y. B. Mahdy, and M. A. Atiea, "Behavior-based features model for malware detection", *Journal of Computer Virology and Hacking Techniques*, vol. 12, pp. 59–67, 2016.
- [71] Y. Fukushima, A. Sakai, Y. Hori, and K. Sakurai, "A behavior based malware detection scheme for avoiding false positive", in *2010 6th IEEE workshop on secure network protocols*, IEEE, 2010, pp. 79–84.
- [72] W. Liu, P. Ren, K. Liu, and H.-x. Duan, "Behavior-based malware analysis and detection", in *2011 first international workshop on complexity and data mining*, IEEE, 2011, pp. 39–42.
- [73] N. Li, Z. Zhang, X. Che, Z. Guo, and J. Cai, "A survey on feature extraction methods of heuristic malware detection", in *Journal of Physics: Conference Series*, IOP Publishing, vol. 1757, 2021, p. 012 071.
- [74] A. Ninyesiga, "Improved heuristics for malware detection", A Working Paper on Research Gate accessible via [https://www. researchgate ...](https://www.researchgate...), Tech. Rep., 2016.
- [75] A. Tang, S. Sethumadhavan, and S. J. Stolfo, "Unsupervised anomaly-based malware detection using hardware features", in *Research in Attacks, Intrusions and Defenses: 17th International Symposium, RAID 2014, Gothenburg, Sweden, September 17-19, 2014. Proceedings 17*, Springer, 2014, pp. 109–129.
- [76] J. M. Vidal, A. L. S. Orozco, and L. J. G. Villalba, "Alert correlation framework for malware detection by anomaly-based packet payload analysis", *Journal of Network and Computer Applications*, vol. 97, pp. 11–22, 2017.
- [77] K. Liu, S. Xu, G. Xu, M. Zhang, D. Sun, and H. Liu, "A review of android malware detection approaches based on machine learning", *IEEE access*, vol. 8, pp. 124 579–124 607, 2020.
- [78] H. Rathore, S. Agarwal, S. K. Sahay, and M. Sewak, "Malware detection using machine learning and deep learning", in *Big Data Analytics: 6th International Conference, BDA 2018, Warangal, India, December 18–21, 2018, Proceedings 6*, Springer, 2018, pp. 402–411.
- [79] T. Sharma and D. Rattan, "Malicious application detection in android—a systematic literature review", *Computer Science Review*, vol. 40, p. 100 373, 2021.
- [80] B. N. Levine, C. Shields, and N. B. Margolin, "A survey of solutions to the sybil attack", *University of Massachusetts Amherst, Amherst, MA*, vol. 7, p. 224, 2006.

- [81] K. Sukhani, S. Sawant, S. Maniar, and R. Pawar, "Automating the bypass of image-based captcha and assessing security", in *2021 12th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, IEEE, 2021, pp. 01–08.
- [82] A. Hassanat, "Bypassing captcha by machine a proof for passing the turing test", *arXiv preprint arXiv:1409.0925*, 2014.
- [83] Chainlink, *Chainlink Documentation*, <https://docs.chain.link/>, 2023.
- [84] Metamask, *A crypto wallet & gateway to blockchain apps*, <https://metamask.io/>, 2023.
- [85] Sepolia, *Sepolia Resources*, <https://sepolia.dev/>, 2023.
- [86] Solidity, *Solidity Documentation*, <https://docs.soliditylang.org/en/develop/>, 2022.
- [87] H. Al-Breiki, M. H. U. Rehman, K. Salah, and D. Svetinovic, "Trustworthy blockchain oracles: Review, comparison, and open research challenges", *IEEE Access*, vol. 8, pp. 85 675–85 685, 2020.
- [88] V. Buterin, *Ethereum danksharding*, <https://ethereum.org/en/roadmap/danksharding/>, 2024.
- [89] V. Doronin, "'lawful interception—a market access barrier in the european union'?", *Computer Law & Security Review*, vol. 51, p. 105 867, 2023.
- [90] R. Walters, "Illegal interception of data", in *Cybersecurity and Data Laws of the Commonwealth: International Trade, Investment and Arbitration*, Springer, 2023, pp. 267–274.
- [91] D. Omand, "Examining the ethics of spying: A practitioner's view", *Criminal Law and Philosophy*, vol. 18, no. 3, pp. 805–818, 2024.
- [92] G. Li et al., "A critical review of cyber-physical security for building automation systems", *Annual Reviews in Control*, vol. 55, pp. 237–254, 2023.
- [93] X. Bultel and C. Onete, "Pairing-free secure-channel establishment in mobile networks with fine-grained lawful interception", in *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*, 2022, pp. 968–970.
- [94] D. Spiekermann, J. Keller, and T. Eggendorfer, "Improving lawful interception in virtual datacenters", in *Proceedings of the Central European Cybersecurity Conference 2018*, ser. CECC 2018, Ljubljana, Slovenia: Association for Computing Machinery, 2018, ISBN: 9781450365154. DOI: 10.1145/3277570.3277578.
- [95] 3. T. W. G. 33, "Security architecture for core network and terminals (release 16)", 3rd Generation Partnership Project (3GPP), Tech. Rep. 3GPP TS 33.108, 2020. [Online]. Available: https://www.3gpp.org/ftp//Specs/archive/33_series/33.108/33108-f10.zip.

