Mariano Di Martino*, Isaac Meers, Peter Quax, Ken Andries, and Wim Lamotte

Revisiting Identification Issues in GDPR 'Right Of Access' Policies: A Technical and Longitudinal Analysis

Abstract: Several data protection regulations permit individuals to request all personal information that an organization holds about them by utilizing Subject Access Requests (SARs). Prior work has observed the identification process of such requests, demonstrating weak policies that are vulnerable to potential data breaches. In this paper, we analyze and compare prior work in terms of methodologies, requested identification credentials and threat models in the context of privacy and cybersecurity. Furthermore, we have devised a longitudinal study in which we examine the impact of responsible disclosures by re-evaluating the SAR authentication processes of 40 organizations after they had two years to improve their policies. Here, we demonstrate that 53% of the previously vulnerable organizations have not corrected their policy and an additional 27% of previously non-vulnerable organizations have potentially weakened their policies instead of improving them, thus leaking sensitive personal information to potential adversaries. To better understand state-of-the-art SAR policies, we interviewed several Data Protection Officers and explored the reasoning behind their processes from a viewpoint in the industry and gained insights about potential criminal abuse of weak SAR policies. Finally, we propose several technical modifications to SAR policies that reduce privacy and security risks of data controllers.

Keywords: subject access request, GDPR policies, authentication issues, social engineering

DOI 10.2478/popets-2022-0037
Received 2021-08-31; revised 2021-12-15; accepted 2021-12-16.

***Corresponding Author: Mariano Di Martino:** Hasselt University - tUL, Expertise Center of Digital Media (EDM), mariano.dimartino@uhasselt.be
Isaac Meers: Hasselt University - tUL, Expertise Center of Digital Media, isaac.meers@uhasselt.be
Peter Quax: Hasselt University - tUL, Expertise Center of Digital Media, Flanders Make, peter.quax@uhasselt.be
Ken Andries: Hasselt University - Law Faculty, Attorney at the Brussels Bar, ken.andries@uhasselt.be
Wim Lamotte: Hasselt University - tUL, Expertise Center of Digital Media, wim.lamotte@uhasselt.be

1 Introduction

One of the key principles in data protection regulations is to give individuals control over their personal data. A common approach to execute this is the *right of access*, which legally empowers individuals (i.e. data subjects) to request all personal information that an organization (i.e. data controller) processes or holds about them. Data subjects are able to exercise the *right of access* by establishing a Subject Access Request (SAR). In the European Union, a SAR has originated from Directive 95/46/EC [21] and was later altered under the General Data Protection Regulation (GDPR) [22]. Around the world, several data protection regulations such as the Californian Consumer Privacy Act (CCPA) [5] and the Personal Data Protection Act of Singapore [1] have implemented similar rights in order to protect the privacy of data subjects.

Several prior works have researched the potential authentication issues in relation to verifying the identity of the data subject (DS) when performing a SAR. In these scientific experiments, the researchers simulate a potential adversary that attempts to perform a SAR under the name of another individual, essentially *stealing* their personal data [3, 8]. Malevolent attackers can use this technique to gain access to sensitive data related to individuals and abuse it for identity fraud or highly targeted phishing attacks, demonstrating a real threat to the privacy of data subjects. Despite the fact that prior work has proposed well-established countermeasures, the main cause of such information leakage found by these works generally point to poorly implemented policies of data controllers (DCs) and the lack of technical security knowledge in handling such request.

However, due to the social engineering nature and different methodologies used, the threat models of these works are often ambiguous to compare as they need to

1 https://www.pdpc.gov.sg/Overview-of-PDPA/The-Legislation/Personal-Data-Protection-Act

anticipate on reactions of the targeted organizations. As organizations rarely communicate about internal SAR policies, it is unknown if they are actually abused by criminal entities in reality. Although some of these works abide by responsible disclosure guidelines in which each DC is notified of their vulnerable policies, it is also uncertain whether these DCs have actually improved their policy accordingly. To better understand these concerns, we propose the following research questions:

- What prior studies have analyzed potential abuse against SAR policies? More specifically, how do they differ in methodologies and how realistic are they?
- What is the effect on SAR policies when potential vulnerable organizations are notified in the form of a responsible disclosure ? Do they improve their policies accordingly?
- How do Data Protection Officers (DPOs) handle SARs internally and how can we improve their policies? Is there a suspicion or proof that criminal adversaries are abusing weak SAR policies in real-life cases?

Providing answers to these questions is accomplished in this paper by the following aspects:

- We compare and analyze the different methodologies, threat models and identification credentials that are utilized by prior work, when performing an adversarial SAR.
- We conduct an experiment in which we review the authentication process of DCs after they were notified of their vulnerable policies two years ago. In other words, we observe whether organizations have correctly implemented the suggestions of prior work [8] from 2019 by developing a novel approach to ethically set up and send fake SARs.
- We interview several DPOs of the targeted organizations and inquire them about their internal SAR policies in order to better comprehend their current authentication methods.
- Based on their interviews and our experiment, we propose several improvements and tips to SAR implementations that will help DCs to reduce their cybersecurity and privacy risks.

The remainder of the paper is structured as follows. Section 2 gives a brief overview of prior work related to SARs in general. Next, Section 3 explains the legal details of a SAR and 'Right of Access'. Furthermore, in Section 4, the threat models and methodologies of prior work that analyzes the identity verification process are examined. In Section 5, we discuss our SAR experiment in which we review the authentication process of DCs after they were notified of their vulnerable policies two years ago. Finally, in Section 6 we explore the reasoning and verification process behind SAR policies as explained by DPOs themselves and in Section 7 we discuss and propose improvements to existing SAR policies.

2 Related work

Since the entry into force of Directive 95/46/EC [21] and later the GDPR [22], several researchers have published empirical studies around the practicality and usability of SARs, from both a technical and legal perspective [1, 9, 14, 24]. More specifically, Herrmann and Lindemann [11] have analyzed the data responses from 150 popular websites and smartphone apps when exercising the 'Right of Access' under the now obsolete Directive 95/46/EC, demonstrating excessive carelessness of DCs when handling those requests as well as incidentally revealing possibilities of criminal abuse. Following the adoption of the GDPR in 2016, Cormack [6] argued that the additional provisions related to the 'Right of Access' (now Article 15 GDPR) may increase the hypothetical risk of leaking personal information to unauthorized third-parties when those parties attempt to impersonate real data subjects. To better conceptualize these risks, Boniface et al. [2] examined potential authentication issues of various popular websites and third-party trackers more in-depth and compared the recommendations of European Data Protection Authorities (DPAs) concerning the transmission of sensitive information such as passports and national ID cards. They further acknowledge the severe impact and possible consequences of some of the weaker recommendations and conclude with several suggestions for improvements when dealing with government-issued IDs and digital identities.

Later in 2019, Di Martino et al. [8] have conducted the first practical experiment in which a simulated adversary sends fake SARs to 50 popular organizations under the pretence of being one of the other authors in the study. The authors accompanied the fake SARs with simple social engineering techniques, resulting in a significant amount of personal information being leaked to the 'supposed' adversary. Following the study of [8], several similar experiments have been performed with a slightly different dataset and methodology [3, 4, 17].

Finally, Kröger et al. [12] observed the GDPR SAR policies of 225 mobile app developers over a longer period of time and point out numerous deceptive state-

ments from developers along with general findings on inadequate data processing practices of some DCs. A similar study is being conducted related to the recent CCPA [19]. Moreover, additional works also suggest potential policy improvements related to authentication but vary widely in terms of security and usability [13, 18]. Interestingly, prior work has also been served as guidance for implementing and refining the recent CCPA statute[2].

3 GDPR 'Right of Access'

The 'Right of Access' of the GDPR gives data subjects the right to request all personal information that a data controller is processing about them [22, Art.15]. The data subject is defined as an *'identifiable natural person'* [22, Art.4(1)], while the data controller is defined as the entity that determines *'the purposes and means of the processing of personal data'* [22, Art. 4(7)]. In a practical context, the data subject (DS) is in fact the person requesting the data, while the term 'data controller (DC)' refers to the organization that is responsible for handling that request. In addition to the raw personal data, the DC should also provide the DS with details about their data such as the retention period and the purpose of the processing [22, Art. 13(1)]. Exercising the 'Right of Access' is usually accomplished by sending a Subject Access Request (SAR) or Data Request to the DC.

