

Analyzing GDPR Compliance of Named Data Networking

Casey Tran
New Mexico State University
Las Cruces, New Mexico
caseywt@nmsu.edu

Reza Tourani
Saint Louis University
St. Louis, Missouri
reza.tourani@slu.edu

Gaurav Panwar
New Mexico State University
Las Cruces, New Mexico
gpanwar@nmsu.edu

Satyajayant Misra
New Mexico State University
Las Cruces, New Mexico
misra@cs.nmsu.edu

Travis Machacek
New Mexico State University
Las Cruces, New Mexico
caprock1@nmsu.edu

ABSTRACT

The popularity of social media platforms, Internet of Things (IoT) devices, and the myriad smartphone applications have created opportunities for companies and organizations to collect individuals' personal data and monetize its sharing at a high rate. A standout example was the *Facebook–Cambridge Analytica* data-sharing arrangement (2018), which allowed Cambridge Analytica to harvest millions of Facebook users' personal data without their consent for political advertisement. In response to such overreach and privacy violations, the European Union introduced the General Data Protection Regulation (GDPR), which mandates data collectors to protect individuals' data privacy and provide the user more control over their personal data. Motivated by this growing interest in personal privacy, we analyze GDPR articles in the context of Named Data Networking (NDN). The context of interest is NDN as the network architecture in a service provider and we investigate GDPR-pertinent NDN features, including naming, caching, forwarding plane, and its built-in trust, for GDPR compliance and present insights on how such compliance can be built, when lacking. We also present experimental results showing compliance overheads and conclude by identifying potential future work.

CCS CONCEPTS

• Security and privacy → Privacy protections.

KEYWORDS

GDPR, Privacy, ICN, NDN, Security.

ACM Reference Format:

Casey Tran, Reza Tourani, Gaurav Panwar, Satyajayant Misra, and Travis Machacek. 2021. Analyzing GDPR Compliance of Named Data Networking. In *8th ACM Conference on Information-Centric Networking (ICN '21)*, September 22–24, 2021, Paris, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3460417.3482979>

1 INTRODUCTION

The permeation of technology in our lives is shifting the traditional definition of user privacy and rights of individuals over their personal data. Popular social media platforms, IoT devices, and numerous applications are harvesting remarkably large volumes of personal data on a daily basis in an attempt to monetize the data. Additionally, there are dedicated businesses and agencies (e.g., data brokers) that collect personal information on the healthcare, education, and finance sectors to sell in data markets. As a result,

we are experiencing an increasing number of privacy violations and unauthorized personal data sharing and tracking [2, 30].

In response to user privacy violations, policies and laws, such as *General Data Privacy Regulation* (GDPR) [11] and the *California Privacy Rights Act* (CPRA) [7] have been proposed to protect human digital rights. The GDPR, introduced in the European Union (EU), is the most extensive regulation with a collection of 99 articles on data protection, privacy, and data owners' rights. Its primary objective is to provide users control in the processing or storage of their personal data and requires compliance from companies storing/user user data. Compliance with the GDPR is non-trivial for most organizations, in part, due to the high cost of retrofitting legacy systems. For instance, the authors in [31] demonstrated that enabling fine-grained logging in a legacy storage system reduces throughput to only 5%!

To ensure GDPR compliance, EU has levied harsh fines—in its first 20 months of being enforced, EU has issued €114 million worth of fines to organizations including Google and Facebook. GDPR is aimed at regulating the way organizations and businesses handle the collected user personal data. With data increasingly being stored in the network infrastructure and internet service providers (ISPs) becoming content delivery networks (CDNs) and with the advent of edge computing where most of the computing of the data will happen at the network edge, GDPR is now applicable to the networking infrastructure—especially when it features handling of personal data, such as receiving, storing, distributing, and processing.

To the best of our knowledge, Cloud Service Providers have yet to be fined for GDPR violations. Companies such as Cloudflare seem to be already meeting the GDPR stipulations, stating explicitly that they do not sell or process personal data other than to provide their networking and content delivery services to customers. Moreover, they give customers control to access, correct, and delete their personal information on the service [6]; more specifically in its Data Processing Addendum as a *Data Processor* in [5]. Currently, several initiatives including the one from NetBrain offer solutions for automation of recurring data collection, indexing, and generation of required audit report for proving GDPR compliance for existing large IP-based enterprise networks [23]. This shows the push for GDPR compliance for IP networks in general.

In contrast to the existing networking architecture, which uses the best-effort IP protocol at the network layer, the Named-Data

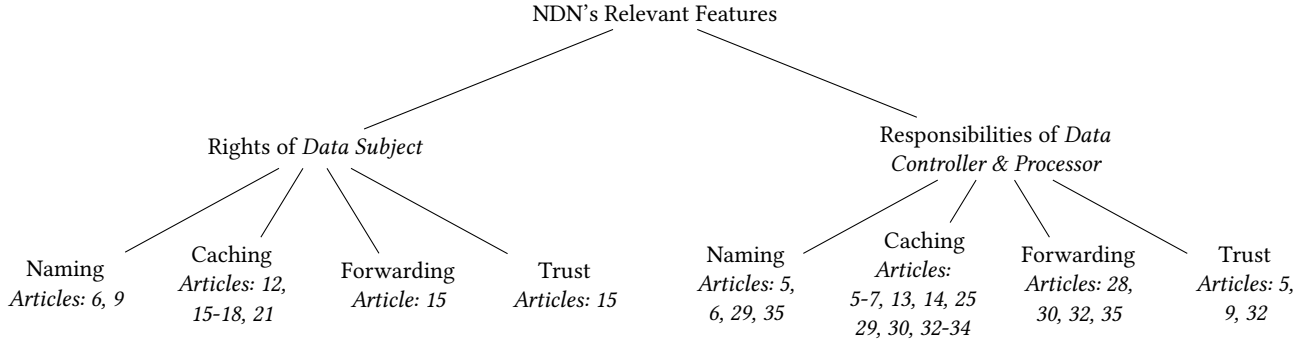


Figure 1: Overview of NDN's features classification and the relevant GDPR articles for each class.

Networking (NDN) [56] architecture promotes in-network intelligence by introducing the application's logic to the network layer. NDN also features a wide range of unique attributes, such as in-network processing and pervasive caching, making privacy and protection laws more relevant to NDN.

Motivated by these observations, and considering that GDPR's articles are often vague and difficult to interpret, this paper is the first in investigating the degree of compliance of NDN features with respect to GDPR. In particular, we identified naming, caching, stateful forwarding, and trust management as NDN's unique attributes that require scrutiny. Our investigation looks at NDN as the network architecture for the network infrastructure (ISP, Cloud, edge) for making the assessment of what it takes to make the network infrastructure GDPR compliant. This work aims to open up a broad view of GDPR for NDN, encouraging the creation of a body of future works from fellow researchers to explore each individual NDN feature in more depth and design approaches to ensure GDPR compliance. This is in contrast to the other dimension of investigation which investigates privacy compliance and violations of collecting user-identifiers based on application-level primitives with automated frameworks [2] and ad exchanges modeling [30].

In what follows we group the GDPR articles into (i) rights of Data Subjects and (ii) responsibilities of Data Controllers categories. Within each category, we discuss the aforementioned attributes and the pertinent GDPR articles as illustrated in Figure 1. Moreover, we performed overhead analysis from building GDPR compliance in NDN by implementing logging for caching, naming, and forwarding operations, as well as encryption of content in caches. Our initial assessments shows that these compliance operations add manageable overheads.

2 RELATED WORK

A few recent initiatives have analyzed GDPR compliance of various systems, such as storage systems and IoT networks [31, 49]. In [31], the authors studied the Redis storage system and concluded that retrofitting a legacy system for compliance may incur unbearable overhead—enabling fine-grained logging reduces throughput to only 5%. The authors in [49] discussed the trade-off between privacy and linkability of identity management in IoT services and concluded that the GDPR fundamentally challenges these design choices for less privacy invasiveness.

