Towards GDPR-compliant data processing in modern SIEM systems

Florian Menges[1], Tobias Latzo[2], Manfred Vielberth[1], Sabine Sobola[4], Henrich C. Pöhls[3], Benjamin Taubmann[3], Johannes Köstler[3], Alexander Puchta[1], Felix Freiling[2], Hans P. Reiser[3], and Günther Pernul[1]

[1] Department of Information Systems, Universität Regensburg, Germany
[2] Department of Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
[3] Institute of IT-Security and Security Law, Universität Passau, Germany
[4] Paluka Sobola Loibl & Partner attorneys at law, Regensburg, Germany

Abstract. The introduction of the General Data Protection Regulation (GDPR) in Europe raises a whole series of issues and implications on the handling of corporate data. We consider the case of security-relevant data analyses in companies, such as those carried out by Security Information and Event Management (SIEM) systems. It is often argued that the processing of personal data is necessary to achieve service quality. However, at present existing systems arguably are in conflict with the GDPR since they often process personal data without taking data protection principles into account. In this work, we first examine the GDPR regarding the resulting requirements for SIEM systems. On this basis, we propose a SIEM architecture that meets the privacy requirements of the GDPR and show the effects of pseudonymization on the detectability of incidents.

Keywords: Security Information and Event Management · SIEM · GDPR · Threat Intelligence · DINGfest.

1 Introduction

1.1 Motivation

The security of the modern information infrastructure is of high importance. In order to detect misuse and attacks at an early stage a lot of information about the events inside IT-infrastructures, e.g. inside computer networks and software applications also across many systems, is required to detect or post-mortem report and document attacks. *Security Information and Event Management* (SIEM) systems help organizations to keep up with the ever increasing complexity by providing a holistic view on IT-infrastructures. Naturally, SIEM systems process enormous amounts of data about security related events, e.g., when specific users login or certain users perform critical actions. It is often argued, that generally

the quality of service depends critically on the quality and detail of the data collected and processed within the system [48], which has been shown for different domains such as threat intelligence [37].

Events like those just described that are processed within the SIEM system are clearly related to concrete users and therefore must be treated as personal information, which require protection under Europe's *General Data Protection Regulation* (GDPR) [13]. Adopted in May 2018, it regulates and harmonizes the protection of personal data in the processing and transfer of data within and between private companies and/or public bodies in the European member states. Although, the GDPR is only compulsory for EU member states, it has evolved into a blueprint for data protection all over the world, as discussions between the US Congress and Mark Zuckerberg in the aftermath of the Cambridge Analytica case indicate[5].

Hence, SIEM systems must also comply to the regulations themselves, which leads to conflicting interests. On the one hand, SIEM systems rely on personal data such as information from the identity and access management (IAM) for providing high detection rates of incidents and thus a high level of protection. On the other hand, the requirements of the GDPR suggest that investigations of data streams as carried out in current SIEM systems may no longer be legally compliant. To complicate things even further, regulations regarding the handling of digital evidence mandate that authenticity and integrity of the data related to an incident should be guaranteed at all times in order to maintain its high legal probative value. It is therefore necessary to find the best trade-off between those two demands. With this work we attempt to fill the resulting research gap and to harmonize legal GDPR requirements with the technical architecture for SIEM systems. To bridge the gap between the disciplines of computer science and law and to produce the most reliable results possible, this paper was written by IT security researchers in collaboration with a lawyer A central idea is the integration of *anonymization* and *pseudonymization* into threat analytics mechanisms. While this makes it necessary to change the original data, it is possible to maintain legal integrity and authenticity by using *redactable and sanitizable signatures*, a cryptographic concept to retain a level of authenticity useful to retain a suitable level of legal evidence even when data gets obfuscated or if certain parts of it are missing. We deploy cryptography to enable balancing authenticity proofs for the collected security-related events with the confidentiality requirements of the information about commercially-relevant internals (trade secrets) and employees' as well as customers' privacy (personal data). Thus, our goal is to minimize the amount of data which is being made accessible to third-parties in every step of the SIEM process. By enforcing this with cryptography the proposed system adheres to the security-by-design principle of least privilege as well as the privacy-by-design principle of data minimization. At the same time we aim to keep the impact on detection as low as possible and thus we provide an audit-able process to gain access to more details if security analysis is needing

[5] https://www.theverge.com/2018/4/11/17224492/zuckerberg-facebook-congress-gdpr-data-protection

it. For the reason of being able to reconstruct original data, leaving a trace in an audit log, we focus on cryptographic methods and support pseudonymization rather than anonymization. Technically, we encrypt and sign events early and store the decryption keys with a party trusted for logging access to stored keys; moreover we employ signatures that allow to slice or redact data.

1.2 Related Work

When looking at the application of privacy mechanisms to threat analytics (e.g. SIEM systems), literature can be divided into a pre-GDPR and a post-GDPR phase, as this regulation still has a big impact on the integration of privacy. In the former phase there are not many results to be found regarding applying privacy to SIEM systems, however the challenges in integrating privacy in forensic and threat analyses has been identified [40,17]. Although the challenges were not solved for SIEM systems, selected works in the IT security domain address it. For example Burkhart et al. [6] describe a privacy preserving solution for secure multi-party computation. Furthermore, a main focus during this era was the application of privacy to intrusion detection systems (IDS), which could be declared as the predecessors of modern SIEM systems and thus in our context are worth a closer look: Sobirey et al. [39] propose an approach for pseudonymizing user related data in IDS and closely examined, which records need to be pseudonymized in audit records. Based on this work, Biskup and Flegel [3] and Park et al. [28] propose an approach which is quite similar to the one presented in this paper as it uses cryptograpic methods to pseudonymize personal data, though these are closely tailored to IDS and not completely adaptable to SIEM. In addition, they were issued before the publication of the GDPR and thus did not have all the requirements in mind and respectively were not evaluated against the new requirements. Our approach also differs,as we have a more abstract view of the whole system and do not focus largely on cryptographic details. Furthermore, Buschkes and Kesdoğan [7] discuss requirements such as data avoidance and reduction of personal data.

Although privacy preserving methods were widely discussed in the past, recently the application of GDPR received an increased amount of attention and new works were published. In relation to SIEM, some work was published covering GDPR compliant data processing. Sgaglione and Mazzeo [38] and Coppolino et al. [8] introduce the COMPACT project, which is a GDPR compliant SIEM. However, they do not go into detail, how this is realised technically. Current research for SIEM systems mainly focuses on the architecture and improvement of such systems and not on the integration of privacy [26,25,27].