- [96] I.-T. W. G. Y.2770, “Requirements for deep packet inspection in next generation networks”, ITU-T (International Telecommunication Union - Telecommunication Standardization Sector), Tech. Rep. ITU-T Y.2770, 2015. [Online]. Available: <https://www.itu.int/>.
- [97] E. (T. S. Institute), “Lawful interception (li); handover interface for the lawful interception of telecommunications traffic”, ETSI, Tech. Rep. ETSI TS 101 671 V3.15.1, 2018. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/101600_101699/101671/03.15.01_60/ts_101671v031501p.pdf.
- [98] V. Turanjanin, “When does bulk interception of communications violate the right to privacy? the limits of the state’s power and the european court of human rights approach”, *International Cybersecurity Law Review*, vol. 4, no. 1, pp. 115–136, 2023.
- [99] M. Gorge, “Lawful interception – key concepts, actors, trends and best practice considerations”, *Computer Fraud & Security*, vol. 2007, no. 9, pp. 10–14, 2007, ISSN: 1361-3723. DOI: [https://doi.org/10.1016/S1361-3723\(07\)70116-1](https://doi.org/10.1016/S1361-3723(07)70116-1).
- [100] A. Tobin and D. Reed, “The inevitable rise of self-sovereign identity”, *The Sovrin Foundation*, vol. 29, no. 2016, p. 18, 2016.
- [101] A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, “A survey on essential components of a self-sovereign identity”, *Computer Science Review*, vol. 30, pp. 80–86, 2018.
- [102] D. Pöhn, M. Grabatin, and W. Hommel, “Analyzing the threats to blockchain-based self-sovereign identities by conducting a literature survey”, *Applied Sciences*, vol. 14, no. 1, p. 139, 2023.
- [103] F. Buccafurri, V. De Angelis, S. Lazzaro, and A. Pugliese, “Enforcing security policies on interacting authentication systems”, *Computers & Security*, vol. 140, p. 103 771, 2024, ISSN: 0167-4048. DOI: <https://doi.org/10.1016/j.cose.2024.103771>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404824000725>.
- [104] M. Naghmouchi, M. Laurent, C. Levallois-Barth, and N. Kaaniche, “Comparative analysis of technical and legal frameworks of various national digital identity solutions”, *arXiv preprint arXiv:2310.01006*, 2023.
- [105] D. Khovratovich and J. Law, “Sovrin: Digital identities in the blockchain era”, *Github Commit by jasonalaw October*, vol. 17, pp. 38–99, 2017.
- [106] D. Ferraiolo, J. Cugini, D. R. Kuhn, et al., “Role-based access control (rbac): Features and motivations”, in *Proceedings of 11th annual computer security application conference*, 1995, pp. 241–48.

- [107] B. Wijnen, R. Presuhn, and K. McCloghrie, *Rfc3415: View-based access control model (vacm) for the simple network management protocol (snmp)*, 2002.
- [108] Sovrin Governance Framework Working Group, *Sovrin glossary appendix d*, <https://sovrin.org/wp-content/uploads/Sovrin-Glossary-V2.pdf>, Accessed online: 30-01-2025, 2025.
- [109] A. Satybaldy, M. S. Ferdous, and M. Nowostawski, "A taxonomy of challenges for self-sovereign identity systems", *IEEE Access*, 2024.
- [110] M. Davie, D. Gisolfi, D. Hardman, J. Jordan, D. O'Donnell, and D. Reed, "The trust over ip stack", *IEEE Communications Standards Magazine*, vol. 3, no. 4, pp. 46–51, 2019.
- [111] A. Preukschat and D. Reed, *Self-sovereign identity*. Manning Publications, 2021.
- [112] Y. Ren, X. Liu, Q. Wu, L. Wang, and W. Zhang, "Cryptographic accumulator and its application: A survey", *Security and Communication Networks*, vol. 2022, no. 1, p. 5 429 195, 2022.
- [113] C. V. Helliar, L. Crawford, L. Rocca, C. Teodori, and M. Veneziani, "Permissionless and permissioned blockchain diffusion", *International Journal of Information Management*, vol. 54, p. 102 136, 2020, issn: 0268-4012. doi: <https://doi.org/10.1016/j.ijinfomgt.2020.102136>.
- [114] D. Hardman, *Aries rfc 0005: Did communication*, <https://github.com/hyperledger/aries-rfcs/blob/main/concepts/0005-didcomm/README.md>, Accessed online: 2022, 2022.
- [115] W. W. W. Consortium et al., "Verifiable credentials data model 1.0: Expressing verifiable information on the web", <https://www.w3.org/TR/vc-data-model/?#core-data-model>, 2019.
- [116] Sovrin Governance Framework Working Group, *Sovrin governance framework*, <https://sovrin.org/library/sovrin-governance-framework/>, Accessed online: 30-01-2025, 2025.
- [117] M. P. Bhattacharya, P. Zavorsky, and S. Butakov, "Enhancing the security and privacy of self-sovereign identities on hyperledger indy blockchain", in *2020 International Symposium on Networks, Computers and Communications (ISNCC)*, IEEE, 2020, pp. 1–7.
- [118] Trinsic, *Welcome to trinsic*, <https://github.com/trinsic-id>, Accessed online: 2022, 2022.
- [119] M. Grinberg, *Flask web development: developing web applications with python*. "O'Reilly Media, Inc.", 2018.
- [120] J. Glöckler, J. Sedlmeir, M. Frank, and G. Fridgen, "A systematic review of identity and access management requirements in enterprises and potential