A recent study shows that most users are naive to data-linkage in the sense that they are willing to share their health and shopping data but not wealth data—while in fact all three are intertwined and linkable [13]. The goal for individuals to know how their data is being used, and in what ways it may be combined and linked to build their online profiles, is fundamental to GDPR. The authors in [21] call for a clear specification of the data collection purposes and its use.

Processing data with the specified intention is one of the primary GDPR objectives. As we shall see, processing does not only target storage of the data but also includes the transferring of personal data. While application developers/providers should consider GDPR compliance of the services they provide, compliance at the application does not eliminate the need for compliance at the network layer. For instance, if a data breach incident happens when users' identifiable data is cached in the network, the corresponding NSP should communicate the incident to the producer and regulators in a timely manner (Articles 33 & 34). Thus, in IP networks, a Software Defined Networking (SDN) system API has been proposed for realizing GDPR-compliance through data trails [48]. Specified actions, purposes, and protection abstractions permit an SDN Controller to monitor the personal data flows over the network; yet, it only addresses compliance records on processes within the SDN without means for protection leaving the network (e.g., via data breach) [48]. Comprehensive security standards by GDPR raises a critical concern on the structural shift which data collectors and processors must become aware of. Hence it is noted that scaling for security and privacy not only calls for retrofitting systems built on conflicting design principles, components and practices, but also addressing compliance for new technologies like NDN in a new security-centric environment [32]. Along this line, in this paper we investigate GDPR compliance of NDN.

3 GDPR OVERVIEW

Out of the 11 chapters of GDPR's 99 articles, we focus on the first four chapters due to their relevance to NDN. *Chapter 1: General Provisions*, discusses GDPR's aim, scope, and definitions (Articles 1-4). *Chapter 2: Principles*, touches on *Processing and Protection* rules on *Personal Data*. (Articles 5-11). *Chapter 3: Rights of the Data Subject*, is focused on the *Data Subject* rights, including the right to

be forgotten, right to rectification, and right to restriction of *Processing*. (Articles 12-23). Finally, *Chapter 4: Controller and Processor*, discusses the necessary security measures to be deployed at the *Data Controller* and *Data Processor* and the obligations of these entities pertaining to *Personal Data* (Articles 24-44).

3.1 Definitions

GDPR's *Article 4* includes a set of 26 basic definitions. Here, we first present a subset of relevant definitions and then elaborate on them in the context of NDN.

- **Personal Data** is any information relating to the identified or identifiable person; *Article 4(1)*.
- **Data Subject** is an identifiable or identified natural person who reveals/discloses *Personal Data* for processing such as a name or location, with regard to one or more physical, genetic, mental, or economic factors; *Article 4(1)*.
- **Processing** refers to any operation on personal data whether for agreed upon (lawful) or unspecified (unlawful) purposes under the data controller or data processor.
- **Data Controller** is an entity that holds the initial and direct contact to the *Processing* agreements between itself and the *Data Subject*; *Article 4(7)*.
- **Data Processor** is an entity that handles/processes the *Personal Data* on behalf of the *Data Controller*; *Article 4(8)*.
- **Protected** is referring to personal data that has been treated with the necessary safeguards such as pseudonymization and encryption to prevent it from affecting the fundamental rights and freedoms of an individual. Protection also obtains from exercising the rights of data subjects.

In the context of NDN, we categorize the data owner (producer in the majority of cases) as *Data Subject* and the network service provider (NSP) as the *Data Controller* since they maintain a direct relationship. Moreover, we categorize the NSPs' components, such as routers and proxies, which participate in the data dissemination, as *Data Processors*. By virtue of the fact that *Data Processors* run the NDN's forwarding daemon (NFD), we define legitimate/legal *Processing* as the fundamental NFD operations. Such operations include, but are not limited to, name-based look-ups, Interest aggregation, content caching, and the rudimentary forwarding strategy. For a better assessment, we revisit these definitions and justify legal *Processing* for each NDN feature.

3.2 Data Controller & Processor Relation

Regarding the relation between the *Data Controller* and *Data Processor*, *Article 28(1): PROCESSOR*, requires the *Data Controller* to only select the *Data Processors* that can ensure legal *Processing*—providing guarantees that *Processing* meets the agreement and protects *Data Subject's* rights. While for NSPs using the NDN architecture offers built-in trust, data integrity, and provenance verification, it falls short in providing data confidentiality, privacy, and access control. As mentioned in Sections 4.4 and 5.4 of the paper, the NDN built-in trust feature itself requires compliance considerations with GDPR regulations.

Additionally, NDN enables in-network intelligent decision-making by introducing the application's logic to the network layer. As such,

the network layer has access to identifying information not available with the traditional network layer (e.g., IP) in the stack. This makes it imperative for network layer compliance with GDPR, even if the application layer may provide privacy guarantees. For instance, if a data breach incident occurs when users' data is cached in the network (cache stores of the routers), the corresponding network service provider (NSP) should communicate the incident to the producer and regulators in a timely manner (*Articles 33 & 34*). Such cases require the direct involvement of NSPs to *Data Subjects* rather than applications or application layer providers.

Moreover, *Article 28(2 and 4)* forbids a *Data Processor* from engaging another *Data Processor* in the *Processing* without authorization and notice from the *Data Controller*. The implications of *Article 28(1, 2, and 4)* in the NDN context is that the NSP should restrict the collaboration of its routers beyond data delivery purposes. If routers' collaboration is inevitable for other *Processing* (e.g., coordinated caching or QoS-aware forwarding), the *Data Subject* agreement should include such terms.

4 RIGHTS OF DATA SUBJECT

The rights and freedoms of data producers impact how transparent and identifiable their *Personal Data* will be in cyberspace. For *Data Subjects* to freely use the Internet, *Data Processors* should ensure that *Personal Data* may not be used for profiling operations without *Data Subjects'* consent while also preventing against *Personal Data* against cyber threats. In order to meet these requirements, GDPR defines rights that *Data Subjects* may exercise to access, control, and receive information about their *Personal Data* from *Data Controllers* in a clear readable format. In this regard, we identified eight articles pertaining to the aforementioned rights. We summarize these articles and identify the corresponding NDN features in Table 1.

Under *Article 6*, *Data Processors* are prohibited from processing of *Personal Data* in a way that is incompatible with the original purpose of data collection. The definition of lawful processing depends on the context and any processing should be done with data owners consent. In the following subsections, we define lawful processing in the context of NDN. Regarding the processing of sensitive information, *Article 9* establishes that the processing of revealing data like racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, as well as the processing of data concerning health, sex life, genetic, or biometric data is prohibited, unless with *Data Subject's* consent. If processing such data is inevitable, the *Data Controller* shall consider any privacy implications that may need to be mitigated. This restriction holds whether the data is explicitly or even implicitly disclosing information in these special categories. For instance, processing data related to shopping habits may invade privacy as it could reveal information about *Data Subject's* state of health [10].

Article 12 requires that all communication between *Data Controller* and *Data Subject* be clearly and easily intelligible using plain language. Moreover, such communication should be in a popular format, whether in writing, electronically, or even orally, as requested by the *Data Subject*. Here, GDPR requires *Data Controllers* to provide information that is comprehensible to the *Data Subject*. However, it

Table 1: The GDPR articles pertaining to the Data Subjects' rights along with the corresponding NDN's features—naming (N), caching (C), forwarding plane (F), and built-in trust (T).