Some important aspects of this right are summarized below:

- The request may be performed in *'writing, or by other means, including, where appropriate, by electronic means'* [22, Art. 12(1)]. In other words, the DS may request their data through at least email or postal mail.
- The DC is required to provide the personal data to the DS within one month. Although the data should be provided without undue delay, if the DC is unable to provide the data within the timeframe (e.g. when the number of requests are substantial), then they may extend the deadline with an extra two months [22, Art. 12(3)]. However, this extension should be communicated to the DS within the first month of receiving their SAR [22, Art. 12(4)].
- The DC must provide the personal data to the DS free of charge, with some exceptions (e.g. excessive repeated requests) [22, Art. 12(5)].
- The DC should *'use all reasonable measures to verify the identity of a data subject who requests access'* [22, Rect. 64].

Finally, under some circumstances, a Data Protection Officer (DPO) is appointed to the organization. This entity has the responsibility, among other things, to *'monitor compliance with this regulation'* [22, Art. 39].

4 Analysis of identity verification

4.1 Introduction

Verifying the identity of the DS is a crucial aspect of having a secure SAR policy. When an organization receives a SAR, it has to verify whether the person who is sending the inquiry is the same person as the one they are asking the personal data from. Numerous authentication methods are utilized by organizations, many of which are not safe for even the weakest threat model possible [3, 4]. To the best of our knowledge, SAR experiments where the main focus is evaluating the authentication process from an adversarial viewpoint have only been conducted by the authors listed in Table 1. In these prior studies, the threat model consists of a malicious actor who attempts to steal personal data from a specifically targeted individual. This is achieved by collecting a minimal amount of personal information about the subject in order to send a fake SAR under the false pretence of being the subject. However, comparing these studies is not trivial due to the different methodologies that are utilized and the manner in which the statistics are being formulated. In this section, we analyze prior work in terms of the attacker methodology and the credentials requested by organizations when they are authenticating the DS.

In Table 1, we present prior works that have conducted SAR experiments in which an adversary is simulated. Here, we list the different credentials that were requested and how many of these DCs were persuaded by social engineering techniques. The first percentage is calculated based on the total number of organizations that require the mentioned credential in their first correspondence with the supposed DS. The second or bold percentage indicates the absolute percentage of DCs that the authors were able to persuade by social

2 https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-15-day-comments-022520.pdf

	Di Martino et al. [8] (2019)	Pavur & Knerr [17] [1] (2019)	Bufalieri et al. [3] (2020)	Our work (2021)
Total number of tested orgs.	55	75	326	54
ID card or passport	24%	13%	10%	N/A
Subject email access	31% / **4%**	15%	16% / **6%**	22% / **5%**
Account login	33% / **2%**	15%	21%	36% / **2%**
User-specific data	16% / **2%**	7%	6%	22% / **2%**
Call subject	5%	1%	1%	2%
Device cookie	0%	11%	1%	0%
Address or region	9%	N/A	N/A	5%
Dataset	Alexa top	Unknown	Alexa top	Same as [8]
Adversarial methodology	MAIL[2], SE[3], SIMUL[4]	MAIL	MAIL, SIMUL	MAIL, SE, SIMUL
Initial request	Email	Email	Email	Registered letter
Number of subjects	2	1	N/A	5
Number of unresponsive orgs.	4 (7%)	17 (23%)	92 (28%)	3 (5%)

Table 1. Prior works that conducted practical SAR experiments where the identity verification process has been assessed and if possible, bypassed. The percentages are calculated based on the total number of tested organizations that required the listed credential in their first correspondence. The second or bold percentage listed in some rows indicate the (absolute) percentage of DCs that the authors were able to persuade by social engineering to not request that credential in exchange for –less safe– credentials, even though that was not requested in their first correspondence.

[1] We only observed the initial 75 organizations in this work as the other 75 organizations are contacted using *daisy chained* information. In addition, this work is not academically peer-reviewed.
[2] MAIL: Email spoofing or a fake email address was used.
[3] SE: Active social engineering attempts were used to convince or persuade the DCs.
[4] SIMUL: ID cards were altered by using image manipulation software.

engineering to not request that credential in exchange for –less safe– credentials, even though that was not requested in their first correspondence. To illustrate an example, we take the '*subject email access*' row. Here, 31% of the organizations require this credential in the work of [8], while 4% of the organizations were persuaded by social engineering to request other weaker credentials, such as the address of the DS, even though 'subject email access' was first required by the DC. The percentages in this table are not necessarily cumulative as some organizations may utilize multiple authentication methods. For clarity, we have also included our experiment of Section 5. Finally, fine-grained statistics from [4] were not available and are therefore excluded from the table.

In the following subsections, we explore the context of the prior SAR experiments and the requested credentials in more detail.

4.2 Threat model and responsiveness

When a potential adversary develops a threat model for abusing SARs, several methodologies are possible, depending on the repertoire of information available to the adversary. For instance, [8] initiated their SARs by providing the name and email address to the DC, while [17] provided additional basic information (including the home address). It is therefore unknown whether some DCs would require email addresses or home addresses in the first place as they are already provided in the initial request. Another inconsistency is that [17] considered organizations to be vulnerable based on the theoretical strength of the authentication methods, without actually attempting to circumvent these. Differently put, they provided the correct (often censored) credentials such as a document that proves the DS address, assuming it would be considered insecure without attempting to manipulate such credential. In contrast, [3] attempted to evade these ineffective methods by creating similar-looking email addresses and (semi) simulating ID cards, while [8] also carried out social engineering techniques which we briefly explain in Section 5.3. In the work of [17], they mention the use of social engineering, however, they did not clarify how such social engineering attempt was performed with the exception of an excuse developed in their initial SAR.

Moreover, as SAR processes are often not public knowledge and mostly unique to each organization, authors are required to make assumptions related to the processing of authentication documents and personal data. The process of authenticating SARs may vary

based on the sensitivity of the personal information requested (see Section 6.3 and 5.5). For instance, health records are more sensitive and hence may be located in different databases and thus require different verification approaches. Or perhaps policies may be biased by the legal wording used in the SAR. For instance, a more aggressive tone in text may incline a DC to be more responsive. Another important element in adversarial SARs is to protect the privacy of the DSs participating in the experiment. To avoid using real data subjects, [3] created a simulated account and populated it with realistically looking data, while [17] and [8] used respectively 1 and 2 real data subjects (which are all voluntarily consenting co-authors).

To this date, Bufalieri et al. [3] performed the largest adversarial SAR experiment with 326 organizations carefully picked from the Alexa top websites (i.e. a list of the most popular websites currently available), but also lists 28% of these organizations not being responsive. However, [8] only lists 7% as unresponsive from a set of 55 organizations in the Alexa top, reinforcing the argument that smaller or less popular organizations (ranked lower in the Alexa top) are often insufficiently equipped with the necessary legal and technical knowledge to correctly handle SARs [20] and may therefore be more susceptible to such attacks.
Although not listed in Table 1, Cagnazzo et al. [4] partially send out SARs by postal mail, observing even less safe authentication methods when compared to sending it by email. This is corroborated further by our interviews in Section 6, which shows that organizations usually do not receive a significant number of SARs through postal mail.

4.3 ID card or passport

To verify the identity of the DS, between 10% and 24% of the organizations requested an ID card or passport in some form, which is considered unsafe by all authors [3, 8, 17]. Many DCs are satisfied with a photocopy of an ID card where everything is censored except for the full name, date of birth and portrait picture, in alignment with the 'data minimization' principle [22, Art. 5(1)(c)]. However, censoring such identifiable information makes it trivial for a potential adversary to use digital manipulation software and alter a photocopy of the ID card to resemble the identity of another individual [8]. Basic information such as the date of birth and portrait picture may be found through Open Source Intelligence (OSINT) techniques such as browsing through the social media profile of the DS. In addition, [3] demonstrated that most DCs would accept a very low-quality image of a real ID card as being valid, justifiably arguing that a professionally simulated ID card would be of better quality and would thus likely be accepted with at least the same success rate.