- contributions of self-sovereign identity”, *Business & Information Systems Engineering*, vol. 66, no. 4, pp. 421–440, 2024.
- [121] R. Soltani, U. T. Nguyen, and A. An, “A survey of self-sovereign identity ecosystem”, *Security and Communication Networks*, vol. 2021, no. 1, p. 8 873 429, 2021.
- [122] A.-S. Shehu, “On the compliance of self-sovereign identity with gdpr principles: A critical review”, *arXiv preprint arXiv:2409.03624*, 2024.
- [123] P. Voigt and A. Von dem Bussche, “The eu general data protection regulation (gdpr)”, *A Practical Guide, 1st Ed.*, Cham: Springer International Publishing, vol. 10, no. 3152676, pp. 10–5555, 2017.
- [124] N. Naik and P. Jenkins, “Your identity is yours: Take back control of your identity using gdpr compatible self-sovereign identity”, in *2020 7th International Conference on Behavioural and Social Computing (BESC)*, IEEE, 2020, pp. 1–6.
- [125] G. Kondova and J. Erbguth, “Self-sovereign identity on public blockchains and the gdpr”, in *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, 2020, pp. 342–345.
- [126] N. Naik and P. Jenkins, “Uport open-source identity management system: An assessment of self-sovereign identity and user-centric data platform built on blockchain”, in *2020 IEEE International Symposium on Systems Engineering (ISSE)*, IEEE, 2020, pp. 1–7.
- [127] Q. Stokkink, G. Ishmaev, D. Epema, and J. Pouwelse, “A truly self-sovereign identity system”, in *2021 IEEE 46th Conference on Local Computer Networks (LCN)*, IEEE, 2021, pp. 1–8.
- [128] A. Grüner, A. Mühle, and C. Meinel, “An integration architecture to enable service providers for self-sovereign identity”, in *2019 IEEE 18th International Symposium on Network Computing and Applications (NCA)*, IEEE, 2019, pp. 1–5.
- [129] R. A. Pava-Díaz, J. Gil-Ruiz, and D. A. López-Sarmiento, “Self-sovereign identity on the blockchain: Contextual analysis and quantification of ssi principles implementation”, *Frontiers in Blockchain*, vol. 7, p. 1 443 362, 2024.
- [130] P. J. Windley, “Sovrin: An identity metasystem for self-sovereign identity”, *Frontiers in Blockchain*, vol. 4, p. 626 726, 2021.
- [131] D. Reed, J. Law, and D. Hardman, “The technical foundations of sovrin”, *The Technical Foundations of Sovrin*, 2016, <https://sovrin.org/wp-content/uploads/2017/04/The-Technical-Foundations-of-Sovrin.pdf>.

- [132] D. López and B. Farooq, "A multi-layered blockchain framework for smart mobility data-markets", *Transportation Research Part C: Emerging Technologies*, vol. 111, pp. 588–615, 2020, issn: 0968-090X. doi: <https://doi.org/10.1016/j.trc.2020.01.002>.
- [133] P. C. Bartolomeu, E. Vieira, and J. Ferreira, "Pay as you go: A generic crypto tolling architecture", *IEEE Access*, vol. 8, pp. 196 212–196 222, 2020. doi: 10.1109/ACCESS.2020.3034299.
- [134] M. d. V. Barros, F. Schardong, and R. F. Custódio, "Leveraging self-sovereign identity, blockchain, and zero-knowledge proof to build a privacy-preserving vaccination pass", *arXiv preprint arXiv:2202.09207*, 2022.
- [135] F. Ahmed and S. A. Hussain, "A privacy-preserving cross-domain network access services using sovrin identifier", in *2021 International Conference on Cyber Warfare and Security (ICWS)*, 2021, pp. 30–37. doi: 10.1109/ICWS53234.2021.9702980.
- [136] L. Helminger, D. Kales, S. Ramacher, and R. Walch, "Multi-party revocation in sovrin: Performance through distributed trust", in *Cryptographers' Track at the RSA Conference*, Springer, 2021, pp. 527–551.
- [137] A. Abraham, K. Theuermann, and E. Kirchengast, "Qualified eid derivation into a distributed ledger based idm system", in *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, 2018, pp. 1406–1412. doi: 10.1109/TrustCom/BigDataSE.2018.00195.
- [138] M. Grabatin and W. Hommel, "Self-sovereign identity management in wireless ad hoc mesh networks", in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2021, pp. 480–486.
- [139] Sovrin foundation, *Use case archives*, <https://sovrin.org/category/use-cases/>, Accessed online: 30-01-2025, 2025.
- [140] A. Ling and S. Butakov, "Trust framework for self-sovereign id in meta-verse health care applications", *Data Science and Management*, vol. 7, no. 4, pp. 304–313, 2024, issn: 2666-7649. doi: <https://doi.org/10.1016/j.dsm.2024.04.003>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666764924000274>.
- [141] M. Shuaib, S. Alam, M. S. Alam, and M. S. Nasir, "Self-sovereign identity for healthcare using blockchain", *Materials Today: Proceedings*, vol. 81, pp. 203–207, 2023.