No.	Article Title	Article Essence	Relevance
6	Lawfulness of processing	Personal Data used only on consent or public interests	N
9	Processing of special data categories	Any revealing data under “special categories” is forbidden	N
12	Transparent information communication	Interactions by Controller to Data Subject is intelligible and readable	C
15	Right of access by the data subject	Provide timely access to data, its processing purposes, and compliance	C, F, T
16	Right to rectification	Data Subject can request Controller to rectify inaccurate data	C
17	Right to erasure	Data shall be punctually removed from the system upon request	C
18	Right to restriction of processing	Data Subject can restrict data processing by Controller if necessary	C
21	Right to object	Automated profiling of Data Subjects is not allowed	C

has been argued that making such information “comprehensible” for the non-technical *Data Subjects* is a challenging task [17].

We identified the following five articles pertaining to *Data Subjects* practicing their autonomy of information and privacy. First, *Article 15* requires that individuals have the right to know (i) the purposes, logic, and existence of any automated processing procedures; (ii) the recipients (categories of recipients) of their *Personal Data*; and (iii) the period and criteria for which the *Personal Data* will be stored. Next, *Articles 16, 17, 18 & 21*, provide *Data Subjects* with the rights to control their *Personal Data* by requiring *Data Controllers* to comply with requests to rectify, erase, restrict, and object to the processing of specific data.

4.1 Naming

From the aforementioned articles, we identified two articles which have implications on naming schemes used in NDN (*Articles 6 & 9*). Although this section is dedicated to the rights of *Data Subjects* as producers, data naming is also a responsibility in their own right. They should avoid including sensitive information in data names as a precaution against profiling by NSPs as specified in *Article 9*. This would help in scenarios where the terms of protection are overlooked by producers while NSPs are profiling or performing unnecessary analysis on naming schemes. For instance, in 2020, a European high school faced a €4,000 penalty for unknowingly publishing health data and sensitive information of teacher rankings due to their application semantics [39].

According to *Article 6*, any processing on the Interest and data packets' names beyond CS, PIT, and FIB look-ups shall be considered as unlawful processing in terms of GDPR unless the identified data producer has previously agreed to such processing. Examples of these unlawful processing (in the absence of producers' agreement) includes user profiling, name-based censorship [15], filtering [33], Distributed Denial of Service (DDoS) countermeasures using per interface or producer rate limiting [1], and monitoring interface-producer statistics to augment users' quality of experience [24]. In summary, any data processing on the names aiming at data flow classification could be perceived as unlawful as it may be used to affect the producer's privacy.

4.2 Caching

There are four articles pertaining to the rights of *Data Subjects* for their cached *Personal Data*. In the context of caching, we define lawful processing as solely insertion and deletion of named-data

to/from the routers' temporary storage while satisfying requests with the cached data.

Following *Articles 15 & 21*, *Data Subjects* can request information on their PD usage and object to any processing that falls outside of the original purpose. For instance, NSPs should inform the producers or obtain their agreement if they decide to employ cross-domain collaborative caching (e.g., Netflix and Comcast working together) [16, 19] or use the cached *Personal Data* to mitigate locality disruption attack [27, 52].

Many collaborative caching mechanisms have been proposed, in which routers in the same domain share their cache states [57]. However, if the *Personal Data* traverses the network to a region where access should be denied, such as outside of the EU's jurisdiction, then the NDN forwarding strategy without geographic blocking would be susceptible to violations against GDPR *Articles 5 & 6*. This is akin to the sanction against Aegean Marine Petroleum Group whose servers holding *Personal Data* were shared with others, resulted in exfiltration of *Personal Data* due to a failure to separate sensitive data from shareable data [38].

Articles 15–18 protect the rights of producers and their *Personal Data* by enforcing NSPs to comply to the requests to rectify, erase, or restrict the *Personal Data* and the *Processing* on it. These actions are meant to maintain the lawfulness of data usage by requiring timely deletion and indexing features for the cached *Personal Data*, in response to *Data Subject's* requests. *Article 17* (the right to be forgotten) is perhaps more pertinent when considering persistent data storage. With data caching and mobility, *Data Controllers* should keep track of the locations of all *Personal Data* copies to be able to comply with *Personal Data* erasure requests. Maintaining the list of all the locations that a *Personal Data* has been cached is a fairly challenging task, which may lead to penalties if *Data Controllers* fail to prove pervasive *Personal Data* erasure in a timely manner. For instance, the Bank of Cyprus Public Company Ltd had lost the ability to access the whereabouts of the *Data Subject's* insurance contract. This event resulted in the bank's failure to uphold *Article 15* and, hence, failure to grant *Data Subject* control over their data. According to the case study in [44], the measured fine of €15,000 was due to the fact that the total violations including *Articles 5, 15, 32 & 33* deemed it more severe.

In another instance [41], the *Data Controller* failed to adequately respond to the *Data Subject's* requests for data erasure as the internal policies concerning deletion and retrieval of *Personal Data* were not properly communicated to the *Data Subject*. In turn this resulted in a fine in the percentage of the annual profit of the *Data*

Controller. Since the policies were not properly informed, this was also a violation of data-transparency according to *Article 5*. In these contexts, compliance requires NSPs to design their systems to adhere to lawful processing of cached *Personal Data*. We discuss this in more detail in Subsection 6.2.

4.3 Stateful Forwarding Plane

In the context of NDN's stateful forwarding plane, we defined legal processing as PIT and FIB operations (e.g., insertion, look-up, and deletion). Any additional processing beyond these operations on *Personal Data* without express permission from the *Data Subject* can be categorized as unlawful. Consent must therefore be predetermined before an NDN network considers using edge routers for enforcing access control on *Personal Data* at the strategy layer [37] on behalf of content providers or custom forwarding strategies requiring performance measurements (e.g., latency, packet drop, and bandwidth consumption) of routers' interfaces using *Personal Data*. As specified in *Article 15(1)(c)*, producers should be informed of the recipients of their *Personal Data*. However, lack of consumer identity in Interest request packets and Interest aggregation along with its benefits and disadvantages [25] hinder this. To alleviate the impact of request aggregation (with respect to data access logging), cross-system mechanisms can be designed to collect data access records through router cooperation in order to share their monitoring and tracing of data across system. However, such systems will still fall short so long as Interest packets do not leak the consumers' identities.

4.4 Built-in Trust

With regard to NDN's built-in trust, we identified one GDPR article that pertains to the rights of producer for the protection of *Personal Data* involving digital certificates. In NDN's *trust schema* [54], the data authentication process follows a set of trust rules and forms a chain-of-trust ending at a trust anchor. Using the trust schema, however, discloses the overall trust model which may include sensitive information (e.g., organizational structures and relationships). To illustrate, consider a user (Charles) interested in publishing an article for CNN news. The published article includes the certificate name `"/CNN/News/PoliticalNews/author/Charles"`. This certificate, in turn, includes the division's certificate `"/CNN/News/PoliticalNews"`, which includes the trust anchor's name `("/Gov/Australia/CNN")`. As illustrated in this example, following a data packet's chain-of-trust may reveal the association of the author to the government that implicitly authorized the publication. Thus, the author's certificate is identifiable and linkable to other domains and identities.

Next, *Article 17* relates to producers when they no longer wish to be identified or linked to some data linked via a trust model. In this regard, the producer may demand its certificate (assuming certificates are considered *Personal Data*) to be erased and forgotten from the network. Erasing certificates allows producers to erase identifying information about them, assuming that certificates can be used to infer what type of data the producer is creating. Evicting the certificates from the network caches, along with any existence of malicious producers who reject the requests for their certificates, will negatively impact NDN's built-in trust. Thus, making data

packets' authenticity and integrity unverifiable. To avoid such circumstances, *Article 17(3)* lists the conditions under which erasure may not be granted. These conditions can be summarized as: (i) freedom of expression and information; (ii) legal obligation; (iii) purposes of public interest such as in-network content poisoning countermeasures [14, 50]; and (iv) conduction of legal claims.