It must also be noted that digital and written data on ID cards and driver's licenses vary greatly over different jurisdictions and countries. For example, European ID cards usually have the National Register Number (NRN) written on the back which is occasionally required by DCs, while others are satisfied with the NRN censored. Yet, a small group of DCs requires the DS to leave the NRN visible due to the fact that it is used as an access identifier for several government organizations. For instance, some insurance companies require the NRN to access a governmental database that contains a historical overview of car accidents in which the DS is involved in order to attach that information to the response of a SAR. It is important to note, however, that most local jurisdictions regulate the processing of certain personal data such as NRNs by law, and thus may result in different approaches. Furthermore, [8] replaced the subject's face on the ID card without actually knowing if the picture is checked by the DC. While [3] argued that DCs generally do not have the necessary tools to detect tampering of photocopied ID card images which we will affirm later during the DPO interviews in Section 6.

Requesting uncensored ID cards or passports is not only vulnerable in a security context, but may also be perceived as an 'abusive' identity check [2]. By sending an ID card to the DC for verifying the DS identity, the DC may receive new personal data from the DS that was not available to them before (e.g. the NRN) and is therefore unable to verify whether or not the NRN actually belongs to the DS performing the request. An extended discussion of this topic is laid out in Section 7.

4.4 Subject email access

16% to 31% of the organizations require the subject to prove they have access to the email inbox corresponding to their own email address known by the DC. This is verified by either sending the personal information to the corresponding inbox or by requiring the initial request to be sent from the email address that is known to the DC. In [8] and [3], respectively 4% and 6% of these organizations could be tricked using social engi-

neering to request other –less safe– authentication credentials or were simply persuaded to send the subject's data right away. Most authors [3, 8, 17] consider this type of credential to be relatively safe. However, it fails to correctly authenticate the DS if the attacker would have control over the subject's inbox. In addition, care should be taken when blindly processing SARs for which the original inquiry comes from an email address known to the DC, as the sender's email address may be spoofed. Differently put, receiving a SAR from the original email address does *not* provide an absolute proof that the person making the request is actually the legitimate subject. The DC should preferably send a verification code to the email address known to be from the DS, which must then be copied by the DS in a reply to the DC. Besides email spoofing, the authors of [12] and [3] observed that some email responses to SARs do not exhibit proper TLS encryption, making them vulnerable to other man-in-the-middle (MitM) attacks. In essence, this indicates that adversaries may potentially intercept such emails and read or alter any communication between DC and DS.

4.5 Account login

15% to 33% of the organizations require the subject to log in to their online platform or website by using the subject's credentials, such as an email and password. This may be the same platform on which the subject has to log in to, in order to access the DC main services of the organization. Alternatively, it may also be a separate part of the platform that is only available after manually performing a SAR through email or by postal mail. This type of authentication method aligns with [22, GDPR Rect. 63], that states that *'Where possible, the controller should be able to provide remote access to a secure system which would provide the data subject with direct access to his or her personal data.'*. Such authentication methods rely on the same security measures as when the subject logs in to the online platform by normal means. Although these credentials are generally considered to be secure, it is important to integrate possible two-factor authentication credentials, if available. For instance, if the user has enabled a 2FA app such as Google Authenticator [10] on their normal login, then it is suggested to apply *at least* the same 2FA step for verifying a SAR. Nevertheless, additional precautions should be taken when dealing with insecure [15] or potentially difficult-to-perform 2FA steps [7].

4.6 User-specific data

6% to 16% of the organizations asked the subject for specific user data that is difficult for an adversary to infer. This includes credentials such as security questions, credit card numbers or ID numbers on utility bills. 2% of the organizations that requested such information were subsequently convinced by [8] to accept credentials that are less safe. It is important to note that [3] and [17] classified the request for the DS home address to this category, while [8] has separated 'user-specific data' and 'address or region of residence'. In one case of [8], the DC requested specific information of the DS that is publicly visible by default on the webpage of the DS. An equivalent hypothetical example would be to ask the DS to send the most recent post made on the DC's online platform, which is publicly visible by default when posted. Moreover, [4] observes that it is problematic when a DC requests user-specific data that is usually not modifiable (e.g. bank account numbers), considering that if such data is leaked externally, the DS may not be able to change that specific information and thus the policy may fail to filter out potential adversaries.

Furthermore, as presented by [17], the specific user data required for some SARs may be gathered from *daisy chains*. In other words, the adversary may use personal information of the subject received from one SAR response to verify their identity of a SAR directed to another DC. Although the addition of daisy chaining appears interesting at first sight, there is an important adversarial risk that makes such a strategy uncertain and difficult to maintain. For instance when performing a fake SAR, numerous organizations will respond to the fake email address of the adversary *and* the original email address of the subject [3, 8]. In this manner, the legitimate subject may be notified when receiving the response to a fake SAR, allowing them to take appropriate actions such as informing other organizations that process their data so they are able to take additional precautionary measurements. It is therefore difficult to assess whether daisy chaining on a large scale is possible at all.

4.7 Call subject

1% to 5% of the organizations called the subject to verify their identity. Most authors agree that this method is relatively safe [8, 17] in the situation that an adversary does not reside in the same household as the subject. However, the threat model becomes more complicated

when the adversary is closely related to the DS. For example, in one case discussed in Section 6.5, the adversary and DS were previously married and involved in a divorce. In this scenario, DCs should be careful when calling the subject, as the partner (or an accomplice) may answer the phone and verify the request for them. Similarly to email spoofing, the phone number of the data subject may be spoofed too if the adversary calls the DC by modifying the Caller ID [16]. It is therefore important for the DC to call the DS and not the other way around. Finally, in case the DS requests the information to be provided orally (e.g. by phone), the DC must verify the identity of the DS by other means [22, Art. 12(1)], for instance, by email.

4.8 Device cookie

This work and [8] did not observe any organization that required a device cookie, while [17] and [3] mentions that respectively 11% and 1% of the DCs requested this credential. We argue that the large discrepancy in percentage is due to the fact that the organizations in the dataset of [17] are not selected by any specific method, resulting in possible biases. This credential is often required by DCs that do not have access to the email address of the DS [17]. For instance, advertising organizations link specific cookies to the identity of the DS as that is often the only unique identifier they have for any given individual or device owner. However, when specifically targetting third-party trackers and ad networks as in [2, 23], such credential is more often requested.

5 SAR experiment

5.1 Introduction

As discussed by prior work and in Section 4, it is clear that a large number of organizations are not well equipped with the necessary knowledge to prevent personal information leakages through SARs. To abide by a responsible disclosure and ethical guidelines, [3] notified all organizations that were vulnerable to their attack and provided technical suggestions and countermeasures, while [17] only notified some vulnerable organizations in order to avoid potential name-shaming the employee who 'fell' for the attack. After the experiment of [8], the authors notified the organizations (in the form of a responsible disclosure) that were vulnerable to their threat model and provided them with suggestions on how to improve their SAR policy. These suggestions are briefly mentioned below:

- If data from person X is requested in a SAR and the DC knows that a specific email address Y belongs to person X, then the inquiry should either originate from email address Y or the requester should be able to verify it has access to email address Y.
- If a phone number of the subject is available, then call the subject to verify their identity.
- If none of the options above are possible, then request pieces of specific non-public personal data that are known to the DC.

Despite their suggestions, it is unclear what impact this responsible disclosure had on these vulnerable organizations. According to the threat model of [8], correctly implementing the first suggestion would already prevent the organization from leaking personal data and would thus not be vulnerable. In this section, we examine whether the same DCs that were vulnerable to the attack of [8] have improved their SAR policies two years later in order to prevent potential data breaches. Likewise, we analyze whether the previously not vulnerable DCs have potentially weakened their policies over time.