- [142] P. Bai, S. Kumar, G. Aggarwal, M. Mahmud, O. Kaiwartya, and J. Lloret, “Self-sovereignty identity management model for smart healthcare system”, *Sensors*, vol. 22, no. 13, p. 4714, 2022.
- [143] S. De Diego, C. Regueiro, and G. Macia-Fernandez, “Enabling identity for the iot-as-a-service business model”, *IEEE Access*, vol. 9, pp. 159 965–159 975, 2021.
- [144] P. C. Bartolomeu, E. Vieira, S. M. Hosseini, and J. Ferreira, “Self-sovereign identity: Use-cases, technologies, and challenges for industrial iot”, in *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, IEEE, 2019, pp. 1173–1180.
- [145] L. Cocco, R. Tonelli, and M. Marchesi, “Blockchain and self sovereign identity to support quality in the food supply chain”, *Future Internet*, vol. 13, no. 12, p. 301, 2021.
- [146] S. Bronitt and J. Stellios, “Telecommunications interception in australia: Recent trends and regulatory prospects”, *Telecommunications Policy*, vol. 29, no. 11, pp. 875–888, 2005, Network Security, Survivability and Surveillance, issn: 0308-5961. doi: <https://doi.org/10.1016/j.telpol.2005.06.010>.
- [147] H. Abelson et al., “Keys under doormats: Mandating insecurity by requiring government access to all data and communications”, *Journal of Cybersecurity*, vol. 1, no. 1, pp. 69–79, 2015.
- [148] I. Brown, “Lawful interception capability requirements”, *Computers & Law*, Aug/Sep, 2013, https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2309413_code892424.pdf.
- [149] M. Christen, B. Gordijn, and M. Loi, *The ethics of cybersecurity*. Springer Nature, 2020, <https://library.oapen.org/handle/20.500.12657/47324>.
- [150] F. Cajani, ““all along the watchtower”: Matters not yet solved regarding communication interception systems and electronic data retained on foreign servers”, *Handling and Exchanging Electronic Evidence Across Europe*, pp. 59–71, 2018.
- [151] M. Pleva, A. Cizmar, and L. Dobos, “Voice quality measuring setup with automatic voice over ip call generator and lawful interception packet analyzer”, *Journal of Electrical and Electronics Engineering*, vol. 5, no. 1, p. 191, 2012.
- [152] F. Buccafurri, A. Consoli, C. Labrini, and A. M. Nesurini, “A solution to support integrity in the lawful interception ecosystem”, in *International Conference on Electronic Government and the Information Systems Perspective*, Springer, 2021, pp. 21–33.

- [153] X. Xu, W.-K. Jia, Y. Wu, and X. Wang, "On the optimal lawful intercept access points placement problem in hybrid software-defined networks", *Sensors*, vol. 21, no. 2, 2021, issn: 1424-8220. doi: 10.3390/s21020428. [Online]. Available: <https://www.mdpi.com/1424-8220/21/2/428>.
- [154] P. Branch, A. Pavlicic, and G. Armitage, "Using mac addresses in the lawful interception of ip traffic", in *Proc Australian Telecommunications Networks & Applications Conference (ATNAC), Sydney, Australia, 2004*, pp. 9–11.
- [155] B. Karpagavinayagam, R. State, and O. Festor, "Monitoring architecture for lawful interception in voip networks", in *Second International Conference on Internet Monitoring and Protection (ICIMP 2007)*, IEEE, 2007, pp. 5–5.
- [156] T. Aljaž, B. Imperl, and A. Švigelj, "Border gateway function performance requirements for the lawful intercept of voice at ims architecture", *AEU - International Journal of Electronics and Communications*, vol. 62, no. 8, pp. 610–621, 2008, issn: 1434-8411. doi: <https://doi.org/10.1016/j.aeue.2007.08.006>.
- [157] A. Muñoz, M. Urueña, R. Aparicio, and G. Rodríguez de los Santos, "Digital wiretap warrant: Improving the security of etsi lawful interception", *Digital Investigation*, vol. 14, pp. 1–16, 2015, issn: 1742-2876. doi: <https://doi.org/10.1016/j.diin.2015.04.005>.
- [158] D. I. Alamillo and S. Schwalm, "Self-sovereign-identity & eidas: A contradiction? challenges and chances of [eidas2]", *European Review of Digital Administration & Law—Erdal2021*, vol. 2, no. 2, pp. 89–108, 2021.
- [159] J. Swart, C. Peters, and M. Broersma, "Sharing and discussing news in private social media groups: The social function of news and current affairs in location-based, work-oriented and leisure-focused communities", *Digital journalism*, vol. 7, no. 2, pp. 187–205, 2019.
- [160] A. K. Ghosh, K. Badillo-Urquiola, S. Guha, J. J. LaViola Jr, and P. J. Wisniewski, "Safety vs. surveillance: What children have to say about mobile apps for parental control", in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–14.
- [161] G. Wang, J. Zhao, M. Van Kleek, and N. Shadbolt, "Protection or punishment? relating the design space of parental control apps and perceptions about them to support parenting for online safety", *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW2, pp. 1–26, 2021.
- [162] S. Ali, M. Elgharabawy, Q. Duchaussoy, M. Mannan, and A. Youssef, "Betrayed by the guardian: Security and privacy risks of parental control solutions", in *Proceedings of the 36th Annual Computer Security Applications Conference*, 2020, pp. 69–83.