5 RESPONSIBILITIES OF DATA CONTROLLER & PROCESSOR

As the complex and lengthy terms and policies of organizations often discourage many non-technical citizens from reading them [22], the GDPR provisions a significantly greater amount of responsibilities for *Data Controllers* to be held accountable in protecting *Personal Data*. In this section, we will first introduce the 13 articles pertinent to the responsibilities of *Data Controller* and *Data Subject* with respect to processing of *Personal Data*. After reviewing these articles (summarized in Table 2), we elaborate on their implications on NDN's features.

Articles 5 & 6 define the lawfulness criteria and constraints for *Data Controllers* to foster compliance as discussed in Section 4. Moreover, *Article 7* mandates *Data Controllers* to prove receipt of consent to the *Personal Data* processing. It further allows *Data Subjects* to demand their consent to be revoked, which should result in *Data Controllers* purging *Personal Data* in a timely manner.

Article 13 lists the information that a *Data Controller* must provide to *Data Subjects* to inform them how their data will be processed and their rights to request erasure, restriction, and objection to *Personal Data* processing. *Data Controller* organizational contact points and its Supervisory Authority contact information are required to be disclosed to *Data Subjects* to lodge complaints or exercise other rights as listed in Section 4. As stated in *Article 14*, *Data Controllers* should provide such information as stated in *Article 13* to *Data Subjects* even if they did not collect the *Personal Data* directly from the *Data Subjects*.

Article 25 mandates the *Data Controllers* to implement and practice *Personal Data* protection principles and integrate safeguards to protect *Data Subjects*' rights. *Article 28* requires *Data Processors* under *Data Controllers* to comply with *Data Subject* rights when handling *Personal Data*, including all security measures listed in *Articles 32-35* and with sufficient guarantees to implement technical and organisational measures. Technical measures refer to pseudonymization and encryption while organizational measures include performing data protection impact assessments for operations that may pose a risk to *Personal Data* [29].

Relating to the transfers of liability, *Article 24* states that any *Data Processor* (other entities) under the control of a *Data Controller* with access to *Personal Data* should follow the *Data Controller*'s instructions. In any of these cases, *Article 30* enforces the *Data Controllers* and their authorized *Data Processors* to maintain records of *Processing* activities concerning *Personal Data*. Such records should contain information, such as *Data Controllers*, purposes of processing, and recipients' categories. *Articles 32-35* include the expected security provisions for GDPR compliance. *Article 32* sets the regulation for the security measures that *Data Controllers* have to implement and practice, including advanced confidentiality, integrity, and availability of the collected data. Moreover, *Articles 33 & 34*

Table 2: The GDPR articles pertaining to the Data Controllers' responsibilities along with the corresponding NDN's features—naming (N), caching (C), forwarding plane (F), and built-in trust (T).

No.	Article Title	Article Essence	Relevance
5	Principles of processing	Data is processed transparently and only on specified purposes	N, C, T
6	Lawfulness of processing	Data used only on consent or public interest	N, C
7	Conditions for consent	Consent must be demonstrated for the processing	C
13	Conditions for data collection	Inform Data Subjects how and where their Data will be used	C
14	Conditions for indirect data collection	Where Data is received indirectly, Article 13 applies	C
25	Protection by design and by default	Safeguard and restrict access to Data	C
28	Processor	Processor shall provides sufficient security guarantees	F
29	Processing under controller's authority	Any entity under Data Controller shall only use Personal data	N, C
30	Records of processing activity	Logs are needed to audit all operations over personal data	C, F
32	Security of processing	Implement security measures pertaining to level of risk	C, F, T
33,34	Data breach notification	Parties shall be punctually informed of any risks or detected breaches	C
35	Data protection impact assessment	Any risk in processing sensitive data is subject to this protocol	N, F

require the *Data Controllers* to respectively notify its Supervisory Authority and *Data Subjects* of *Personal Data* breaches within 72 hours and without undue delay. Finally, *Article 35* describes the responsibility of *Data Controllers* to carry out assessments [51] of their operations and seek advice from designated data protection officers to demonstrate GDPR compliance.

5.1 Naming

In this subsection we list the GDPR articles that outline the responsibilities of the *Data Controller* and *Data Subject* in the context of naming. While it is obvious that the applications and the services that generate an Interest or Data packet should avoid using identifiable names, compliance with GDPR regulation is not only about the application layer but also the network layer as it obtains, processes, and stores such personally identifiable data. In this context, legal processing may be defined as using names for basic forwarding operations—cache, PIT, and FIB look-ups as well as executing routing protocols [12]. We refer the readers to several prior results on name obfuscation and privacy for naming studied in the general context (not GDPR) [9, 28, 34, 35].

Based on *Article 6*, using the packet's name to process outside the bounds of the aforementioned NDN operations is unlawful such as in DDoS mitigation schemes which use name prefix for rate limiting, name-based censorship, and QoE aware forwarding. But since not all unlawful processing is malicious (e.g., DDoS mitigation), NSPs wishing to use packets' names for improving their consumers' QoE or deploying customized network protection schemes must inform and obtain consent from consumers.

NDN DeLorean [55] is an example of a non-malicious processing that requires producers' consent as it has implications on the producers' right to be forgotten (*Article 17*). NDN DeLorean aims at enabling data verifiability beyond the certificate expiry period by sealing the existence of signed named data to a chronological and immutable ledger. We argue that the applications which permanently store *Personal Data* in an immutable fashion, such as NDN DeLorean or distributed ledger based applications, cannot comply with the GDPR as producers' request for *Personal Data* erasure breaks the applications' premise unless blockchain redaction solutions [8, 26] are implemented by *Data Controllers*. Such rejection of an application to dereference outdated or stale certificates could face a penalty much similar to Google Belgium SA where they

were unable to dereference outdated articles by a particular *Data Subject* [42]. Finally, we note that NDN is susceptible to security flaws [36]. Hence, following *Article 35*, NSPs have to carry out a Data Protection Impact Assessment when adopting new features to assess and demonstrate compliance.

5.2 Caching

We identified twelve articles on the responsibilities of NSPs with respect to cached *Personal Data*. In the context of content caching, we define the legitimate processing as cache insertion, deletion, and lookup for satisfying the subsequent requests. For an NSP or its authorized entities (e.g., routers and proxies), any operation beyond what is legitimate requires the producer's consent as regulated in *Article 6*. For instance, *Personal Data* pre-fetching prior to consumers requests [18] (similar to content delivery networks functionality) is unlawful unless in agreement with the producer.

Personal Data traversing independent NSPs increases its possibility of being cached across multiple networks and increases the risk of it being used unlawfully. Under *Articles 13* and *14*, all the collaborating NSPs with the intention of caching *Personal Data* should follow the particular guideline for *Personal Data* processing and protection while informing the producers. Such a guideline should be defined by the NSP (of course with producer's consent) that directly obtains the *Personal Data* from a producer. Further precedent from [40] shows that transferring *Personal Data* to third parties requires adequate security measures especially in the processing to prevent against data-processing misuses.

We point out that NDN is currently missing a few functionalities that are needed to comply with *Articles 15*, *30*, *32*, and *34*. First, to prove and maintain accountability and compliance with *Articles 5* and *6*, there lacks a logging mechanism to record all operations performed on *Personal Data* as specified in *Article 30*. Such logging can be used for satisfying compliance requirements like when the users' data need to be removed from the NSP's caches or in case of a data breach incident, in which the NSP should detect the scope of the breach and inform the data owner and regulators. Such requirements cannot be satisfied by the application providers since they don't have any view as to how the data traverses the NSP's network. NDN's inherent security requires the data packets to be signed for integrity and provenance purposes, and hence, comply with *Article 25*. However, as specified in *Article 32*, NSPs should

protect cached *Personal Data* (beyond integrity and provenance) and guarantee an appropriate level of security to mitigate the risks of caching.