In cryptography, digital signatures are used to ensure authenticity and integrity of data, i.e., they guarantee that upon inspection data is unchanged and comes from an attributable source. Special techniques of *redactable signature schemes* (RSS) by Steinfeld et al. [41,18] allow subsequent deletions in the data, while *sanitizable signature schemes* (SSS) as proposed by Atieniese at al. [1] even allow subsequent edits by dedicated authorised parties while maintaining authenticity of the remaining data. Both RSS and SSS allow to balance authenticity

with privacy protection, because they allow retaining the integrity and authenticity protection for the unedited or not-removed parts of the document and at the same time keep the confidentiality protection for the overwritten parts of the document. In cryptography the latter property is intuitively termed privacy. While many schemes have appeared in the literature [2], only some of which uphold privacy and only those schemes that additionally fulfil detectability, known as non interactive public accountability [5] can be used for eIDAS[6] compliant signatures [30,31,45]. While the legal compliance of such signatures has been subject to research, the integration of such schemes into privacy protection of SIEM have not yet been investigated. In particular, in the application scenario of SIEM we want to be able to later reveal previously not-shared content. For this, a special form of digital signatures is needed which has the property of *mergeability*, i.e., the ability to re-add signed content to previously redacted but still signed content and re-generate a valid signature over the merged content [34].

1.3 Contribution and Outline

To the best of our knowledge we are not aware of an approach, that integrated GDPR into SIEM in a comprehensive way. Given the fact that these regulations need to be applied by all companies that operate within the European Union, there appears to be high demand for systems that are GDPR compliant. In this paper we present the first privacy-friendly – and thus GDPR-compliant – SIEM architecture that protects the confidentiality as well as the authenticity of security-relevant events starting at their collection, keeping the protection during the analysis and finally sending an incident report.

The presented architecture allows the deployment of a SIEM that meets the regulatory requirements under the EU data protection. It protects personal information in the data sets from unnecessary visibility using pseudonymization and encryption techniques without a significant reduction in detectability. Hence, we balance data-quality (detection of incidents) with legal obligations from privacy legislation and thus also protect trade secret by sharing only the minimum necessary information in any step of the SIEM process. Thus we strongly adhere to the GDPR's data minimization principle. Still, we achieve the highest level of confidence that the security-relevant events initially recorded and reported into the SIEM process are protected from tampering by using redactable and sanitizable signature schemes to proof authenticity. This allows us to balance the need for generating data with a high legal evidence with the need to protect privacy (and trade-secrets).

The legal analysis carried out for this architecture and presented in this paper shows that even potentially invasive data can be collected in a GDPR-compliant manner as our proposed system balances the necessity of the collection (detection and reporting of actual security incidents) with the protection of users

[6] eIDAS is short for the current legislation which defines the technical functionalities to allow electronic signatures to be legally equivalent to handwritten signatures within the EU [12].

privacy and customer's trade-secret needs. The paper shows the actual influence of pseudonymization on incident detection mechanisms and the results of the performed legal evaluation.

The remainder of this paper is structured as follows. First, we give some background on the GDPR (legal) and SIEM (technical) in Section 2. In Section 3 we develop the research questions that arise from integrating GDPR into SIEM. On this basis, we describe our GDPR-compliant architecture for SIEM systems in Section 4. The architecture is the evaluated on both technical and legal level in Section 5. The paper concludes in Section 6.

2 Background

This section provides the background information that is needed to understand the approach presented in this paper. Thus, we first give an overview of the functionality and properties of SIEM systems, as this serves as a basis for our architecture. Subsequently, we give an overview of the requirements the GDPR defines with special attention to the processing of personal data.

2.1 Security Information and Event Management (SIEM)

In general, SIEM was first mentioned by Gartner [49]. It originated from the initially separate systems Security Information Management (SIM) and Security Event Management (SEM) [15]. SIEM must fulfill several requirements, which are all connected: Log collection, enrichment with context data, log normalization, event correlation, and analysis as well as long- and short-term storage of log data, reporting, monitoring, alerting, and incident response [24,47,14].

A SIEM system as described in [47] is essentially designed for collecting relevant log data in a central place from arbitrary systems such as network devices or operating systems. This among other things enables the detection of incidents and in this way gaining situational security awareness. On a high level of abstraction, a SIEM system consists of the three main steps *data acquisition*, *processing* and *reporting*, which are elaborated in the following in more detail.

Data acquisition: Hereby, it first collects relevant event information, in most cases in the form of log data, which gets enriched with additional context data. There are basically two approaches for data acquisition: First, the data can be pushed into the SIEM by the data generating system. Thereby, the SIEM does not influence the generated data. Second, the data can be pulled by the SIEM from the observed system, which grants more control over the generated data enabling for example the assurance of integrity. This data then is translated into a uniform representation during the normalization step.

Processing: The core of the system is the correlation and analysis component, wherein information from various sources is correlated and incidents get detected by methods such as pattern matching. Real-time threat detection

enables fast reactions in case of an incident, whereas forensic analysis pursues the goal of analyzing the whole extent of the event in the aftermath in order to secure evidence. A distinction can, therefore, be made between short-term and long-term storage of relevant data. For long-term storage, it is particularly important to preserve the data in a tamper-proof way in order to be able to use it as evidence in court. Monitoring and visual security analytics enable security analysts to be actively involved in the analysis process. In the case of a detected incident, alerting and incident response triggers necessary reactions to mitigate further harm.

Reporting: An essential part of modern SIEM systems is reporting occurred incidents for compliance reasons (e.g. critical infrastructure providers) or enables the participation in established threat intelligence sharing platforms between participating organizations.

2.2 GDPR

In the following we present some background on general and SIEM-specific GDPR demands and later (see Sect.3) detail, which problems arise for SIEM systems to be built to comply.

Since 25 May 2018, the GDPR has been in force throughout the European Union (EU) to ensure the protection of the "natural person" in data processing. This regulation is directly applicable to all member states of the EU. The GDPR does not contain any immediate legal requirements for software developers but if they intend to sell their product to customers, they must be aware of the legal requirements.

Data protection is the protection of the natural person from privacy impairments through the processing of data concerning the person. Everyone should be free to decide who, when and how their data should be accessible. The term of personal data is therefore defined as "all information relating to an identified or identifiable natural person", Art. 4 (1) GDPR.