- [163] J. Fardouly, N. R. Magson, C. J. Johnco, E. L. Oar, and R. M. Rapee, "Parental control of the time preadolescents spend on social media: Links with preadolescents' social media appearance comparisons and mental health", *Journal of youth and adolescence*, vol. 47, pp. 1456–1468, 2018.
- [164] P. Dumar, H. Atashpanjeh, and M. N. Al-Ameen, "" it's hard for him to make choices sometimes and he needs guidance": Re-orienting parental control for children", *Proceedings of the ACM on Human-Computer Interaction*, vol. 8, no. CSCW1, pp. 1–51, 2024.
- [165] J. Lundberg and O. Marklund, *Adolescents in control: Promoting adolescents autonomy in parental control applications*, 2023.
- [166] Y. Li, Y. Li, Q. Yan, and R. H. Deng, "Privacy leakage analysis in online social networks", *Computers & Security*, vol. 49, pp. 239–254, 2015, issn: 0167-4048. doi: <https://doi.org/10.1016/j.cose.2014.10.012>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167404814001588>.
- [167] S. Ramezani, T. Meskanen, and V. Niemi, "Parental control with edge computing and 5g networks", in *2021 29th Conference of Open Innovations Association (FRUCT)*, 2021, pp. 290–300. doi: 10.23919/FRUCT52173.2021.9435552.
- [168] N. Zakaria, L. K. Yew, N. M. A. Alias, and W. Husain, "Protecting privacy of children in social networking sites with rule-based privacy tool", in *8th International Conference on High-capacity Optical Networks and Emerging Technologies*, 2011, pp. 253–257. doi: 10.1109/HONET.2011.6149828.
- [169] S. A. Suchaad, K. Mashiko, N. B. Ismail, and M. H. Z. Abidin, "Blockchain use in home automation for children incentives in parental control", in *Proceedings of the 2018 International Conference on Machine Learning and Machine Intelligence*, 2018, pp. 50–53.
- [170] W. Fuertes, K. Quimbiulco, F. Galárraga, and J. L. García-Dorado, "On the development of advanced parental control tools", in *2015 1st International Conference on Software Security and Assurance (ICSSA)*, 2015, pp. 1–6. doi: 10.1109/ICSSA.2015.011.
- [171] W. Fuertes, K. Quimbiulco, F. Galárraga, and J. L. García-Dorado, "On the development of advanced parental control tools", in *2015 1st International Conference on Software Security and Assurance (ICSSA)*, IEEE, 2015, pp. 1–6.
- [172] G. Giorgi, A. La Marra, F. Martinelli, P. Mori, and A. Saracino, "Smart parental advisory: A usage control and deep learning-based framework for dynamic parental control on smart tv", in *Security and Trust Management*:

- 13th International Workshop, STM 2017, Oslo, Norway, September 14–15, 2017, Proceedings 13*, Springer, 2017, pp. 118–133.
- [173] G. A. Leon-Paredes, O. G. Bravo-Quezada, P. P. Bermeo-Aguaysa, M. J. Pelaez-Currillo, and L. L. Jimenez-Gonzalez, “Preventing cyberbullying on social networks with spanish parental control nlp system”, *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 11, 2023. doi: 10.14569/IJACSA.2023.01411141. [Online]. Available: <http://dx.doi.org/10.14569/IJACSA.2023.01411141>.
 - [174] S. Lazzaro, V. De Angelis, A. M. Mandalari, and F. Buccafurri, “Is your kettle smarter than a hacker? a scalable tool for assessing replay attack vulnerabilities on consumer iot devices”, in *2024 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2024, pp. 114–124. [Online]. Available: <https://doi.org/10.1109/PerCom59722.2024.10494466>.
 - [175] R. T. El-Maghraby, N. M. Abd Elazim, and A. M. Bahaa-Eldin, “A survey on deep packet inspection”, in *2017 12th International Conference on Computer Engineering and Systems (ICCES)*, IEEE, 2017, pp. 188–197.
 - [176] C. Xu, S. Chen, J. Su, S.-M. Yiu, and L. C. Hui, “A survey on regular expression matching for deep packet inspection: Applications, algorithms, and hardware platforms”, *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2991–3029, 2016.
 - [177] M. Trevisan, I. Drago, M. Mellia, and M. M. Munafo, “Automatic detection of dns manipulations”, in *2017 IEEE International Conference on Big Data (Big Data)*, IEEE, 2017, pp. 4010–4015.
 - [178] D. Chen et al., “Privacy-preserving encrypted traffic inspection with symmetric cryptographic techniques in iot”, *IEEE Internet of Things journal*, vol. 9, no. 18, pp. 17 265–17 279, 2022.
 - [179] F. Buccafurri, V. De Angelis, S. Lazzaro, and A. Pugliese, “Enforcing security policies on interacting authentication systems”, *Computers & Security*, p. 103 771, 2024.
 - [180] F. Buccafurri, V. De Angelis, and S. Lazzaro, “Mqtt-i: Achieving end-to-end data flow integrity in mqtt”, *IEEE Transactions on Dependable and Secure Computing*, pp. 1–18, 2024. doi: 10.1109/TDSC.2024.3358630.
 - [181] Netgear, *Netgear r8000*, <https://www.netgear.com/home/wifi/routers/r8000/>.
 - [182] OpenWrt, *Openwrt 23.03.5*, <https://openwrt.org/toh/netgear/r8000>.