5.2 Methodology

In [8], the organizations were contacted through email considering that this is the most used method of performing SARs. However in this experiment, we sent the same set of organizations a registered letter (without return address) containing a SAR and included limited information to 'verify' the identity of the subject. This way, we were able to track the letter and know with certainty when the organizations had received our request by monitoring the tracking numbers. The limited information in the letter consisted of a simulated copy of the ID card, the name, original and *fake email address* of the subject. In this letter, we developed an excuse in which we explained that – *SARs sent from the original email address of the subject were being bounced back, while SARs from the fake email address did not. But overall, no response was received and a letter is therefore sent.* – (Appendix A.1). In contrast to prior work, no email spoofing was performed. Similar to [8], the simulated ID is produced by digitally manipulating a photocopy of a random ID card (included in the letter) and replacing the date of birth, profile picture and full name of the subject, which we found publicly through OSINT. Fur-

thermore, we censored all other visible information on the ID card such as the NRN and signature.

In [8], there is a total of 55 tested organizations in the dataset of which the majority is extracted from the Belgian Alexa top 50 in combination with additional organizations that the subjects had a consumer relation with. In this experiment, we analyzed the same organizations as in [8], however, we left out one organization that went bankrupt. Additionally and similar to the approach of [8], we did not re-evaluate 14 organizations that had (and still have) a fully automated process, resulting into 40 organizations for our dataset.

In total, we assumed the identity of 5 subjects which are all employees of Hasselt University and all gave full consent during all the stages of the experiment (see Section 5.6). To avoid the authors from receiving potential sensitive information of the subjects during the study, we set up a process that puts the subjects as a controlling MitM, demonstrated in Figure 1.

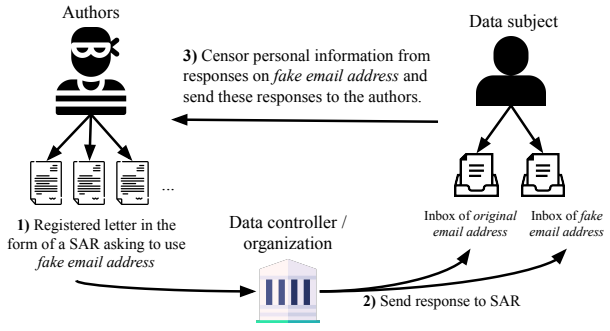

Fig. 1. The three-step flow of our SAR experiment. The authors never observe personal information from the DS as the DS censors all personal information before sending it to the authors. Note that only the censored responses from the fake email address are sent to the authors.

Each subject was first asked which of the organizations they think had personal information related to them. Afterwards, we randomly assigned the organizations to them so each subject had approximately the same number of organizations, which is further specified in Appendix A.2. Next, for each subject, we utilized a free domain provider to create an email address that is supposed to act like the *fake email address* of the adversary. Each fake email address was constructed in relation to the full name of the test subject (*firstname.lastname@provider.com*). Finally, the password of each fake email address was given to the appropriate subject and they were subsequently required to change the password so that the authors could not access it. To maintain a consistent methodology, we also made sure that the publicly found OSINT data that is used on the simulated ID card was correct by requesting each subject to verify their data.

It is important to note again that the initial experiment of [8] was performed by email, while our experiment is executed by postal mail using registered letters. However, if an organization would have correctly implemented the suggestions of [8], then our experiment would have failed for that specific organization. In other words, the organization would have not been vulnerable according to the threat model. We chose to send out registered letters instead of emails to prevent DCs from directly detecting the attacker methodology of [8] instead of implementing safe authentication methods.

5.3 Communication

Starting from the setup in the prior section, we sent the registered letters to each organization in December 2020. In case the DC did *not* respond to our letter or did *not* provide us the necessary data within 30 days (plus possible 60 days extension when permitted), we sent them an email from the fake email address of the subject, reminding the DC of the deadline and pressuring them to respond.

Although the registered letter contained the fabricated story that illustrated the core concept of our social engineering attack, some DCs required additional persuasion by email. Therefore, we utilized similar social engineering techniques as in [8]. A brief overview is shown here:

- **Dismiss access to the DS email address**: When the DC requires proof that the subject has access to the original email address, we state that 'we are unable to send or receive emails from the DC email domain as described in our letter'.
- **Deliberately omitting unknown credentials**: When the DC requests specific information that an adversary does not know, we simply ignore that request and provide other information that is available to the adversary. We also censored the NRN on the simulated ID card since we consider that to be unknown to a potential adversary.

In addition, we consider an organization to be *vulnerable*, *unresponsive* or *safe* in the following cases:

- **Vulnerable**: *if and only if* we received the personal data of the subject from the DC through the *fake email address* of the adversary within 90 days. Ex-

amples of such data include financial transaction details, location data, online shopping behavior and sensitive job application information.

– **Unresponsive**: *if and only if* we did not receive any response to the SAR either through the original email address of the subject or the fake email address of the assumed adversary within 90 days. According to Art. 12(6) and Art 12(4) of the GDPR [22], in case the DC has doubts about the identity of the DS, it should request additional information to verify the identity. In any case, refusing to act upon the request would require the DC to at least notify the DS about this event. Not responding to the request, either to the original or fake email address is therefore not an option.
– **Safe**: in all other cases. The organization either sent the personal data only to the original email address or they required authentication methods that are difficult to manipulate for an adversary. Moreover, we were not able to persuade these organizations using social engineering techniques. For our threat model, we consider all requested credentials except for a censored ID card, home address, name and date of birth to be safe and therefore not vulnerable to our attack.

Although the participants/subjects of this study are positioned as a MitM, they were, however, not permitted to change the email contents written by the authors in any way, and thus had no direct influence on any of the communication between the authors and DCs. In addition, all email communication related to the SAR that was sent from the DC to the *original email address* of the subject, was ignored. However after the experiment ended, we asked the subjects to report us whether they received a response on their original email address or home address from every organization. This is necessary to decide between the 'safe' or 'unresponsive' category mentioned above. Finally, all subsequent communication with the organizations after sending the letters, was carried out between December 2020 and April 2021. In this paper, the organizations are anonymized by categorizing them according to their market sector as proposed by [8] (Appendix A.3).

5.4 Results

In Table 2, we show that the number of vulnerable organizations have increased from 15 to 17 after a period of two years. We note that 3 out of 4 organizations that did not respond in 2019 to the authors of [8], did respond in 2021 to our request and were vulnerable to the attack. Ignoring these three organizations, there are still 6 organizations left that were not vulnerable in 2019, but are now in 2021. It is unknown whether some of these organizations were simply tricked using social engineering techniques or if they altered their SAR policies by allowing weaker authentication methods in general. On a more positive note, 4 organizations of which 2 are financial businesses, have improved their policy such that they are no longer vulnerable to our attack. In Table 3, we indicate the change of authentication credentials requested over the years in more detail and lay them out in the same structure as [8]. *Fin_A* was the only organization that sent the personal data to the subject by physical mail, a change from their policy in 2019. This was their standard policy as they were not willing to provide the personal data by email. Furthermore, *Fin_C* did not significantly change their policy, but were simply instructed internally to not deliver the data by email, as they provided the data using their authenticated online platform. Furthermore, 5 organizations either requested a verification code or sent the personal information through the original email of the subject in 2019, however, they were now all persuaded by our letter to send the data to the attacker-controlled fake email address. In comparison to [8], an additional 5 organizations were persuaded to *not* request additional information such as the address of the subject, calling the subject, or even requesting specific data that would be difficult for an adversary to know.

	2019	2021
Total number of orgs. in dataset.	55	55
Total number of manually contacted orgs.	41	40
Undetermined organizations	N/A	3
Safe organizations	22	17
Unresponsive organizations	4	3
Vulnerable organizations	15	17
Vulnerable both years	8	
Vulnerable only in 2019	7	
Vulnerable only in 2021	9	

Table 2. Comparison of [8] from 2019 and our work in 2021, in terms of vulnerable organizations. Between the organizations in the dataset and the contacted organizations, we ignored 14 organizations that had a fully automatic process in 2019 and 2021. In addition, we filtered out one organization that went bankrupt prior to this study. The 'undetermined' row indicates the organizations where the methodology has failed outside of our control.