While the GDPR's requirement for security is rather vague, we argue that if *Personal Data* are going to be cached, its encryption might be more desirable. This is particularly important since data encryption can reduce the severity and cost of data breaches;¹ breaches should be communicated to the producer and regulators in a timely manner as stated in *Article 33 & 34*. For example, in [45], *Data Controllers* were penalized as they failed to fulfill data breach obligations simply due to a lost USB flash drive containing *Personal Data*. However, *Personal Data* encryption before caching and its decryption for each subsequent request (every cache hit) may impose computational overhead on the NSPs. Thus, NSPs might be in favor of avoiding *Personal Data* caching.

To avoid performance degradation as a result of in-network logging, efficient mechanisms need to be designed to avoid redundant logging as data traverses the NSP's infrastructure yet allowing reconstruction of complete data trail. Similarly, a more efficient data protection framework can be designed to encrypt the data as it enters the NSP's network and decrypt it as it leaves the NSP's network rather than per cache hit encryption/decryption. Thus, drastically reducing the cost of data protection at rest.

5.3 Stateful Forwarding Plane

For our discussion on the stateful forwarding compliance with the GDPR, we have identified four articles. We note that the legal processing, in this context, includes the look-ups, insertions, and deletion to/from CS, PIT, and FIB.

In reference to *Article 28*, NSPs should only permit the *Data Processors* (entities such as routers and proxies) that can meet the expected security standards (to keep the NSPs compliant according to *Article 32*) to be involved in the *Personal Data* forwarding process. In particular, routers and proxies involved in *Personal Data* forwarding should satisfy the demands of confidentiality, integrity, availability, and resilience under the same provisions that an NSP is subjected to as in *Articles 32-35*.

Moreover, they should collect, maintain, and provide records of their packet forwarding as stated in *Article 30*, which may incur significant cost and degrade packet forwarding throughput in order to comply with and avoid the penalties of GDPR. Defining unique node identifiers (an identifier can be a node id with the NSP's prefix) under NDN will be a necessity for these records which we note does not contradict the architecture's fundamentals since the content producers as *Data Subjects* remain decoupled from the location but not the NSP nodes themselves.

5.4 Built-in Trust

Regarding the NDN's built-in trust compliance with GDPR, we identified two articles pertaining to the challenges that arise from NDN's signed data packets. In NDN, the packet signature and the certificate name expose sensitive and identifiable information concerning the structure of organizations and trust relationships as we discussed in Subsection 4.4).

¹Data breach fines can run up to €20 million or 4% of a company's global annual turnover.

In this regard, one possible threat to consumer privacy would be profiling the producers through their named certificates which consumers interact with in the data packets. We note that even using semantically meaningless [3] or encrypted data names [47] for data packets cannot prevent such threat as it relies on producers' potentially static certificates. Such information can be used to infer a consumer's ethnic origin, political opinions, religious, state of health, and sexual orientation as stated in *Article 9*.

Moreover, signing data packets hinders producers' privacy and anonymity [4] and is in stark contrast to *Article 32(1)*, which requires *Personal Data* pseudonymization—processing *Personal Data* such that it cannot be linked/attributed to a particular *Data Subject* without supplementary metadata. As shown in [43], the Mayor's Office of Kescemét was penalized for the unlawful disclosure of the details and *Personal Data* of a whistleblower. While approaches have been proposed to achieve k-anonymity for producers through various signature schemes [4, 28], further investigation is needed to design a holistic anonymous data publication framework.

Despite the privacy and anonymity challenges, which arise from signed certificates, these features allow the NSPs to remove inaccurate *Personal Data* from the network in compliance with *Article 5(1)(d)*. In this context, inaccurate *Personal Data* is a data packet with an invalid signature. To this end, the NSPs should communicate the aim and scope of such processing—using certificates to perform in-network signature verification as a content poisoning countermeasure (refer to discussion in Subsection 4.4)—to the producers to obtain their consent.

6 NDN GDPR COMPLIANCE

In this section, we review the imperative features for NDN to be GDPR-compliant. One such feature is logging NFD critical operations. To obtain a lower bound on latency overhead of logging, we ran a set of experiments using a simple testbed by connecting two machines running NFD over a wireless access point. This paper assesses the fundamental compliance requirements on a per node basis. There are other aspects to assess: design of a scalable logging mechanisms for network-wide deployment, name obfuscation solutions for naming privacy and their compliance for GDPR. Such evaluations will on one hand assess the challenges in NDN arising from multiple paths, node specific forwarding and caching strategies and the inherent system dynamics, and on the other hand, they will also look at the compliance of proposed privacy and security solutions.

We ran the producer application on a desktop running Ubuntu 18.04, 4.0 GHz single core CPU with 8 GB RAM and ran the consumer's application on a MacbookPro that runs OSX Catalina, 1.4 GHz quad-core CPU with 8 GB RAM. In our experiments, we used ndn-cxx library v0.70.

6.1 Personal Data Distinguishability

As defined in *Article 4(1)*, *Personal Data* is any piece of information relating to an identified/identifiable natural person who can be identified by a unique identifier such as a name, location information, or combination of genetic, economic, and cultural factors. Such *Personal Data*, to be treated in accordance with the GDPR, should

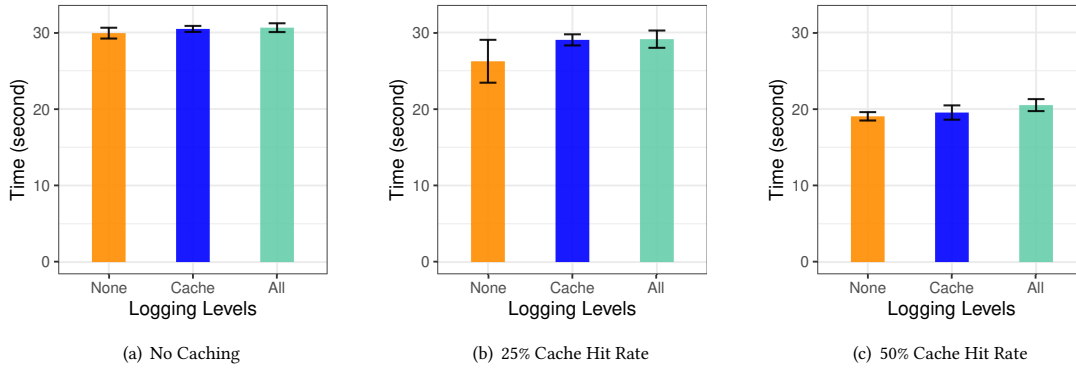


Figure 2: The impact of multi-level logging on the consumer’s end-to-end latency. The logging levels include no logging (None), logging cache operations (Cache), and logging all NFD operations (All); time is total running time.

be identified by the NSPs and their routers and proxies. Thus, features need to be integrated into the Data packet’s format to enable *Personal Data* distinguishability.

Conventional NDN naming suggests hierarchical human-readable names. Using the semantics in the naming scheme, NDN promotes flexible packet forwarding [24], augments security [58], and simplifies content caching. Thus, a naive approach would be integrating an identifier into *Personal Data* names, which is known and accepted across all NSPs. Nevertheless, augmenting the *Personal Data* names with the identifier require a unified naming scheme to specify which name component is the identifier. This is challenging due to the consistency requirement across all producers and application developers. Such unified naming scheme will further put the application developers in charge of deciding the *Personal Data*, while in applications, producers should make such decisions. A more viable approach would be including a flag in the NDN packets, allowing the producers to make fine grained decisions on per packet basis.