General principles of data processing In order to achieve the goal of high personal protection, the GDPR pursues the regulation of all basic principles which are regulated in Art. 5 (1) GDPR. According to the GDPR, only lawful, fair and transparent data processing is permitted in a transparent manner, Art. 5 (1) lit. a GDPR, in order to serve the principle of *good faith*. Furthermore, data processing shall be lawful only if and to the extent that it is applied by Art. 6 GDPR. But even then, it must be done in a transparent way which is traceable for the data subject. Another important principle is the *purpose limitation*, Art. 5 (1) lit. b GDPR. Thus, the clear purpose of processing must be previously established and legitimate. Furthermore, the *principle of data minimization* is applicable, Art. 5 (1) lit. c GDPR. This means that data collection is only allowed within limits for specified, explicit and legitimate purposes and not further. A further principle is *accuracy*, Art. 5 (1) lit. d GDPR. Only correct data may be collected. Even after processing it must be ensured that personal data is

accurate. If this is not the case, data must be erased or rectified with delay. One further fundamental principle is the storage limitation, Art. 5 (1) lit. e GDPR. In the course of data processing it must be ensured that an identification is only possible in case of it being necessary for the purpose. *Integrity and confidentiality* have recently become further important principles, Art. 5 (1) lit. f GDPR. This means that processing must occur in compliance with general safety standards. Compliance with these principles must be proven at any time by the controller, Art. 5 (2) GDPR.

Processing on a legal basis and transparency obligations Generally, there must be a legal basis for the processing, otherwise all collection of personal data is considered unlawful. The exceptions are regulated by Art. 6 GDPR. First of all, processing is lawful if the data subject has given *consent* to data processing, Art. 6 (1) point (a) GDPR. Processing is also allowed for the performance of contracts, Art. 6 (1) point (b) GDPR. If the processor is subjected to a legal obligation, data processing is also lawful without further requirements, Art. 6 (1) point (c) GDPR. This also applies if the protection of vital interests is pursued, Art. 6 (1) point (d) GDPR. The processing is also possible for the performance of a task carried out in the *public interest* or in the *exercise of official authority*, Art. 6 (1) point (e) GDPR. Lastly, processing is lawful for the purposes of the *legitimate interests pursued by the controller or by a third party*, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, Art. 6 (1) point (f) GDPR. While a lot of the exceptions might be triggered by the need and want to have a SIEM to protect from security breaches, the collection must meet a balance test, e.g. collection must not overshoot the goal, and must be transparent, e.g. clearly communicated to the data subjects.

Data security In addition, the GDPR regulates a close interconnection of data protection with data security (especially Art. 32 GDPR) to the effect that technical and organizational measures in the data-processing company enable the highest degree of data security (availability, confidentiality, integrity, and resilience). In our scenario, the two security goals confidentiality and integrity are particularly relevant. We must ensure that information about events that could relate to incidents is transferred from the source to the sink, and is not altered or made accessible to unauthorised persons. To protect integrity an unauthorized modification must be detected if it happened, which is especially important to use non-tampered recorded data as evidence. Thus, protecting integrity and authenticating the data's origin provides legal value. Further, confidentiality protection guarantees that no unauthorized party is able to obtain information not intended for them, e.g. we must securely communicate the personal data to have them reach only the right recipients.

Redaction The term "redaction" itself is not found in the GDPR directly; it refers to the irreversible removal of the information [43]. This process is explicitly

mentioned in guidance documents that explain how to remove information that is not subject to the information to be released under laws for the freedom of access to information, e.g. UK FOIA [29] and thus is also applicable as a technique in the context of GDPR's data minimization [44,43].

Pseudonymization The term "pseudonymisation" of the GDPR means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organizational measures to ensure that the personal data are not attributed to an identified or identifiable natural person, Art. 4 (5) GDPR. In order to make a pseudonymization, the data subject must first be assigned to pseudonyms. These can be user IDs. Thereafter the necessary data for identification must be kept separately. It must be ensured that they are strictly separated from the pseudonyms. It is important that features that can only indirectly lead to identification must be removed in the event of pseudonymization. The pseudonymization can be made by the data subject himself, the controller or an independent third party, a data trustee. An assignment rule must be created, e.g. through a reference table. The pseudonymization must be performed without the knowledge of the data subject. The data subject must be informed about the pseudonymization and it must be clarified who generates the pseudonym, who owns the assignment rule and under what circumstances an identification may take place. This is because pseudonymized data continues to be personal data. Particularly in the area of monitoring software development compliance with the GDPR is mandatory because it extracts its results from a data stream that contains personal data. The pseudonymization protects the data of the data subjects. At the same time, it is an opportunity to still being able to detect security incidents effectively and identify the responsible user in this regard.

3 Problem Statement and Research Questions

In SIEM systems very large amounts of data are processed from various sources, while at the same time a GDPR compliant data protection must be guaranteed. This can be achieved by protecting all data relevant to data protection against unauthorized access using techniques such as encryption or pseudonymization. Working on protected data, however, brings different additional problems with it. On the one hand, it needs to be ensured that incident recognition is still possible despite the data protection. On the other hand it also needs to be possible to remove the protection in case of an actual incident. These aspects, which we have identified as essential for a GDPR-compliant analysis process, translate into the following three specific research questions. In this work we use the pseudonymization for the realization of the data protection, since this represents a valid procedure according to both GDPR and different reporting regulatory environments.

3.1 Data Protection considerations and attacker model:

SIEM systems work with data from highly heterogeneous sources. As a result, different requirements need to be met in order to enable data protection in accordance with the GDPR. Establish data protection through the full encryption of all data would be the most intuitive and legally compliant way to process the data. However, since the GDPR only requires the protection of personal data, the data can also be classified according to protection requirements and partially peudonymized in this context. In this way it can be achieved to still be legally compliant, while more meaningful data is available for analysis at the same time. To achieve this, all acquired data needs to be available in a standardized form to allow the identification of information that needs to be protected. More specifically, the data acquired can be differentiated into information that is not relevant for data protection, data that may be relevant for data protection (e.g. path information in folder structures) and specific information relevant for data protection (e.g. e-mail addresses or contents of e-mails). In addition to this, the pseudonymization mechanism also needs to be protected. It must be ensured in a technical and organizational way for each data processing step within the SIEM system.