In another instance, one organization (*Ret_I*) stated that they did not receive our registered letter because the postal code was supposedly wrong. When we followed up with a comparison of the address stated in their privacy policy and our tracking number, they proceeded to completely remove the personal data of the subject without consent, acting as if no data from the subject was available. This type of reaction is not new and is in line with prior studies [12]. Later, we discovered that some translations of the privacy policy of *Ret_I* had a different postal code, resulting in the letter *potentially* being delivered to the wrong address. At the end, it was unclear as we could not decisively indicate whether the DC lost the letter or if it had never received it in the first place. Fortunately, *Ent_L* and *Trl_C* improved their policy significantly by calling the subject or requesting specific user data (such as non-public account or serial numbers) that is usually not known by an adversary, rendering them safe in our study. Due to the fact that registered letters entail a certain risk when it comes to losing the letter in transit, we observed two cases (*Ent_B* and *Oth_G*) in which the letter could not be delivered. In one of these cases, the mailing company lost the letter completely (which contained a copy of the simulated ID card) and compensated us with an amount of € 39 ($46). Finally, 4 organizations explicitly stated that they did not find or received our registered letter even though our tracking number reported otherwise. The whereabouts of these letters, including the ones of the organizations that remain unresponsive, are unknown. Situations in which the post office lost our letter or when we were not absolutely sure whether these organizations have actually received our letters, are classified as 'undetermined' in Table 2.

During our study, we observed several other interesting aspects related to our requests:

- Two organizations packed the personal information in a password-protected ZIP archive, sent it to the fake email address and proceeded to send the password of that ZIP file in another email to the same fake email address. It is unclear whether this is a mistake made by the individual handling the request or an additional authentication method that provides (questionable) data security guarantees. Similar cases occurred in [3, 12].
- One organization requested us to log in with the subject's existing account on a company-owned website in order to verify the identity. Since our threat model in this study does not attack this specific method, we requested to cancel the SAR. However, upon that request, the organization suddenly stated *'the data is ready to be sent. We can send it to you now if you want?'*, which we accepted resulting in receiving the subject's personal data without having to verify our identity in any way. Slightly similar cases occurred in [12].
- We observed at least 3 cases in which a data protection employee of the organization visited the LinkedIn profile of the DS shortly after we sent a SAR. We assume this is related to verifying the identity of the DS, however, the exact reason remains unclear.

Considering the number of organizations that remain vulnerable after the responsible disclosure in 2019 and having had two years to implement the suggested countermeasures, we are worried that some SAR policies still entail a significant privacy risk for the DS. Overall, 53% of all vulnerable organizations in [8] have not yet improved their policies, while an additional 27% of the non-vulnerable organizations have potentially weakened their policy as we were able to gain a significant amount of personal information from our participants. We can conclude that even after a responsible disclosure, a significant number of organizations have not implemented these necessary precautions to prevent potential data breaches through SARs

5.5 Sensitivity of leaked and requested data

The type of personal data received through the fake email address of the DS was quite similar to the data received by the study of [8]. That should not be a surprise considering that each organization usually handles a specific set of personal information (e.g. financial). However logically, not every organization has an equal amount of personal information from each DS as some may use their services to a lesser or larger extent. Specifying and analyzing which exact data was leaked by each organization is therefore not a meaningful discussion. Nevertheless, it is still important to discuss what type of personal data was received and which relation it has to the requested identification data used in the authentication procedure of the DC. Note that the personal data discussed below is not exhaustive of all data that we received as an adversary in this experiment.

Common pieces of personal information that were received from the vulnerable organizations are email addresses, the home address and phone number of the DS. Although that type of information is considered to be

Organization	Account login		Subject email access		Address or region of residence		Back ID card (NRN)		Call subject		Specific user data		Vulnerable	
	'19	'21	'19	'21	'19	'21	'19	'21	'19	'21	'19	'21	'19	'21
Fin_A							*	*					✓	
Fin_B													✓	✓
Fin_C	*	*											✓	
Fin_D				*									✓	✓
Ret_A													✓	✓
Ret_B	✓	*	*	*								*		
Ret_C		✓	✓											
Ret_E					✓	✓							✓	✓
Ret_F					✓	✓							✓	✓
Ret_G		*									✓			✓
Ret_H			✓	✓										
Ent_A			*	*							*	*		
Ent_C											✓	✓		
Ent_D		*	✓								*	*		
Ent_E			✓	✓										
Ent_F											✓	✓		
Ent_H			✓	✓							✓	✓		
Ent_I	✓	*	✓	*								*		
Ent_J			✓	✓										
Ent_K			✓											✓
Ent_L			✓								*	✓	✓	
Trl_A			✓										✓	✓
Trl_B			✓											✓
Trl_C			*							✓		✓	✓	
Trl_D			✓										✓	✓
New_A					✓								✓	✓
New_B			*	*							*	*		
Oth_A	✓			*		*								✓
Oth_B									✓					✓
Oth_D			✓									*		✓
Oth_E				✓							✓	*		
Ent_M	U		U	*	U		U		U		U		?	✓
Ret_J	U		U		U		U		U		U		?	✓
Oth_F	U		U		U		U		U		U		?	✓
Oth_C	Unresponsive in 2021													?
Oth_G	Letter lost in transit in 2021												?	?
Fin_E	Unresponsive in 2021													?
Ent_G	Letter not picked up by organization in 2021												✓	?
Ent_B	Letter address unreadable in 2021													?
Ret_I	Removed data without consent in 2021												✓	?

Table 3. Overview of requested credentials for 40 organizations in 2019 and 2021, including the ones that were unresponsive (U). A '✓' indicates that the corresponding organization used the mentioned credential to verify the identity of the DS. A '*' shows that the corresponding credential was not forced, by either accepting an alternative credential or by being able to persuade the DC using social engineering. A '?' indicates that it is unknown whether the organization was vulnerable in that particular year. A summary of this experiment is listed in the last column of Table 1 and 2.

'personal data' as stated in [22][Art.4(1)], our threat model considers this to be data that may often be found through the use of OSINT. In one case, an organization leaked financial records, consisting of credit card numbers, payment logs and financial service details that were used by the DS in question. Another organization leaked the whereabouts such as the exact GPS coordinates and routes taken by the DS at different points in time, going back for more then 3 years. Furthermore, various organizations leaked website logging data such as the visited web URLs or the predicted interests in certain news articles. A final organization leaked details about retail purchases offline and online, together with serial numbers of products bought by the DS. In any case, all organizations leaked at least some pieces of information of the DS that we, as authors, were not able to infer through OSINT. We conclude that all data responses of the vulnerable organizations contained sensitive personal information that would moderately or sometimes severely affect the privacy of the DS. Especially with financial information, identity theft may, unfortunately, be a realistic consequence in case this would leak to a potential (adversarial) third party.

5.6 Notes on ethical research

The Ethical Research Committee (ERC) and legal council of Hasselt University have explicitly approved and authorized the experiment discussed in Section 5. The individuals that were involved as data subjects gave full consent for their participation during the whole study and have approved that the authors initiated SARs under their name. Each individual was given the option to end their participation in the study at any given time. We should also point out that the authors could not directly read any email communication between the subjects and the corresponding DCs as the individuals were set up as a MitM as described in Section 5.2. Therefore, the authors did not receive any personal information from the data subjects through the DC.

To abide by responsible disclosure guidelines, all DCs that we considered to be 'vulnerable' to our threat model were notified of our research and were given suggestions on how to improve their 'Right of Access' policies accordingly. In this paper, all names of the DCs are anonymized to protect the organizations from reputational damage and potentially criminal abuse, in case the organizations would not improve their policy. The vulnerable organizations were also not reported to the DPAs. Every organization that responded to our responsible disclosure appreciated our suggestions and answered positively to any follow-up correspondence. Finally, as discussed in Section 5.2, our experiment involved altering the individual's proof of identity which was then included in the registered letter and sent to the appropriate DC. We should stress that no official government documents were modified during this study, only a digitally scanned photocopy.