- [183] zetago, *Ipssec modern ikev2 road-warrior configuration*, <https://openwrt.org/docs/guide-user/services/vpn/strongswan/roadwarrior>.
- [184] amirhosseinchoghaei, *Increase openwrt disk space*, <https://github.com/amirhosseinchoghaei/Increase-openwrt-disk-space>.
- [185] Creative Commons Attribution-NonCommercial-ShareAlike 2.5, *Scapy-python*, <https://scapy.readthedocs.io/en/latest/index.html>.
- [186] P. M. Chanal and M. S. Kakkasageri, "Security and privacy in iot: A survey", *Wireless Personal Communications*, vol. 115, no. 2, pp. 1667–1693, 2020.
- [187] E. M. Rudd, A. Rozsa, M. Günther, and T. E. Boult, "A survey of stealth malware attacks, mitigation measures, and steps toward autonomous open world solutions", *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 1145–1172, 2016.
- [188] E. T. Michailidis, D. G. Kogias, and I. Voyiatzis, "A review on hardware security countermeasures for iot: Emerging mechanisms and machine learning solutions", in *Proceedings of the 24th Pan-Hellenic Conference on Informatics*, 2020, pp. 268–271.
- [189] A. M. Mandalari, H. Haddadi, D. J. Dubois, and D. Choffnes, "Protected or porous: A comparative analysis of threat detection capability of iot safeguards", in *2023 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2023, pp. 3061–3078.
- [190] Baker, et al., *Cisco architecture for lawful intercept in ip networks*, <https://www.rfc-editor.org/rfc/rfc3924.html>, 2004.
- [191] European Telecommunications Standards Institute, *Lawful interception (li); concepts of interception in a generic network architecture*, https://www.etsi.org/deliver/etsi_tr/101900_101999/101943/02.01.01_60/tr_101943v020101p.pdf, 2004.
- [192] European Telecommunications Standards Institute, *Digital cellular telecommunications system (phase 2+) (gsm); universal mobile telecommunications system (umts); lte; 5g; lawful interception (li) architecture and functions (3gpp ts 33.127 version 16.12.0 release 16)*, https://www.etsi.org/deliver/etsi_ts/133100_133199/133127/16.12.00_60/ts_133127v161200p.pdf, 2004.
- [193] F. Boeira, M. Asplund, and M. Barcellos, *Provable Non-Frameability for 5G Lawful Interception*, 2023.
- [194] F. Intoci, J. Sturm, D. Fraunholz, A. Pyrgelis, and C. Barschel, "P3li5: Practical and confidential lawful interception on the 5g core", in *2023 IEEE Conference on Communications and Network Security (CNS)*, 2023, pp. 1–9. doi: 10.1109/CNS59707.2023.10288872.