For being able to design a compliant and secure system, it is conducive, to define an attacker model, that determines the necessary measures. Thereby, the role, the goal, behaviour and the resources of the attacker are delimited:

- **Role:** The attackers role against which we consider our system protected can either be an outsider or an insider. An outsider is any person who has only access to interfaces of the system, which are open to the public. The outsider can however utilize a breach to gain access to certain parts or data of the system. Any third party who is involved in the SIEM system can also be referred to as an outsider. In contrast, an insider is any person who is directly involved into the system, such as analysts or server-admins.
- **Goal and behaviour:** The attacker can be either passive or active. The passive attacker only lists to the data without any intervention, whereas the active attacker tries to gain access to the data or the system by actively interacting with the system. For our approach, the considered goal of the attacker is to gain access to private data, since we design a SIEM system which is GDPR compliant.
- **Resources:** Since we utilize measures which are based on common asymmetric or symmetric cryptography, we can only consider attackers, with limited resources.

In summary, this raises the first question:

Q1: How must data that is processed in SIEM systems be protected to be GDPR compliant?

3.2 Impairment of Incident Identification through Data Concealment:

The GDPR stipulates that data protection must be applied as early as possible within the analysis process. Considering the data management of SIEM systems, this translates into a data protection obligation at the time of data acquisition. As a result, incident detection always needs to be performed on pseudonymization data. In this context, the relationship between pseudonymized data and plain text data within the data stream is a significant factor influencing possible analyses. This may impair both automated and manual analyses due to possible losses in the meaningfulness of the data analyzed. This leads to the second question:

Q2: Does the recognition of security incidents function properly despite of data pseudonymization or may losses and trade-offs be expected here?

3.3 Lifting the Pseudonymization while retaining Data Authenticity:

To enable the utilization of information about detected security incidents, while being compliant with the legislation, two main conditions need to be met. On the one hand, information about incidents must be available in the long term in an integrity preserving manner. On the other hand, a de-pseudonymization of the data needs to be possible at any time after detection. This warrants that the information can be used as a reliable means of evidence in trials that might take place in the future. Please note that the GDPR proposes both anonymization and pseudonymization techniques as possible data protection measures. However, since the use of anonymization would prevent the data from being used as evidence, this technique will not be considered further in this paper.

To achieve this, appropriate technical and organizational measures need to be in place ensuring that de-pseudonymization is only possible in case of actual incidents. Additionally, legal compliance also needs to be ensured for the data after lifting the pseudonymization for further processing. This requires protecting the data's integrity including origin authenticity.

Thus, the principle of data minimization complicates especially the goal of integrity. This principle has always been at the center of data protection and can be found in European and member state legal texts, e.g. already in the former Directive 95/46/EC and thus also in the GDPR. In detail Art. 5 GDPR describes that personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

In general, there are two ways to conform with data minimization: (a) not collect or not forward personal data if it unnecessary, e.g. deleting it from data sets or blacken it out, i.e. redact it, or (b) making it harder to restore the personal data. For the latter, an important measure in this regard is the pseudonymization of data because it reduces the risks for the data subject and simultaneously

it helps the controller to fulfil his data protection obligations (see also recital 28). By, for example, ensuring that only pseudonymized data is used during the incident analysis and that the person is only revealed if an anomaly is detected, would make a legally compliant processing of personal data in a SIEM conceivable. For the former, personal data, e.g. fields that contain this information, could just be removed before forwarding them.

On the other hand, both mechanisms for data minimization (pseudonymization or removal) result in a modification of the initially gathered data; an intended modification but a modification nevertheless. This results in standard cryptographic mechanisms to protect the data's integrity, such as digital signatures or message authentication codes, to fail. Thus, they are unsuitable to protect the integrity end-to-end. The more recent cryptographic mechanisms, known as redactable [18,41] or sanitizable signatures [1] are capable of allowing our architecture to authorize modifications such as removal of unnecessary data points from authentic data set gathered by the SIEM. From their initial versions these algorithms have evolved (see [2] for an overview). Most recently they undergo the process of becoming an internationally recognized signature standard[7]. While this process takes time, the current status shows that the cryptographic mechanisms have the needed maturity to be backed and accepted in the cryptographic community. Once becoming recognised through such a standard, legal argumentation for compliance becomes a lot easier as legislators and judges will find the algorithms in lists of known mechanisms. Even if not standardised (or not yet) the provided authenticity offerings are technically equivalent to normal signatures [32,16] and in any case much better than having none and also non-standard algorithms are suitable to win legal arguments in court cases – bearing the need for technical expertise appointed by court. This leads to the third question:

Q3: Which conditions need to be met to ensure that incident information can be de-pseudonymized in case of an incident and how can it be used as means of evidence?

4 Conceptualizing a GDPR-compliant SIEM System

Although SIEM systems have grown to mature security tools, privacy has largely been neglected in this area. Thus, we have previously defined central research questions that arise when applying the GDPR regulation. To answer these questions, we propose a SIEM architecture that is compliant with GDPR, while largely preserving its functionalities in this section. Therefore, we propose concrete solutions for each of the individual research questions based on an extended, GDPR-compliant architecture.

[7] ISO/IEC 23264 Redaction of Authentic Data https://www.iso.org/standard/78341.html [last accessed: Jan. 2020]

4.1 DINGfest base Architecture

This section gives an overview on our general security monitoring architecture and assigns the previously defined research questions to the respective areas of the architecture. The presented architecture is based on the general DINGfest architecture as presented in [23] and extends it by data protection measures and the resulting GDPR compliance. DINGfest is a research project that aims at improving the detection, forensic analysis and the reporting of detected incidents. The project started June 2016 and will finalize at the end of 2019 and is funded by the German Federal Ministry of Education and Research. The general system monitoring architecture is illustrated in Fig. 1. It consists of three main modules – namely data acquisition, data analysis and incident reporting located within an organization and shows an external authority possible counterpart for the receipt of detected incidents. The counterpart is intentionally included in the architecture design, since its role and the management of the data flows generated there are one central factor in ensuring the legal compliance of the system. This concerns data protection requirements according to the GDPR on the one hand and may also concern existing statutory reporting obligations of the organization.