We strongly recommend that future studies should take these ethical considerations into account when deploying a similar experiment.

6 Replies and interviews

6.1 Replies

	2019	2021
Number of vulnerable organizations	15	17
'Taking suggestions into account'	9	7
'Forwarded to appropriate channels'	2	3
No reply or automatic reply	4	7

Table 4. Number of replies from vulnerable organizations when performing a responsible disclosure. 2019 is linked to the study of [8], while 2021 is linked to this work.

To the best of our knowledge, Urban et al. [23] were the first to interview various privacy employees regarding transparency and overall GDPR compliance in online advertising, demonstrating the need of a clear guidance from executive powers on how to properly assess and improve such compliance. However due to the different focus of the study, the interview questions and answers related to authenticating the DS for SARs are sparse. In addition and despite prior work demonstrating several issues in the authentication and data processing practices of SARs, there has been limited information available on how DCs react to responsible disclosures and how they handle their SAR policies internally.

Table 4 shows a summary of the type of responses that we received when notifying the DCs with a responsible disclosure and suggestions on how to improve their vulnerable SAR policies. In 2019, 9 of the vulnerable organizations stated that they 'would take the improvement suggestions into account'. However, 6 of these 9 organizations are still vulnerable to our SAR experiment in 2021. Furthermore, there is an increase from 4

to 7 organizations that did not reply to the responsible disclosure. Finally, 3 organizations stated that they 'would forward our responsible disclosure to the appropriate channels', as many privacy inquiries were initially received by customer service which often forward them to the internal email address of the DPO.

Moreover, some of the vulnerable DCs responded to us with interesting aspects of their policy which we briefly mention below:

- One DC finds it difficult to strike the right balance between security and abusive identity checks (see Section 7 for a more in-depth discussion). They also argue that the amount of personal data their organization holds, does not require a stricter policy for verifying the identity of the DS. Moreover, the DC states that they do verify that the email address where the SAR originates from, is in fact from the legitimate DS. Although our experiment in Section 5 has demonstrated this verification was unsuccessful, as the organization was vulnerable in our study.
- Similarly, another DC responded that requesting an ID card is sufficient for the amount of personal data they hold. In addition, they argued that this was their first SAR ever. However, this was not the case for their organization as [8] already initiated a SAR (by email) to the same organization in 2019.
- One DC responded that this was their first SAR by postal email.

In addition to the replies above, we also questioned several DPOs of organizations that were *not* vulnerable in our experiment. In these interviews discussed in the next sections, we examine if and how they have implemented our suggestions of [8], their internal identification checks and potential cases of SAR abuse.

6.2 Interview methodology

In the final phase of our experiment in Section 5, we contacted every organization that was not vulnerable to our experiment, –in addition– to the responsible disclosure emails in Section 5.6. However in this email, we mentioned both experiments ([8] and this work) and requested whether it was possible to have a live video meeting with the DPO of the organization to gain more information about their SAR policies. In case we found a direct email address of the DPO in the privacy policy of the organization, then we used that one. For some organizations, we also used the email address of the DPO that we already possessed from our prior study [8]. From the 17 organizations, only 3 organizations answered positively to our request for a meeting, while 2 organizations either provided us with a support/privacy employee or required us to send the questions by email without having a live meeting. Since our goal was to exclusively interview DPOs and to have a global understanding of the organization SAR policies, we politely declined the last 2 organizations. We explicitly did not request interviews from organizations that were vulnerable in this experiment as we were uncertain whether these organizations would feel *attacked* or *ashamed* as their response may then be biased to avoid potential legal implications.

For the remaining 3 organizations, we devised and asked the following interview questions in the exact order below:

- **Q1**: How many SARs have been received *before* and *after* the adoption of the GDPR?
- **Q2a**: Have there been any policy changes to handling a SAR since the research of [8]? If yes, what are these? (only for organizations that were vulnerable in '19).
- **Q2b**: What is the specific information that is used or required to identify or authenticate the DS in a SAR?
- **Q3**: Has there been any suspicion that third parties are attempting to falsify SARs to request data from other individuals?
- **Q4**: How many employees of your (or an external) organization have access to the personal data required to handle a SAR?

It is important to note that some organizations provided extra information to questions that we did not envision or that we had to censor partially to avoid deanonymization. Although this would make their additions primarily anecdotally, we still think they are interesting for future work and have therefore, included them in the interview answers. Finally, some questions were not answered in their entirety, as some did not want to disclose specific information or the information was simply not available.

In the following subsections, we provide the answers (mapped as A1 to A4) of each organization on the interview questions.

6.3 Interview of organization X

Organization X has more than 9 million customers and a revenue of more than 5 billion euro per year. The organization was vulnerable in 2019, but not in 2021.

- **A1**: Since the adoption of the GDPR in 2018, we have received between 55 and 75 SARs per year. The annual average number of SARs received has almost doubled since the adoption of the GDPR in comparison to when the now repealed Directive 95/46/EC was in effect. Only 1 to 2 SARs per year are received by postal mail. Also, we often receive SARs in combination with a complaint unrelated to privacy. Around 10% of the SARs are sent utilizing the same external tool[3].
- **A2a**: The work of [8] has induced small policy changes in relation to confirming and verifying the email address of the DS. We are much more careful now to manually check whether the email address is correct. Unfortunately, larger policy changes did not occur due to budgetary constraints.
- **A2b**: A copy of an ID card is still required, although we have no tools to verify the authenticity of a photocopied ID card. Besides the email check and ID card, we do not have a fixed policy but *'we do take several additional measures to make sure we are sending the data to the right person'* (literal quote). Also, the process of verifying the identity of the DS varies depending on how the inquiry is performed. For instance, digitally vs via postal mail. Usually, the response to a SAR is communicated digitally, but if requested, postal mail responses will be allowed as well.
- **A3**: There are no signs of potential abuse. Yet, some DSs will stop responding after asking for credentials to verify their identity.
- **A4**: In total, 4 internal employees have access to all the personal data necessary to handle a SAR.

6.4 Interview of organization Y

Organization Y has a revenue of more than 4 billion euro per year and has an international market share of at least 3%. The organization was vulnerable in 2019, but not in 2021.

- **A1**: The work of [8] had initiated the first SAR to this organization since the enforcement of the GDPR. Since then, less than 15 SARs have been received.
- **A2a**: The responsible disclosure of [8] has induced a major change to their policy. Several of the suggestions related to identity verification were taken into account, including calling the subject and asking for specific user data of the DS. In addition, an external legal organization was consulted and a DPO was appointed in order to further improve their privacy policy.
- **A2b**: The email address of the DS is always verified. Initially in 2018, we requested a photocopy of the identity card to verify a SAR due to the recommendation of the DPA. The DPA later dropped this recommendation, upon which we followed suit.
- **A3**: Due to the low number of SARs, no potential SAR abuse has been discovered.
- **A4**: SARs are received by the customer relation department and are thus collectively handled. *Authors: the exact number of employees was not available to the DPO.*

6.5 Interview of organization Z

Organization Z has a revenue of more than 1 billion euro per year and has a Belgian market share of at least 30%. The organization was not vulnerable in 2019, nor in 2021.

- **A1**: Since the enforcement of the GDPR in May 2018 up until March 2021, we have received between 100 and 200 SARs. Combining all possible requests under [22, Art. 15-17] (i.e the 'Right of Access', 'Right to rectification' and 'Right to be forgotten'), there is a total of 1200 to 1600 requests.
- **A2b**: The email address of the DS is the main identifier and verification emails are always sent to verify the identity of the requester. Sometimes, additional information is requested, such as identifier numbers or the data is sent by postal mail if the email address of the DS is not known. We initially requested a photocopy of the identity card to verify a SAR due to the recommendation of the DPA, which was regarded by some of our DSs as being 'abusive'. The DPA later dropped this recommendation, upon we followed suit.
- **A3**: We observed one case in which someone repeatedly performed SARs with the email address of another individual *(Authors: the organization did not*

3 https://ministryofprivacy.eu/privacy-tools/show-me-my-data/

provide more information related to the exact context.). Moreover, we think it is difficult to identify potential abuse. Many inquiries expire after the requester fails to deliver the correct credentials. However, we also observed another case of a divorce in which one person of the previously married couple abused the SAR process in an attempt to receive personal information of their former partner.
– **A4**: In total, 5 internal employees have access to all the personal data necessary to handle a SAR.