- [195] K. Michail-Alexandros, X. George, C. Ioannis, and K. Anastasios, “Quantum computing and lawful interception applications”, in *International Conference on Information Technology & Systems*, Springer, 2023, pp. 363–373.
- [196] A. Mileva, A. Velinov, and D. Stojanov, *New covert channels in Internet of Things*, 2018.
- [197] K. Harris, W. Henry, and R. Dill, “A network-based iot covert channel”, in *2022 4th international conference on computer communication and the internet (ICCCI)*, IEEE, 2022, pp. 91–99.
- [198] P. Bedi and A. Dua, “Network steganography using the overflow field of timestamp option in an ipv4 packet”, *Procedia Computer Science*, vol. 171, pp. 1810–1818, 2020.
- [199] S. Qadir and B. Noor, “Applications of machine learning in digital forensics”, in *2021 International Conference on Digital Futures and Transformative Technologies (ICoDT2)*, IEEE, 2021, pp. 1–8.
- [200] M. Zamir et al., “Face detection & recognition from images & videos based on cnn & raspberry pi”, *Computation*, vol. 10, no. 9, p. 148, 2022.
- [201] M. H. Mofrad and D. Mosse, “Speech recognition and voice separation for the internet of things”, in *Proceedings of the 8th International Conference on the Internet of Things*, 2018, pp. 1–8.
- [202] W. Eddy, Ed. MTI Systems, *Rfc 9293 transmission control protocol (tcp)*, <https://www.ietf.org/rfc/rfc9293.html>, 2004.
- [203] R. Sobti and G. Geetha, “Cryptographic hash functions: A review”, *International Journal of Computer Science Issues (IJCSI)*, vol. 9, no. 2, p. 461, 2012.
- [204] Mikrotik, *Mikrotik rb750gr3*, <https://mikrotik.com/product/RB750Gr3>.
- [205] Raspberry Pi Foundation, *Raspberry pi 4 model b*, <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/>.
- [206] Apache License 2.0, *Kafka*, <https://kafka.apache.org/>.
- [207] Confluent Inc, *Confluent-kafka*, <https://pypi.org/project/confluent-kafka/>.
- [208] Adam Geitgey, *Face-recognition*, <https://pypi.org/project/face-recognition/>.
- [209] Apache Software License (Apache 2.0), *Opencv-python*, <https://pypi.org/project/opencv-python/>.
- [210] Gordon Lyon, *Nmap*, <https://nmap.org/>.
- [211] Wikipedia contributors, *Magic number (programming)* — *Wikipedia, the free encyclopedia*, [https://en.wikipedia.org/wiki/Magic_number_\(programming\)](https://en.wikipedia.org/wiki/Magic_number_(programming)), 2024.

- [212] Copyright Bitcoin Project 2009-2020., *Bitcoin*, <https://developer.bitcoin.org/>.
- [213] janoside, *Bitcoinexplorer*, <https://bitcoinexplorer.org/>.
- [214] L. Reid, *Gen ai in search: Let google do the searching for you*, <https://blog.google/products/search/generative-ai-google-search-may-2024/>, May 2024.
- [215] Y. Mehdi, *Reinventing search with a new ai-powered microsoft bing and edge, your copilot for the web*, 2023. [Online]. Available: <https://blogs.microsoft.com/blog/2023/02/07/reinventing-search-with-a-new-ai-powered-microsoft-bing-and-edge-your-copilot-for-the-web/>.
- [216] F. Landymore, *Just running chatgpt is costing openai a staggering sum every single day*, <https://futurism.com/the-byte/chatgpt-costs-openai-every-day>, Apr. 2023.
- [217] P. Team, *Why we are experimenting with advertising?*, <https://perplexity.ai/hub/blog/why-we-re-experimenting-with-advertising>, 2024.
- [218] R. Whitmam, *Google is quietly testing ads in ai chatbots*, <https://arstechnica.com/ai/2025/05/google-is-quietly-testing-ads-in-ai-chatbots/>, 2025.
- [219] F. Times, *Openai explores advertising as it steps up revenue drive*, <https://www.ft.com/content/9350d075-1658-4d3c-8bc9-b9b3dfc29b26>, 2025.
- [220] A. Rana and R. Nagda, "A security analysis of browser extensions", *arXiv preprint arXiv:1403.3235*, 2014.
- [221] A. Fass, D. F. Somé, M. Backes, and B. Stock, "Doublex: Statically detecting vulnerable data flows in browser extensions at scale", in *Proceedings of the 2021 ACM CCS*, 2021, pp. 1789–1804.
- [222] C. Obimbo, Y. Zhou, and R. Nguyen, "Analysis of vulnerabilities of web browser extensions", in *2018 International Conference on Computational Science and Computational Intelligence (CSCI)*, 2018. DOI: 10.1109/CSCI46756.2018.00029.
- [223] A. Aggarwal, B. Viswanath, L. Zhang, S. Kumar, A. Shah, and P. Kumaraguru, "I spy with my little eye: Analysis and detection of spying browser extensions", in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, 2018, pp. 47–61. DOI: 10.1109/EuroSP.2018.00012.

- [224] D. Bui, B. Tang, and K. G. Shin, “Detection of inconsistencies in privacy practices of browser extensions”, in *2023 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2023, pp. 2780–2798.
- [225] M. Xie, J. Fu, J. He, C. Luo, and G. Peng, “Jtaint: Finding privacy-leakage in chrome extensions”, in *Information Security and Privacy: 25th Australasian Conference, ACISP 2020, Perth, WA, Australia, November 30–December 2, 2020, Proceedings 25*, Springer, 2020, pp. 563–583.
- [226] H. Shahriar, K. Weldemariam, M. Zulkernine, and T. Lutellier, “Effective detection of vulnerable and malicious browser extensions”, *Computers & Security*, vol. 47, pp. 66–84, 2014.
- [227] S. Das, “A cloud-based framework for security analysis of browser extensions”, Ph.D. dissertation, Queen’s University, 2016.
- [228] A. Nayak, R. Khandelwal, E. Fernandes, and K. Fawaz, “Experimental security analysis of sensitive data access by browser extensions”, in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 1283–1294.
- [229] S. Bandhakavi, N. Tikku, W. Pittman, S. T. King, P. Madhusudan, and M. Winslett, “Vetting browser extensions for security vulnerabilities with vex”, *Communications of the ACM*, vol. 54, no. 9, 2011.
- [230] Q. Chen and A. Kapravelos, “Mystique: Uncovering information leakage from browser extensions”, in *Proceedings of the 2018 ACM SIGSAC CCS*, 2018, pp. 1687–1700.
- [231] Q. Xie, M. V. K. Murali, P. Pearce, and F. Li, “Arcanum: Detecting and evaluating the privacy risks of browser extensions on web pages and web content”, in *33rd USENIX Security Symposium (USENIX Security 24)*, 2024, pp. 4607–4624.
- [232] D. M. Martin Jr, R. M. Smith, M. Brittain, I. Fetch, and H. Wu, “The privacy practices of web browser extensions”, *Communications of the ACM*, vol. 44, no. 2, pp. 45–50, 2001.
- [233] Y. Ling, K. Wang, G. Bai, H. Wang, and J. S. Dong, “Are they toeing the line? diagnosing privacy compliance violations among browser extensions”, in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2022, pp. 1–12.
- [234] S. Karami, P. Ilia, K. Solomos, and J. Polakis, “Carnus: Exploring the privacy threats of browser extension fingerprinting”, in *27th Annual Network and Distributed System Security Symposium, NDSS 2020, San Diego, California, USA, February 23–26, 2020*, The Internet Society, 2020.