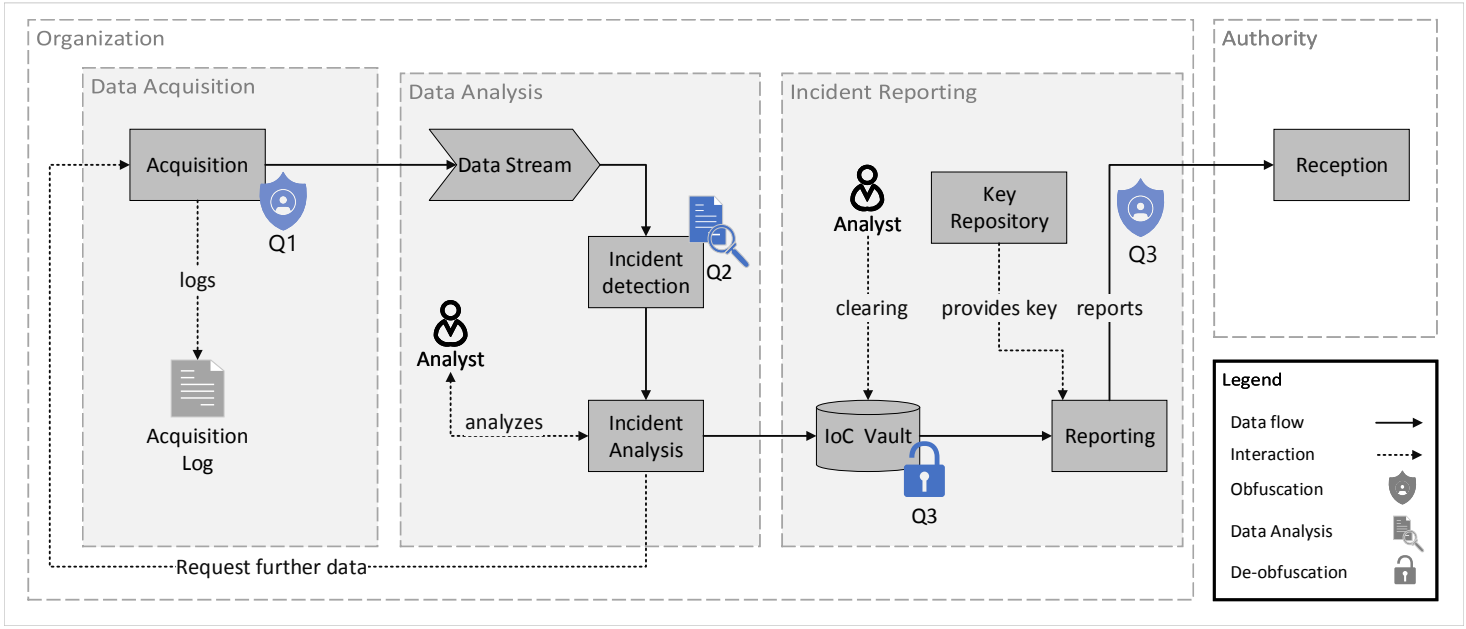

Fig. 1. DINGfest Base Architecture

The **data acquisition** module collects data from all monitored computing resources in the company. This data may contain personal data of employees and customers that needs to be protected. The monitored resources are not only computing devices like workstations, servers and mobile devices, but also network devices like routers and switches. The actual data is obtained from various sources. This includes, for example, data extracted with the help of Virtual Machine Introspection (VMI). This also includes data obtained from system log files or incident information provided by human sensors [46]. Moreover, all data extractions within the acquisition area are stored in the acquisition log. This enables a later auditability of all the information obtained. The extracted data is finally pushed into a larger data stream, which serves as data basis for the data analysis section. Data acquisition is the starting point at which all data (including personal data) is transferred to the system. As a result, research question Q1 must be addressed within this module to show how data must be protected or pseudonymized to ensure GDPR-compliant data handling. This additionally generates the required prudential value for the gathered evidence.

The **data analysis** module analyzes the whole data stream and tries to detect security incidents using a combination of fingerprinting and pattern recognition. If the detection engine discovers a potential security violation it generates an incident alert that contains a description of the assumed violation and the related data records. The alert is then received and analyzed by a forensic analyst. The analyst can use a visual analysis interface and request additional data from the data acquisition module. Should the suspicion be confirmed, the incident is forwarded to the reporting module. Otherwise, the incident alert is deleted right away. As shown above, the data acquired during data acquisition must be protected. This makes data analysis more difficult, since information is lost as a result of pseudonymization. Therefore, the research question Q2 will be addressed within this module to show how far incident detection with disguised data is still possible.

The tasks of the **reporting module** include the long-term storage of analyzed incidents (usually between several weeks and several year depending on the local legislation) and the reporting to local authorities in accordance with the legislation in force. Arriving incidents are therefore stored in a database and processed by an incident reporter. During this process the reporter might query the database to contrast the current incident with past incidents. Eventually a report is generated and forwarded to local authorities, in order to inform them or comply with regional regulations. Such reports may contain information about innocent individuals or company assets which also need to be protected. Within this module the research question Q3 will be addressed. The aim is to ensure that information can be de-pseudonymized in the event of an actual incident, while preserving its integrity. This is necessary to enable the use of the data as evidence in possible later court cases. Furthermore, it must also be ensured that the data can be reported to the appropriate authorities in compliance with the law and data protection regulations.

4.2 GDPR compliant Data Processing

In the following we present our approach in more detail, especially relevant parts, which enable GDPR compliant data processing inside SIEM (Q1). To this end, we propose an approach that pseudonymizes personal data at relevant points and at the same time allows to de-pseudonymize this data in case of a detected incident (Q3) in compliance with GDPR regulations. Fig. 2 shows the basic structure of this approach. To achieve a pseudonymization of the information, cryptographic methods are used. These on the one hand prevent access to personal data by encrypting it and on the other hand allow decrypting it under certain constraints specified by the GDPR. However, the decisive question is where the data must be encrypted and how to implement the key management for encryption.