7 Balancing the SAR authentication procedure

Despite that our threat model consists of an adversary that has access to basic information such as the name, home address and date of birth of the DS, some organizations may not have access to that data in the first place and thus have no way to check whether the provided data is correct when verifying the DS identity of a SAR. For instance, as discussed in Section 4.3, some organizations request a photocopy of the ID card or passport of the DS without explicitly stating that they should censor the NRN. The DC may not know the NRN of the subject in the first place, and therefore unintentionally introduces an additional 'risk' of processing the newly given data, namely the NRN. The term for this type of check is coined by [2] as an *'abusive identity check'* and may also affect the data minimization principle, stated in [22][Art.5(c)]. In their paper, they provide a potential solution to this issue by creating a watermarked ID card that contains the validity period and name of the DC in order to prevent the DC from directly sharing it with other third parties. Interestingly, the authors of [8] provided us with a reply that they received of the Belgian DPA when notifying them about their research. In this reply, they advice to use a similar concept of a watermarked ID as in [2] by simply writing the name of the DC and current date on the photocopy. Unfortunately, later work has indicated that such a watermarked ID may be altered by extracting the data on the ID card and then digitally manipulating a photocopy to construct a simulated ID, thereby nullifying the watermark [3, 8]. Further in the response, the DPA also states that *'the probability of a third party having access to an ID card of someone else is much lower than the possibility of taking over someone's normal email address.'*, missing the point of a simulated ID card, despite the authors providing counterarguments. In the end, we argue that requesting a photocopy ID card to verify the identity of the DS is an objectively weak method, primarily due to the abusive identity check and the possibilities of simulated ID cards.

However, this still raises the question to which extent the authentication procedure may require specific data to verify the identity of the DS. For instance, when the DC has no knowledge about an email address or online account of the DS, the list of potential authentication credentials to verify is then reduced significantly. Although we have demonstrated the need for more stringent identification procedures, requesting too much data from the DS may be frustrating for the DS as it then becomes too difficult to exercise their 'Right of Access'. For example, some organizations require multiple pieces of user-specific information, such as the date and name of the last purchase or the serial number of a recently bought product. Although that information may be challenging for an adversary to find, it may also be burdensome for some DSs. The balance between user satisfaction, legal requirements and a prevention of a data breach by malicious SARs is arguably narrow. The question in how far the DC may go to prove the SAR is legitimate, is difficult to answer from both a technical as well as a legal perspective.

However based on prior work and our study, focused from a (primarily) security perspective, we advise DCs to request the following information (in the specific order listed) to authenticate a DS:
– If the DC has an online account of the DS on their website, then use that one to provide an automatic method integrated into the website to request their personal data. Alternatively, the website may assign a unique randomly generated identifier to each online account and show that identifier on the website of the DC. The DS should then provide that secret identifier when manually performing a SAR through for instance, email or postal mail. This way, implementing an automated process is not necessary, while still keeping the advantage of having an online account to identify against.
– If the DC has access to a valid email address of the DS, then require the DS to reply a verification code that is sent to the email inbox of the DS. This way, the DS has proven it has access to the email inbox, even if the initial SAR would have a spoofed email address. This method could also work by calling the subject by phone as it is quite similar to how most two-factor authentications are performed.
– If none of the above is possible, then specific user-data is required to verify the identity. This piece of

data may differ widely depending on the organization. Although, device cookies or serial numbers are considered to be relatively safe in case multiple of these pieces are required. Yet, user satisfaction has to be taken into consideration when performing this type of authentication as the data requested should be proportional and reasonable.

In addition, we note that care should be taken when sending the personal data of the DS by postal email. The DC has to make sure the home address is up-to-date. This way, we cover the fact that the DS may have moved their home address. Finally, reply-to headers [4] in emails may be abused to send a spoofed email which then contains a reply-to email header with the fake email address of the adversary. If the DC does not notice or remove the reply-to header, then the response to a SAR may be sent to the adversary.

8 Limitations & future work

Our threat model consists of an adversary that has access to a collection of basic personal information such as the name, email address and home address of the DS. In this paper, we argue that this information may be gathered from OSINT. However, the extraction of such information at scale is not trivial as it often requires manual labor to look for this info on social media and alike. A large automated exploitation of our attack is therefore, fortunately, difficult to perform. Future work may look more closely to the relation between personal information of a DS that is available in public data leaks and the potential abuse of SAR under the identity of that DS.

As discussed previously in Section 5.2, our SAR experiment was conducted between December 2020 and April 2021, in the midst of the COVID-19 pandemic. As many organizations still handle SARs manually and employees are often working from home, there may have been external circumstances that made our attack easier to perform. These difficulties are especially reflected when social engineering comes into play.

The reasons why we specifically chose to only replicate our prior study [8] are as followed: 1) we were the first to send out adversarial SARs. The time between the initial responsible disclosure and the study in this paper is almost 2 years, therefore giving each DC enough time to improve their policies. Nevertheless, not all studies have carried out a responsible disclosure to the vulnerable organizations. 2) we already had access to all information that was necessary to perform this study and lastly, 3) adding replications of other prior work with larger datasets would increase the amount of time considerably as our experiment is conducted by sending out postal mails. In addition, social engineering each organization is also relatively time-consuming and thus would require more human resources. Therefore, additional replications are left for future work.

Finally, the number of replies that we received on the responsible disclosures in respectively 2021 and 2019, were minimal in comparison to the organizations that were actually vulnerable. In addition, the DPO interviews in Section 6 are only a small subset. Therefore, we were careful not to make general assumptions about their replies and considered their replies to be individual cases. Despite the small dataset, we argue that the interviews still contain a significant amount of information, useful for future work. Nevertheless, future work may look into interviewing a larger set of DPOs (or other privacy employees) in order to have broader view of the internal policies and issues these organizations have when handling SARs.

9 Conclusion

In this paper, we have discussed the different threat models and methodologies of potential attacks to SAR policies and the requested credentials that are present in such processes. Here, we have shown that some methodologies of prior work have vastly different assumptions, making them difficult to compare. Furthermore, we have demonstrated that adhering to a responsible disclosure guideline does not significantly improve the SAR policies of most vulnerable organizations as 53% of them are still vulnerable to leaking personal information of their data subjects. This data includes, but is not limited to, financial transaction details, location history and online shopping behavior. In addition, the replies received from some vulnerable DCs indicate room for improvement in the DC's technical knowledge to properly assess cybersecurity risks. Therefore, we have proposed several changes in SAR processes to reduce their risks. Finally, we gained detailed insights into individual SAR policies of some organizations by interviewing several DPOs and observed potential signs of malicious SAR abuse.

4 https://datatracker.ietf.org/doc/html/rfc4021, section 2.1.4

Acknowledgements

This research was funded in part by the Bijzonder Onderzoeksfonds (BOF). We thank all organizations that responded to our responsible disclosure and policy suggestions. In addition, our sincere gratitude goes to the DPOs and privacy employees that were willing to participate in our interview and the subjects who have participated in this study. Finally, we thank the reviewers for their feedback and interesting discussions.

References

[1] Ausloos, J., and Dewitte, P. Shattering one-way mirrors - data subject access rights in practice. International Data Privacy Law 8, 1 (03 2018), 4–28.

[2] Boniface, C., Fouad, I., Bielova, N., Lauradoux, C., and Santos, C. Security Analysis of Subject Access Request Procedures How to authenticate data subjects safely when they request for their data. In Annual Privacy Forum (2019).