Although protecting *Personal Data* through distinguishability provides attackers knowledge of which data is personal, not distinguishing personal data keeps the NSPs blindsided and prevents them from taking proper measures in protecting such information, which may result in more devastating consequences. For instance, an NSP might not report a data breach if it does not know the leaked data includes personal information or might cache a piece of personal data in plaintext. Name obfuscation schemes can be applied alongside distinguishability techniques to further protect data privacy.

6.2 Multi-level Logging and Monitoring

In compliance with *Articles 5 & 30*, NSPs need to provide auditable and trustworthy records of all the operations that have been executed on *Personal Data*, including but not limited to caching and forwarding. This requires the design of a system-wide multi-level logging mechanism for secure collection, aggregation, and access enforcement to logs across all entities of an NSP. Thus, we conducted a set of experiments (averaged over ten runs) to assess the impact of logging on the system performance and user experience. In these experiments, the consumer application requests 1000 data packets from the producer in a stop and wait fashion. The choice

of stop and wait application is for illustrating the performance metrics, such as latency and resource utilization, and is not intended to represent any real world application. In these experiments, we enabled caching on the producer to minimize the impact of cache lookup operations on our measurements.

In the first set of experiments, we measured the total content download time under three logging levels—no logging (None), logging caching operations (Cache), and extensive NFD logging (All)—for 0%, 25%, and 50% cache hit ratios. As shown in Figure 2, logging all operations slightly increased the total content download time across all experiments. We also noticed that the content retrieval time decreases as the cache hit ratio increases, which is in part due to the fact that the producer needs to generate and sign a smaller number of data packets in scenarios with higher cache hit ratios. In summary, we note that the latency cost of logging in the current NFD version is negligible.

We also assessed NFD’s CPU utilization during our experiments (Figure 3). In calculating the CPU utilization, we gathered information from the process information filesystem by sampling the number of jiffies the process accumulated in user mode and kernel mode. A jiffy is a timer interrupt which is used by the kernel as a mechanism to keep track of time intervals. These intervals are incremented whenever the process has time on the CPU. We sampled NFD’s jiffy counts in the user and kernel spaces before and




Figure 3: CPU utilization across all logging levels.

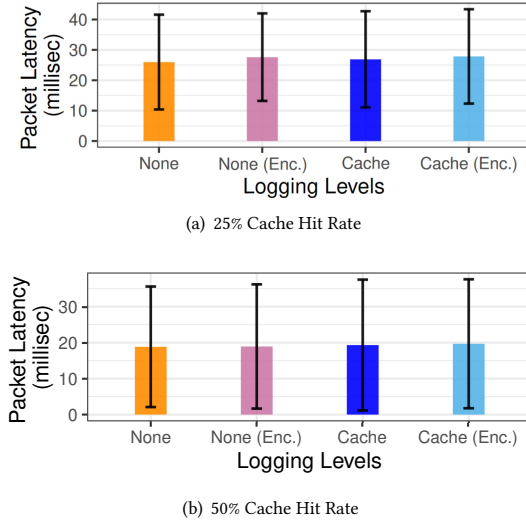


Figure 4: The impact of data decryption during cache hit on the consumer's end to end latency.

after each experiment. We then divided the total number of *jiffies*, accumulated by the NFD process, by the total number of *jiffies* accumulated by the CPU for that time interval. This gave us an approximation of the CPU utilization of NFD for each experiment. One can observe that enabling extensive NFD logging (“All”) has increased the CPU utilization to 8%, up from 4% where there is no logging (“None”). This increase in CPU utilization is consistent across all cache hit rates. Cache operations logging increased CPU utilization by 1%.

6.3 Encryption/Decryption of Cached Data

Personal Data protection is another critical feature for GDPR compliance; particularly when the controller is storing *Personal Data*. *Articles 25 & 32* establish that access to *Personal Data* should be protected by design and appropriate security measures should be implemented in accordance to the risk of storing *Personal Data*. NDN's built-in security and caching provide *Personal Data* integrity, provenance, and availability, leaving confidentiality as the potential risk. For confidentiality, cached *Personal Data* should be encrypted—the routers encrypt the *Personal Data* before cache insertion and decrypt it per cache hit. We assume that the application generating the data encrypts it before transmission and hence, the NSP's data encryption is another layer of protection regardless of what the application layer service has done concerning the *Data Subjects*.

To evaluate the overhead of providing cached *Personal Data* protection, we ran a set of experiments by extending NFD with data encryption (using AES-128 in CBC mode)—data is encrypted before caching and decrypted per cache hit. We assess the packet retrieval latency for the cached data at 25% and 50% cache hit rates, with no logging and no encryption, no logging but use of encryption, logging of cache operations with content not encrypted, and logging of cache operations with content encrypted, respectively, shown in Figure 4. Encrypting/decrypting the protected cached data introduced negligible latency overhead across the board. This is due to the small size of the packets (1kB) and given the native acceleration of AES operations (Intel AES-NI). We also noticed that




Figure 5: CPU utilization of Different Configurations.

using bigger cache sizes resulted in lower retrieval latency. This is because the producer now has to generate and sign a smaller number of data packets. Overall, we argue that protecting cached data for GDPR-compliance imposes an acceptable overhead. Figure 5, which shows the corresponding CPU utilization overhead for the different configurations and different caching rates reiterates our assertion of low overhead of compliance.

7 CONCLUDING DISCUSSIONS AND INSIGHTS

7.1 Consumer Identification

The lack of consumer identity in Interest packets is perhaps one of the NDN's features most against GDPR compliance. As in *Article 15(1)(c)*, *Data Subjects* have the right to know the recipient (categories of recipients) of their *Personal Data*. However, NDN Interest packets do not carry consumers information. So, when a *Data Subject* seeks that information from a router, the router has no way of identifying which locally logged Interest corresponded to which *Data Subject* and hence, cannot satisfy the question. We argue that including some form of consumer identification that does not exactly identify the consumer in the Interest packets will not violate NDN's principles. In fact, the benefits of including consumers specific metadata in requests has been discussed in the literature for various applications, such as access control delegation [37, 53] and in cyber-physical systems [20, 35]. These applications use pseudonym based consumer identification.

Moreover, a group of applications such as online banking may require the Interests to be signed by the consumers for authentication and non-repudiation. This is also the case when the Interests include application related information (*Application-Parameter*). Such signed Interests should provide the corresponding certificate name (similar to the Key Locator field in data packets), allowing the requests' recipients to validate the requests' signatures and integrity. The certificate name can form the consumer's unique identifier.

7.2 Timely Request Processing

The GDPR's *Articles 16-18* and *21* are among the most prominent articles, which grant *Data Subjects* rights to request rectification

and erasure of *Personal Data* as well as restriction and objection to processing. *Data Controllers*, in turn, should comply with such requests without undue delays. However, GDPR does not provide any timeline for the *Data Controllers* to act upon the *Data Subjects*' requests. In this scenario, the GDPR vagueness may benefit *Data Controllers* and disadvantage *Data Subjects*. The authors in [31] defined *Real-time* compliance, in which the *Data Controllers* have to satisfy the *Data Subjects*' requests on demand with quick turnaround (GDPR does not clearly stipulate a time period); the compliance will be *eventual*, if it cannot be satisfied in real-time. The authors concluded that real-time compliance may impose significant overhead on the *Data Controllers*. Moreover, the cost of real-time compliance has a direct relationship with the size of *Data Controllers*' organizations.²

In our context, ISPs are among the *Data Controllers* that should be compliant. Considering the size of ISPs, we envision significant overhead (both computation and communication) in providing real-time compliance. Such real-time compliance will be more challenging where organizations (*i.e.*, NSPs and ISPs) share *Personal Data*. As stated in *Article 14*, the *Data Processor* shall provide necessary information to *Data Subjects* and comply with the origin *Data Controller* (with direct relation with the *Data Subject*) even if the *Personal Data* is not directly obtained from the *Data Subject*. Thus, a *Data Subject*'s request for a *Personal Data* deletion may need to be propagated across multiple service providers, resulting in significant overhead unless carefully designed. To illustrate, ISPs like Comcast and CenturyLink could decide to sell consumers' habits information and edge computing insights. The compliance of a *Personal Data* deletion request by all involved entities would be a significant operation.