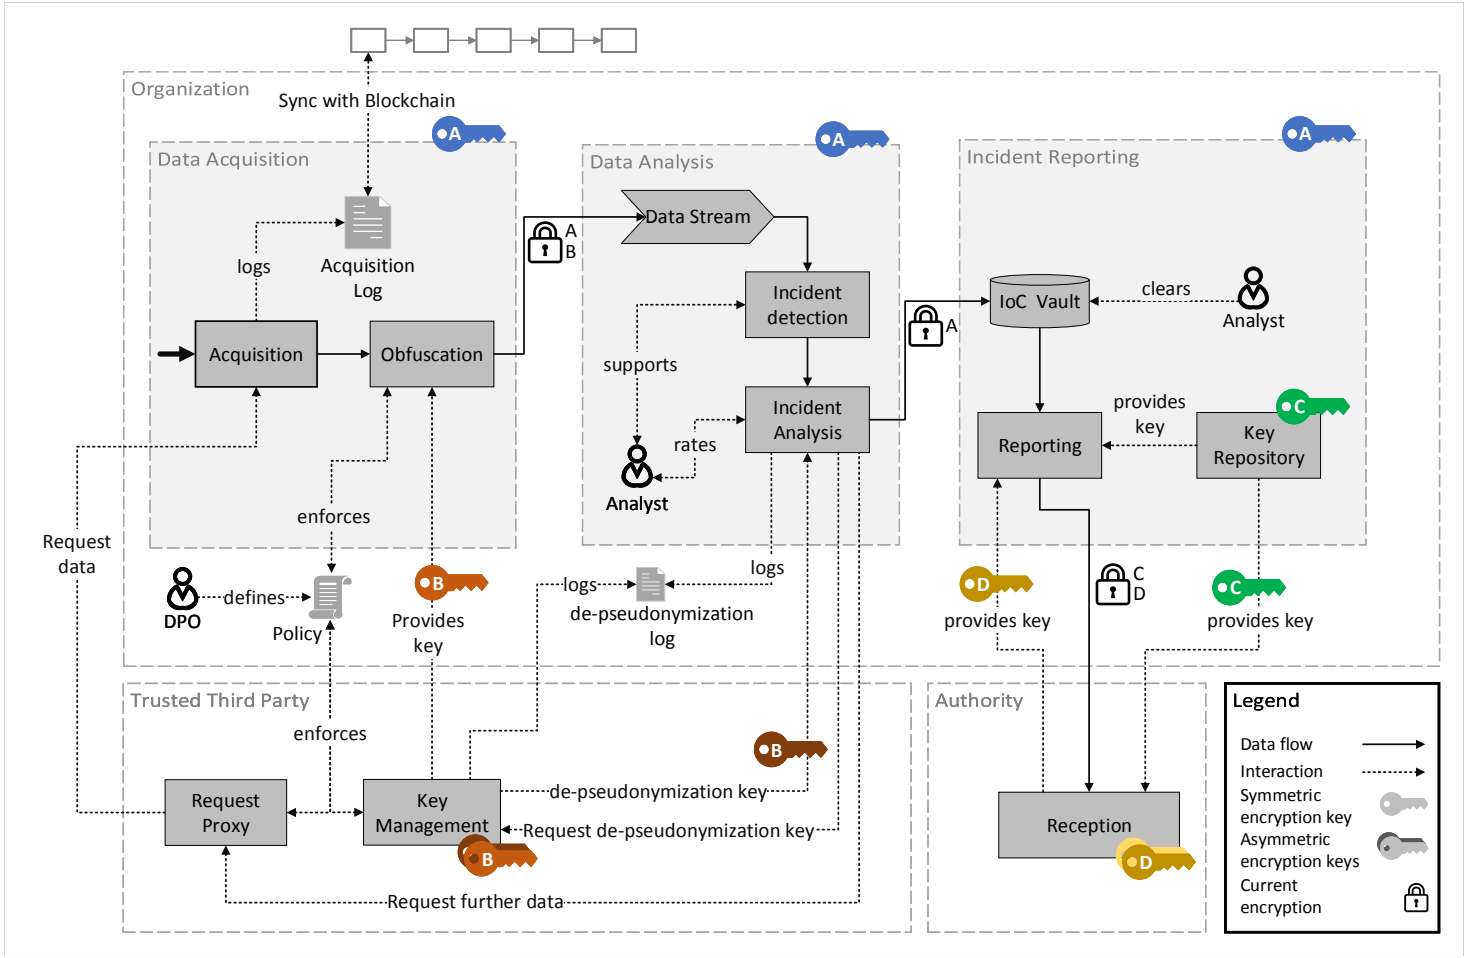

Fig. 2. DINGfest GDPR Architecture

The GDPR demands the protection of personal data as it is processed. Thus, we argue that personal data must be pseudonymized as soon as possible in the system. In the case of SIEM this is the case directly after or ideally during the data is acquired. To achieve this, a public key (B) is provided by a *TTP (Trusted Third Party)*, which is responsible for *key management*. With this key, the fields containing personal data are encrypted asymmetrically and the unencrypted personal data is deleted. In order to be able to comprehend and proof, that all personal data has been pseudonymized, an *Acquisition Log* is kept. For this purpose, we propose to use a tamper proof logging scheme, which synchronizes all logged data with an external blockchain as presented by Putz et al. [35]. For proper use of public key cryptography, we refer to López et al. [22].

In order to determine, which fields must be encrypted, a *Policy* is followed. For each logging system type an individual mapping must be defined that specifies the fields containing personal data. This policy is defined by the *Data Protection Officer (DPO)* or an equivalent position inside the organization. A DPO is responsible for compliance and data protection within an organization and should have the necessary expertise to make the required decisions.

In addition to the data protection aspects shown above, the model presented is also intended to provide protection against attacks resulting from the attacker model defined in Section 3. According to the assumptions made, possible dangers from insiders and outsiders are examined more detailed in the following.

Outsider: According to the model, outsiders can be divided into two main groups. Common outsiders, which have no specific reference to the system, and the TTP as an outsider, which is partially involved in the analysis process. The authority as the third participant is not considered in detail here, as it is supposed to have access to the data it receives in the context of a report.

- **Common outsider:** The major problem is to prevent data flows to outsiders. Specifically, the areas of data acquisition, data analysis and incident reporting must be protected. This is essentially guaranteed by a consistent use of the internal, symmetric key A. This ensures that the data remains protected even in the event of an unwanted extraction. The data may only be possibly unprotected in the case of an extraction during the data acquisition process. The security of this data mainly depends on the level of protection of the underlying source system.
- **TTP:** The TTP used in the present data model also represents a specific outsider, who is integrated into the SIEM process. However, the TTP only receives the key B from the data flow for custody, but does not have access to data from the data stream at any time. It is also worth noting that the TTP does not have access to the key A at any time. Accordingly, the TTP must be considered the equivalent of other outsiders in the case of data leaks.