[3] Bufalieri, L., Morgia, M. L., Mei, A., and Stefa, J. GDPR: When the Right to Access Personal Data Becomes a Threat. In 2020 IEEE International Conference on Web Services (ICWS) (2020), pp. 75–83.

[4] Cagnazzo, M., Holz, T., and Pohlmann, N. GDPiRated – Stealing Personal Information On- and Offline. In Computer Security – ESORICS 2019 (Cham, 2019), K. Sako, S. Schneider, and P. Y. A. Ryan, Eds., Springer International Publishing, pp. 367–386.

[5] CCPA. California Consumer Privacy Act, 2018. Cal. Legis. Serv. Ch.55 (A.B. 375).

[6] Cormack, A. Is the Subject Access Right Now Too Great a Threat to Privacy? European Data Protection Law Review 2 (2016), 15–27.

[7] Das, S., Kim, A., Jelen, B., Streiff, J., Camp, L. J., and Huber, L. Towards Implementing Inclusive Authentication Technologies for Older Adults. In Who Are You?! Adventures in Authentication Workshop (Santa Clara, California, USA, Aug. 2019), WAY '19, pp. 1–5.

[8] Di Martino, M., Robyns, P., Weyts, W., Quax, P., Lamotte, W., and Andries, K. Personal Information Leakage by Abusing the GDPR "Right of Access". In Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (2019), SOUPS'19, USENIX Association, p. 371–386.

[9] Galetta, A., Fonio, C., and Ceresa, A. Nothing is as it seems. The exercise of access rights in Italy and Belgium: dispelling fallacies in the legal reasoning from the 'law in theory' to the 'law in practice'. International Data Privacy Law 6 (11 2015), ipv026.

[10] Google Inc. Stronger security for your Google Account. https://www.google.com/landing/2step/, accessed on April 21st 2021.

[11] Herrmann, D., and Lindemann, J. Obtaining personal data and asking for erasure: do app vendors and website owners honour your privacy rights? In Sicherheit 2016 - Sicherheit, Schutz und Zuverlässigkeit (Bonn, 2016), M. Meier, D. Reinhardt, and S. Wendzel, Eds., Gesellschaft für Informatik e.V., pp. 149–160.

[12] Kröger, J. L., Lindemann, J., and Herrmann, D. How Do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on IOS and Android Apps. In Proceedings of the 15th International Conference on Availability, Reliability and Security (New York, NY, USA, 2020), ARES '20, Association for Computing Machinery.

[13] Kutyłowski, M., Lauks-Dutka, A., and Yung, M. Gdpr – challenges for reconciling legal rules with technical reality. In Computer Security – ESORICS 2020 (2020), L. Chen, N. Li, K. Liang, and S. Schneider, Eds., Springer International Publishing, pp. 736–755.

[14] Mahieu, R. L. P., Asghari, H., and van Eeten, M. Collectively exercising the right of access: individual effort, societal effect. Internet Policy Review 7, 3 (2018).

[15] Markert, P., Farke, F., and Dürmuth, M. View The Email to Get Hacked: Attacking SMS-Based Two-Factor Authentication. In Who Are You?! Adventures in Authentication Workshop (Santa Clara, California, USA, Aug. 2019), WAY '19, pp. 1–6.

[16] Mustafa, H., Xu, W., Sadeghi, A. R., and Schulz, S. You Can Call but You Can't Hide: Detecting Caller ID Spoofing Attacks. In 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (2014), pp. 168–179.

[17] Pavur, J., and Knerr, C. GDPArrrrr: Using Privacy Laws to Steal Identities. CoRR abs/1912.00731 (2019).

[18] Petrlic, R. Identitätsprüfung bei elektronischen Auskunftsersuchen nach Art. 15 DSGVO. Datenschutz und Datensicherheit - DuD 43, 2 (Feb. 2019), 71–75. (German).

[19] Samarin, N., Kothari, S., Siyed, Z., Wijesekera, P., Fischer, J., Hoofnagle, C., and Egelman, S. Investigating the Compliance of Android App Developers with the CCPA. In 5th Workshop on Technology and Consumer Protection (ConPro '21) (2021), Association for Computing Machinery.

[20] Syrmoudis, E., Mager, S., Kuebler-Wachendorff, S., Pizzinini, P., Grossklags, J., and Kranz, J. Data Portability between Online Services: An Empirical Analysis on the Effectiveness of GDPR Art. 20. Proceedings on Privacy Enhancing Technologies 2021, 3 (2021), 351–372.

[21] The European Parliament and the Council. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data . OJ L 281 (November 1995).

[22] The European Parliament and the Council. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). OJ L 119 (May 2016), 1–88.

[23] Urban, T., Degeling, M., Holz, T., and Pohlmann, N. "Your Hashed IP Address: Ubuntu.": Perspectives on Transparency Tools for Online Advertising. In Proceedings of

the 35th Annual Computer Security Applications Conference (New York, NY, USA, 2019), ACSAC '19, Association for Computing Machinery, p. 702–717.

[24] Urban, T., Tatang, D., Degeling, M., Holz, T., and Pohlmann, N. A Study on Subject Data Access in Online Advertising After the GDPR. In Data Privacy Management, Cryptocurrencies and Blockchain Technology (Cham, 2019), C. Pérez-Solà, G. Navarro-Arribas, A. Biryukov, and J. Garcia-Alfaro, Eds., Springer International Publishing, pp. 61–79.

A Appendix

A.1 Contents of registered letter

As we would like to prevent real malicious adversaries to potentially copy and subsequently abuse our methodology, we only provide a summary of the contents of the registered letter below.

- Introduction with name of DS.
- We started by explaining (with random but logically chosen dates) that we have attempted to contact the organization several times by sending a SAR through the *original email address of the DS*, but no response was received after 30 calendar days since the emails bounced back.
- Next, we explained that, because the emails bounced back, we have sent an email to the DC through the *fake email address of the DS* and stated that this email did not bounce back.
- We closed our explanation with the fact that the first SAR was requested more than 30 days ago, pressuring them to immediately provide answers to the SAR.
- We 'repeated' our original SAR in this letter. This inquiry consists of formal requests for several pieces of information such as the personal data, retention period and the existence of automated profiling, including the references to the appropriate GDPR articles.
- Finally, we provided information from the DS to verify the identity. This information consists of the name and a simulated copy of the ID card. In addition, we requested to send the reply to the fake email address of the adversary or the original email address of the DS. We explicitly included the permission to send the answers to the original email address, in order to not raise any suspicion (yet, our threat model does not require an adversary to have access to the original email address of the DS. See the 'safe' classification in Section 5.2).
- We closed our letter by stating, again, that the period in which the DC should have answered our SAR has already expired.

Please note that there was no email communication with the DC and the DS (or adversary) *before* sending the registered letter. The excuse related to the 'emails bouncing back' is an attempt to socially engineer the DC.

A.2 Organizations appointed to subjects

In Table 5, we 1) specify the number of organizations that has personal data from each subject and 2) specify the final number of organizations that were contacted through our experiment by each subject. Each organization was contacted by exactly one subject. The assignment of each organization to a subject is chosen as follows:

- If the organization only had personal data of one subject, then the organization was assigned to that subject.
- Otherwise, no specific method was applied to assign the organization to the subject as the organization was then randomly assigned to any of the 5 subjects, while making sure that the final number of contacted organization for each subject is approximately equal.

Since these organizations are relatively popular, they all had personal data from *at least* one of the subjects.

	Number of accounts	Number of contacted orgs.
Subject A	13	7
Subject B	11	7
Subject C	13	8
Subject D	12	7
Subject E	20	11
		= 40

Table 5. Specifies how many organizations were appointed to and contacted by each subject. The number of accounts indicates the number of organizations from the dataset where the mentioned subject has *at least* some personal data (e.g. an online account).

A.3 Organization classification

In our experiment, we classified each organization into a category, according to the market sector they are most active in. This classification is proposed by [8].

- Financial (Fin_x)
- Retail (Ret_x)
- Entertainment (Ent_x)
- Transport and Logistics (Trl_x)
- News Outlets or Publishers (New_x)
- Any other organization that does not pertain to any of the above categories (Oth_x)