- [235] M. Gupta, C. Akiri, K. Aryal, E. Parker, and L. Praharaj, "From chatgpt to threatgpt: Impact of generative ai in cybersecurity and privacy", *IEEE Access*, 2023.
- [236] H. Li et al., "Privacy in large language models: Attacks, defenses and future directions", *arXiv preprint arXiv:2310.10383*, 2023.
- [237] Q. Li et al., "Llm-pbe: Assessing data privacy in large language models", *arXiv preprint arXiv:2408.12787*, 2024.
- [238] H. Kibriya, W. Z. Khan, A. Siddiq, and M. K. Khan, "Privacy issues in large language models: A survey", *Computers and Electrical Engineering*, vol. 120, p. 109 698, 2024.
- [239] B. Wang et al., "Decodingtrust: A comprehensive assessment of trustworthiness in gpt models.", in *NeurIPS*, 2023.
- [240] C. Peris et al., "Privacy in the time of language models", in *Proceedings of the sixteenth ACM international conference on web search and data mining*, 2023, pp. 1291–1292.
- [241] J. Yu, X. Lin, Z. Yu, and X. Xing, "Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts", *arXiv preprint arXiv:2309.10253*, 2023.
- [242] D. Bowen, B. Murphy, W. Cai, D. Khachaturov, A. Gleave, and K. Pelrine, "Data poisoning in llms: Jailbreak-tuning and scaling laws", *arXiv preprint arXiv:2408.02946*, 2024.
- [243] Z. Zhou et al., "Model inversion attacks: A survey of approaches and countermeasures", *arXiv preprint arXiv:2411.10023*, 2024.
- [244] B. Yan et al., "On protecting the data privacy of large language models (llms): A survey", *arXiv preprint arXiv:2403.05156*, 2024.
- [245] R. Staab, M. Vero, M. Balunović, and M. Vechev, "Beyond memorization: Violating privacy via inference with large language models", *International Conference on Learning Representations (ICLR)*, 2024.
- [246] N. Newman, R. Fletcher, K. Eddy, C. T. Robertson, and R. K. Nielsen, *Digital news report 2023*, 2023. [Online]. Available: <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023>.
- [247] MitmProxy, <https://mitmproxy.org/>.
- [248] Anthropic, *Claude*, <https://anthropic.com/claude>, 2025.
- [249] OpenAI, *Making requests*, <https://platform.openai.com/docs/api-reference/making-requests>.
- [250] E. Wallace, K. Xiao, R. Leike, L. Weng, J. Heidecke, and A. Beutel, "The instruction hierarchy: Training llms to prioritize privileged instructions", *arXiv preprint arXiv:2404.13208*, 2024.

- [251] O. for Civil Rights, *The use of online tracking technologies by hipaa covered entities and business associates*, <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/hipaa-online-tracking/index.html>, Jun. 2024.
- [252] U. D. of Education, *Protecting student privacy*, <https://studentprivacy.ed.gov/ferpa>.
- [253] MixPanel, *Session replay demo: Watch and learn*, <https://mixpanel.com/releases/session-replay/>.
- [254] FTC, *A few key principles: An excerpt from chair khan's remarks at the january tech summit on ai*, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/02/few-key-principles-excerpt-chair-khans-remarks-january-tech-summit-ai>, 2024.
- [255] FTC, *Artificial intelligence and the risk of consumer harm*, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2025/01/ai-risk-consumer-harm>, 2025.
- [256] E. Commission, *The ai act*, <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>, 2024.
- [257] P. AI, *Comet: A browser for agentic search by perplexity*, <https://www.perplexity.ai/comet>, 2025.
- [258] T. B. C. of New York, *Dia browser*, <https://www.diabrowser.com/>, 2025.
- [259] G. DeepMind, *Project mariner: Research prototype exploring the future of human-agent interaction, starting with browsers*, <https://deepmind.google/models/project-mariner/>, 2025.