During data analysis, an incident detection approach is followed. This approach is mainly automated but can also be supported by human analysts. Thereby, the incident detection is conducted solely on pseudonymized data and

thus is GDPR compliant. The thereby used event detection approach is elaborated further in the following chapter.

Insider: In the present model, only two groups of people have access to the internal data. These are analysts in the areas of data analysis and incident reporting on the one hand and data protection officers who define the corresponding policies on the other hand.

- **Analyst:** The analysts involved are only provided with specific data extracts and personal data under certain circumstances. For this purpose, an approval for specific data components must be granted according to the policy defined by the data protection officer. If such an approval does not exist, analysts always work only with pseudonymized data.
- **DPO:** The data protection officer is never given access to the data within the data stream and thus has access to resources that are equivalent to an outsider. On the other hand, the DPO has a protective influence on the data stream by defining the respective policy. This ensures that the protection of the data stream is always split between two different roles within the company.

4.3 Event Detection on protected Data

Different software usually comes with different log formats that is often loose text. In our case, we use standard Linux logs like *syslog* and *auth.log* that usually come with a Linux distribution. Furthermore, we use *access.log* of Apache's HTTP Server [42]. Since system call traces are a rich source of behavioral information [36,19], we also use system calls traces that are obtained via virtual machine introspection. System call tracing has a negative impact on performance, but especially enterprise environments can benefit since some events cannot be detected using common logs.

Table 1. The unified log message format [21]

Name	*Description*
source	The source from where the message comes from.
type_id	Describes the type of the message, e.g., the system call number.
date	Timestamp of the message when it was generated.
path	A path, e.g., which path was opened.
user	The user who performs the event.
process_name	The name of the process that performs the event.
...	...
misc	Can be used for random things (no personal data) that do not fit into that format.

Log messages are transformed into a unified structured log format. An excerpt of the message format that we use in DINGfest can be seen in Table 1 [21]: The entry *source* specifies from which of our sources log the message comes from. Thereby, we assume that it is not possible to deduce from the source to the user, i.e., in server scenarios. One of the most important attributes of the unified log message is *type_id*. This ID specifies what kind of message it is. In case of system calls the *type_id* is the system call number. *Misc* may contain arbitrary information that does not fit into the unified message format, e.g. command line option. We assume, that this field does not contain personal data. For the evaluation we checked manually that this field does not contain personal data. Another useful feature is *path*. However, the path may contain personal data such as the user name.

An example of a unified log message is shown in Listing 1.1

```
{
    "source"       : "syscalls",
    "type_id"      : 59,
    "process_name" : "ls",
    "user"         : "alice",
    "pid"          : 103,
    "path"         : "/home/alice/topsecret/"
    ...
}
```

Listing 1.1. Example of a unified message.

So we can distinguish between three kinds of log file entries:

1. Those that definitely contain personal data (e.g., user),
2. those that may contain personal data (e.g., path), and
3. those that do definitely not contain personal data (e.g., source, type_id, process_name, misc)

The classification may vary from system to system. For example, it is also possible that in a specific scenario a process name or a source name may also contain personal data. The classification, however, determines which features can be used for privacy friendly event detection, namely only features from the third category. We use this idea in our evaluation in Section 5.1 to assess the impact of privacy protection on event detection quality.

4.4 De-pseudonymization in case of an incident

In the previous section we presented an approach that allows us to perform incident detection on data that is pseudonymized according to the GDPR regulation. On this basis, this section describes how security incidents can be de-pseudonymized after detection in order to analyze and process them further and to prepare them for a legally compliant report. The complete process for these descriptions is additionally shown in Fig. 3.

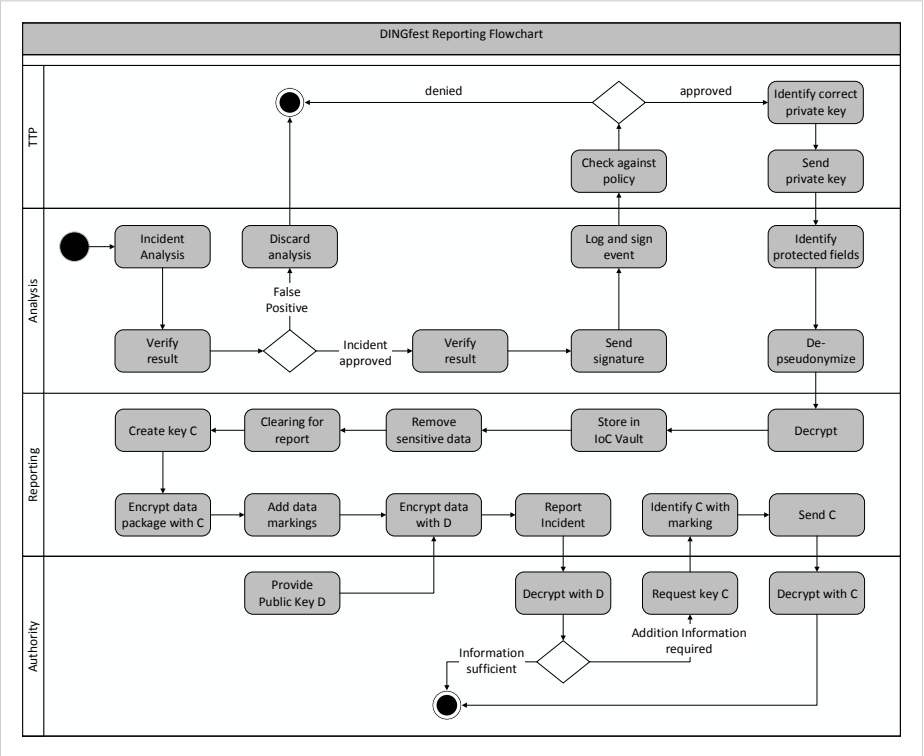

Fig. 3. DINGfest Reporting Flowchart

When the automated data analysis found indications for a possible incident, it is first necessary to verify the result. For this purpose, the data is revised by an analyst to ensure that it is an actual incident and to avoid false positives. Once the analyst approves the incident within the *Data Analysis* module, the *Trusted Third Party (TTP)* is contacted, which initially provided the public key B for the data pseudonymization within the *Data Acquisition* module. The *TTP* receives the signature of the data packet concerned in order to be able to identify the appropriate key and checks the request against the decryption policy specified by the DPO. If the check is negative, the request is denied. If the check is positive, the *TTP* determines the correct private key B for the signature provided and sends it to the *Data Analysis* module to allow the de-pseudonymization of the data. When de-pseudonymizing incident data, it is also important to enable auditing of de-pseudonymizations. Therefore, we store every de-pseudonymization request and sign it, in order to be able to provide proof of data access afterwards.