7.3 Distributed Auditing Framework

Supervisory authorities doing continuous or frequent record gathering and compliance checks becomes increasingly difficult as the number and size of disparate network organizations increase. This requires the design of a holistic and unified auditing framework to enable DSs to provide their consent and DCs to provide the relevant records in a secure, immutable, and comprehensible manner. The authors in [46] proposed a Blockchain based platform, in which a trusted third party, DSs, DCs (including their DPs) participate to guarantee GDPR-compliant data processing. But, with the large amount of traffic passing through NSPs, extensive logging would result in tens of thousands of transactions per second which does not scale. This requires holistic designs and use of efficient cryptographic techniques.

7.4 Compliance at Application Layer

In this work we investigated NDN for GDPR compliance, with the focus on the networking infrastructure rather than the applications and services at the application layer. We argue that GDPR compliance of the application layer services is a different dimension and requires per-application assessment. Our two cents is that solely relying on data encryption and secure tunneling approaches is

not sufficient in this regard. There have been a few recent compliance initiatives with privacy acts and regulations such as GDPR, Children's Online Privacy Protection Act, and California's Online Privacy Protection Act [2, 30]. In the future, we plan to explore each individual NDN feature in more depth and design approaches to ensure GDPR compliance. Considering the scope of the work, we believe building custom solutions to ensure GDPR compliance for NDN architecture is fairly challenging and requires effort from the research community.

This paper, along another dimension, is a first step towards building a regulations-compliant network infrastructure. We demonstrate that the compliance, particularly for GDPR, can be attained with minimal overhead.

8 ACKNOWLEDGEMENTS

Research supported by Intel Labs, US NSF awards #1800088, #2028797, #1914635, EPSCoR Cooperative agreement #OIA-1757207. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or Intel Inc. We also thank our shepherd Prof. Edmund Yeh for the helpful shepherding and the anonymous reviewers for their insightful feedback and suggestions.

REFERENCES

- [1] Alexander Afanasyev, Priya Mahadevan, Ilya Moiseenko, Ersin Uzun, and Lixia Zhang. 2013. Interest flooding attack and countermeasures in Named Data Networking. In *2013 IFIP Networking Conference*. IEEE, 1–9.
- [2] Ahmad Bashir and Christo Wilson. 2018. Diffusion of User Tracking Data in the Online Advertising Ecosystem. *Proceedings on Privacy Enhancing Technologies* 2018 (10 2018), 85–103. <https://doi.org/10.1515/popets-2018-0033>
- [3] Mark Baugher, Bruce Davie, Ashok Narayanan, and Dave Oran. 2012. Self-verifying names for read-only named data. In *2012 Proceedings IEEE INFOCOM Workshops*. IEEE, Orlando, Florida, 274–279.
- [4] Abdelber Chaabane, Emiliano De Cristofaro, Mohamed Ali Kaafar, and Ersin Uzun. 2013. Privacy in content-oriented networking: Threats and countermeasures. *ACM SIGCOMM Computer Communication Review* 43, 3 (2013), 25–33.
- [5] Cloudflare. 2020. Cloudflare Data Processing Addendum. Retrieved May, 2021 from https://www.cloudflare.com/cloudflare_customer_DPAv3.pdf
- [6] Cloudflare. 2021. Cloudflare and GDPR compliance. Retrieved May, 2021 from <https://www.cloudflare.com/privacy-and-compliance/gdpr/>
- [7] CPRA. 2019. California Privacy Rights Act (CPRA). Retrieved May, 2021 from https://oag.ca.gov/system/files/initiatives/pdfs/19-0021A1%20%28Consumer%20Privacy%20-%20Version%203%29_1.pdf
- [8] David Derler, Kai Samelin, Daniel Slamanig, and Christoph Striecks. 2019. Fine-Grained and Controlled Rewriting in Blockchains: Chameleon-Hashing Gone Attribute-Based. In *26th Annual Network and Distributed System Security Symposium, NDSS*.
- [9] Steven DiBenedetto, Paolo Gasti, Gene Tsudik, and Ersin Uzun. 2012. ANDaNA: Anonymous Named Data Networking Application. In *Annual Network & Distributed System Security Symposium (NDSS)*.
- [10] Charles Duhigg. 2012. How companies learn your secrets. *The New York Times* 16, 2 (2012), 1–16.
- [11] GDPR. 2016. Complete guide to GDPR compliance. Retrieved May, 2020 from <https://gdpr.eu/>
- [12] AKM Mahmudul Hoque, Syed Obaid Amin, Adam Alyyan, Beichuan Zhang, Lixia Zhang, and Lan Wang. 2013. NLSR: named-data link state routing protocol. In *Proceedings of the 3rd ACM SIGCOMM workshop on Information-centric networking*. 15–20.
- [13] Maria Karampela, Sofia Ouhbi, and Minna Isomursu. 2019. Exploring users' willingness to share their health and personal data under the prism of the new GDPR: implications in healthcare. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (embc)*. IEEE, 6509–6512.
- [14] Dohyun Kim, Sunwook Nam, Jun Bi, and Ikjun Yeom. 2015. Efficient content verification in named data networking. In *Proceedings of the 2nd ACM Conference on Information-Centric Networking*. 109–116.
- [15] Tobias Lauinger, Nikolaos Laoutaris, Pablo Rodriguez, Thorsten Strufe, Ernst Biersack, and Engin Kirda. 2012. Privacy risks in named data networking: What is the cost of performance? *ACM SIGCOMM Computer Communication Review*

²Google Cloud has a period of around 6 months for complete deletion of data from their backup systems (<https://cloud.google.com/security/deletion>).

- 42, 5 (2012), 54–57.
- [16] Jun Li, Hao Wu, Bin Liu, and Jianyuan Lu. 2012. Effective caching schemes for minimizing inter-ISP traffic in named data networking. In *2012 IEEE 18th International Conference on Parallel and Distributed Systems*. IEEE, 580–587.
 - [17] Dominique Machulet and Rainer Böhme. 2019. Multiple Purposes, Multiple Problems: A User Study of Consent Dialogs after GDPR. *CoRR* abs/1908.10048 (2019). arXiv:1908.10048 <http://arxiv.org/abs/1908.10048>
 - [18] Giulia Mauri, Mario Gerla, Federico Bruno, Matteo Cesana, and Giacomo Verticale. 2016. Optimal content prefetching in NDN vehicle-to-infrastructure scenario. *IEEE Transactions on Vehicular Technology* 66, 3 (2016), 2513–2525.
 - [19] Travis Mick, Reza Tourani, and Satyajayant Misra. 2016. Munc: Multi-hop neighborhood collaborative caching in information centric networks. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking*. 93–101.
 - [20] Travis Mick, Reza Tourani, and Satyajayant Misra. 2017. LAsER: Lightweight authentication and secured routing for NDN IoT in smart cities. *IEEE Internet of Things Journal* 5, 2 (2017), 755–764.
 - [21] Lokke Moerel and Corien Prins. 2016. Privacy for the homo digitalis: Proposal for a new regulatory framework for data protection in the light of Big Data and the internet of things. Available at SSRN 2784123 (2016).
 - [22] Lokke Moerel and Corien Prins. 2016. Privacy for the homo digitalis: Proposal for a new regulatory framework for data protection in the light of Big Data and the internet of things. Available at SSRN 2784123 (2016).
 - [23] NetBrain. 2018. GDPR Compliance: 6 Steps to Get your IT Network Ready for GDPR. Retrieved May, 2021 from <https://www.netbraintech.com/blog/gdpr-compliance-6-steps-to-get-your-it-network-ready-for-gdpr/>
 - [24] Gaurav Panwar, Reza Tourani, Travis Mick, Abderrahmen Mtibaa, and Satyajayant Misra. 2017. DICE: Dynamic multi-RAT selection in the ICN-enabled wireless edge. In *Proceedings of the Workshop on Mobility in the Evolving Internet Architecture*. 31–36.
 - [25] Gaurav Panwar, Reza Tourani, Satyajayant Misra, and Abderrahmen Mtibaa. 2017. Request aggregation: the good, the bad, and the ugly. In *Proceedings of the 4th ACM Conference on Information-Centric Networking*. 198–199.
 - [26] Gaurav Panwar, Roopa Vishwanathan, and Satyajayant Misra. 2021. ReTRACE: Revocable and Traceable Blockchain Rewrites using Attribute-based Cryptosystems. Cryptology ePrint Archive, Report 2021/568. <https://eprint.iacr.org/2021/568>.
 - [27] Hyundo Park, Indra Widjaja, and Heejo Lee. 2012. Detection of cache pollution attacks using randomness checks. In *2012 IEEE International Conference on Communications (ICC)*. IEEE, 1096–1100.
 - [28] Sanjeev Kaushik Ramani, Reza Tourani, George Torres, Satyajayant Misra, and Alexander Afanasyev. 2019. NDN-ABS: Attribute-Based Signature Scheme for Named Data Networking. In *Proceedings of the 6th ACM Conference on Information-Centric Networking*. 123–133.
 - [29] Thomas Reuters. 2020. Technical and Organisational Measures. Retrieved June, 2020 from [https://uk.practicallaw.thomsonreuters.com/w-014-8211?transitionType=Default&contextData=\(sc.Default\)&firstPage=true&bhcp=1](https://uk.practicallaw.thomsonreuters.com/w-014-8211?transitionType=Default&contextData=(sc.Default)&firstPage=true&bhcp=1)
 - [30] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. 2018. “Won’t Somebody Think of the Children?” Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies* 2018 (06 2018), 63–83. <https://doi.org/10.1515/popets-2018-0021>
 - [31] Aashaka Shah, Vinay Banakar, Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram. 2019. Analyzing the Impact of {GDPR} on Storage Systems. In *11th {USENIX} Workshop on Hot Topics in Storage and File Systems (HotStorage 19)*.
 - [32] Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram. 2019. The Seven Sins of Personal-Data Processing Systems under GDPR. In *11th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 19)*. USENIX Association, Renton, WA. <https://www.usenix.org/conference/hotcloud19/presentation/shastri>
 - [33] Junxiao Shi, Teng Liang, Hao Wu, Bin Liu, and Beichuan Zhang. 2016. Ndn-nic: Name-based filtering on network interface card. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking*. 40–49.
 - [34] Reza Tourani, Satyajayant Misra, Joerg Kliewer, Scott Ortel, and Travis Mick. 2015. Catch me if you can: A practical framework to evade censorship in information-centric networks. In *Proceedings of the 2nd ACM Conference on Information-Centric Networking*. 167–176.
 - [35] Reza Tourani, Satyajayant Misra, and Travis Mick. 2016. Application-specific secure gathering of consumer preferences and feedback in ICNs. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking*. 65–70.
 - [36] Reza Tourani, Satyajayant Misra, Travis Mick, and Gaurav Panwar. 2017. Security, privacy, and access control in information-centric networking: A survey. *IEEE communications surveys & tutorials* 20, 1 (2017), 566–600.
 - [37] Reza Tourani, Ray Stubbs, and Satyajayant Misra. 2018. TACTIC: Tag-based access control framework for the information-centric wireless edge networks. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 456–466.
 - [38] CMS.Law GDPR Enforcement Tracker. 2020. ETID-118. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-118>
 - [39] CMS.Law GDPR Enforcement Tracker. 2020. ETID-226. Retrieved October, 2020 from <https://www.enforcementtracker.com/ETID-226>
 - [40] CMS.Law GDPR Enforcement Tracker. 2020. ETID-326. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-326>
 - [41] CMS.Law GDPR Enforcement Tracker. 2020. ETID-34. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-34>
 - [42] CMS.Law GDPR Enforcement Tracker. 2020. ETID-344. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-344>
 - [43] CMS.Law GDPR Enforcement Tracker. 2020. ETID-36. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-36>
 - [44] CMS.Law GDPR Enforcement Tracker. 2020. ETID-422. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-422>
 - [45] CMS.Law GDPR Enforcement Tracker. 2020. ETID-74. Retrieved November, 2020 from <https://www.enforcementtracker.com/ETID-74>
 - [46] Nguyen Binh Truong, Kai Sun, Gyu Myoung Lee, and Yike Guo. 2019. GDPR-Compliant Personal Data Management: A Blockchain-based Solution. *CoRR* abs/1904.03038 (2019).
 - [47] Gene Tsudik, Ersin Uzun, and Christopher A Wood. 2016. Ac3n: Anonymous communication in content-centric networking. In *2016 13th IEEE Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 988–991.
 - [48] Benjamin E Ujcich and William H Sanders. 2019. Data Protection Intents for Software-Defined Networking. In *2019 IEEE Conference on Network Softwarization (NetSoft)*. IEEE, 271–275.
 - [49] Sandra Wachter. 2018. Normative challenges of identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR. *Computer law & security review* 34, 3 (2018), 436–449.
 - [50] Yi Wang, Zhuyun Qi, Kai Lei, Bin Liu, and Chen Tian. 2017. Preventing “bad” content dispersal in named data networking. In *Proceedings of the ACM Turing 50th Celebration Conference-China*. 1–8.
 - [51] Ben Wolford. 2020. How to conduct a Data Protection Impact Assessment (template included). Proton Technologies AG. Retrieved June, 2020 from <https://gdpr.eu/data-protection-impact-assessment-template/>
 - [52] Mengjun Xie, Indra Widjaja, and Haining Wang. 2012. Enhancing cache robustness for content-centric networking. In *2012 Proceedings IEEE INFOCOM*. IEEE, 2426–2434.
 - [53] Kaiping Xue, Peixuan He, Xiang Zhang, Qudong Xia, David SL Wei, Hao Yue, and Feng Wu. 2019. A Secure, Efficient, and Accountable Edge-Based Access Control Framework for Information Centric Networks. *IEEE/ACM Transactions on Networking* 27, 3 (2019), 1220–1233.
 - [54] Yingdi Yu, Alexander Afanasyev, David Clark, KC Claffy, Van Jacobson, and Lixia Zhang. 2015. Schematizing trust in named data networking. In *Proceedings of the 2nd ACM Conference on Information-Centric Networking*. 177–186.
 - [55] Yingdi Yu, Alexander Afanasyev, Jan Seedorf, Zhiyi Zhang, and Lixia Zhang. 2017. NDN DeLorean: An authentication system for data archives in named data networking. In *Proceedings of the 4th ACM Conference on Information-Centric Networking*. 11–21.
 - [56] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, Patrick Crowley, Christos Papadopoulos, Lan Wang, Beichuan Zhang, et al. 2014. Named data networking. *ACM SIGCOMM Computer Communication Review* 44, 3 (2014), 66–73.
 - [57] Meng Zhang, Hongbin Luo, and Hongke Zhang. 2015. A survey of caching mechanisms in information-centric networking. *IEEE Communications Surveys & Tutorials* 17, 3 (2015), 1473–1499.
 - [58] Zhiyi Zhang, Yingdi Yu, Alexander Afanasyev, Jeff Burke, and Lixia Zhang. 2017. NAC: Name-based access control in named data networking. In *Proceedings of the 4th ACM Conference on Information-Centric Networking*. 186–187.