After the analysis module has received the private key B from the *TTP*, the initially appended data markings are used to identify all fields that contain pseundonymized information within the data package and to de-pseudonymize

it. The resulting data package is then transferred to the Incident Reporting module where it is decrypted using the key A and stored within the *IoC Vault*, which is an integrity proof, long term storage for incident data as shown by Boehm et al. [4]. This data can then be used for further analyses and incident reports. If an incident must be reported, the data needs to be cleared by an analyst first. This is important to prevent both privacy violations of personal data and the publication of confidential company data. More specifically, the analyst must decide, which data is to be excluded from the report and which is to be secured. This additional information about performed pseudonymizations and exclusions is appended to the data using further data marking definitions. The actual extent of data protection, however, may strongly depend on the use-case of the report. While in the case of reports within the scope of a reporting obligation, the statutory requirements must be complied with, in the case of voluntary reports significantly more data can be concealed or removed.

After an analyst has cleared an incident and prepared it for a report within the *IoC Vault*, it is transferred to the *Reporting* module in the next step. A symmetrical key C is created in the key repository and assigned to this very data package to pseudonymize the previously chosen contents. Furthermore, the package is extended by an additional data marking that contains the signature of the utilized key C key to enable later attributions. The key C is then used to pseudonymize the data according to the analyst's specifications. In addition to this, the data is also encrypted with the public key D, provided by the recipient (for example, an authority). This ensures that the reported data can only be opened by the correct recipient, i.e., a legal authority. In a final step, the data (secured with keys C and D) is transferred to the recipient. If the data was reported due to a reporting obligation, such as European NIST directive [11] or German IT-security law [9], the receiving authority may request a decryption afterwards under certain conditions. In this case, the authority needs to be able to request the key C from the organization. In order to receive the key C, the authority transmits the key signature contained in the data package to the organization. This enables the assignment of the correct key. If the correct key C is assigned, it can be transferred to the authority.

4.5 Evidence generation using malleable signatures which withstands pseudonymizations

We can positively answer the second half of Q3, i.e. we can ensure that we are able to use the pseudonymized or partial reported events as means of evidence. Assume we add protection of integrity and origin authentication during the data gathering inside the data acquisition module. Inside DINGfest's base architecture, the evidence could be protected by standard electronic signature schemes, e.g. the different acquisition modules would sign the data they gathered. When a cryptographic signature algorithm complies with the requirements of common legal frameworks for electronic signatures its signature provides a high probative value for the data being signed. This said, any subsequent modification for data-protection compliance would destroy any evidence guarantees for the

remaining data, e.g. removing the full path from a signed full file name of a malicious executable as it contains a personally-identifiable user name also removes the evidence protection for the name of the executable. Thus, we propose use redactable or sanitizable (malleable) signature schemes to retain the authenticity and integrity of the data gathered from the data acquisition module towards the final report. This means that, if wanted, the digital signature protects the authenticity of the data provided even till the incident report, i.e. so in the final report one can verify that the event data has not been modified in unauthorized ways –not tampered with– and that it originated from a trusted data acquisition module.

By omitting the cryptographic details of other malleable signature schemes, the privacy statement describes, which parts of the data could be removed or pseudonymized. This allows to control the signature scheme for which subsequent changes are made and parts are authorized. Thus, removing the original data or encrypting these parts, would allow the subsequent steps to always verify the authenticity of the remaining data. When data is de-pseudonymized the best-suitable malleable signature schemes are those that offer mergeability [33]. This allows to put data parts back into the signed data set and thus the original malleable signature would now verify over all remaining parts, the ones previously readable plus those added by the de-pseudonymization. If not added, private malleable signature schemes [2] retain the confidentiality of all those parts that have been removed, i.e. even though one can successfully verify the signature on partial data, the information contained in the signature itself does not allow an attacker to gain information on the data parts removed and thus not shared. Hence, the added value of a retain-able private malleable signature, like [33], does not violate any GDPR requirements [30].

The legal analysis of these private accountable redactable signature scheme shows that they increase the legal probative value for the signed reported data as eIDAS compliant electronic signatures could provide [30,16,31,45]. Hence, the DINGfest GDPR architecture protects the records such that the remaining information can be used as means of evidence; further after the de-pseudonomization steps at any later time the data's origin and originality is attested.

5 Evaluation

In the previous sections we presented an architecture for a GDPR compliant SIEM system. The central elements of our approach are to guarantee a GDPR compliant data processing, to enable the recognition of security incidents on pseudonymized data as well as the de-pseudonymization of the data in the case of an incident. In this section we evaluate the validity of our approach in two ways. First, we conduct a technical evaluation of the impact of pseudonymization on the detectability of events. Subsequently, we carry out a legal evaluation for our proposed solution. To achieve this, we investigate the individual components of our architecture presented in Section 4.2 on conformity with the specifications of the GDPR as shown in Section 2.2.

5.1 Impact on Detectability

Evaluation Methodology The evaluation of the impact of privacy protection on the quality of SIEM is performed based on the theory for forensic fingerprint calculation of Dewald [10].

In this theory, all interactions of interest with the system (e.g., by users) are called *events*. An example event is the login of a user. Many events either directly or indirectly leave digital traces within the system (e.g., in log files on disc or in main memory). These traces are formalized as *feature vectors*. Generally, a feature is a quantifiable attribute of a system that can be observed by the SIEM. In our study we concentrate on feature vectors that can be extracted from log files system call traces [21]. Obviously, tracing all system calls is very expensive in terms of performance, there are ways to get rid of most overhead caused by system call tracing. The theory [10], which we now explain, defines conditions under which an event is detectable based on the features traces it leaves in the log files of a system.

The set of features that we consider in our evaluation is based on an abstract representation of log file entries and attempts to harmonize many log files in modern systems. Our format represents every log file entry using the following four *features*:

- a *source* from what log the message comes from,
- a generic *type_id* that describes the kind of log message,
- a *path*, and
- a *misc* field that may contain arbitrary content (e.g., the name of a network adapter).

A *feature vector* is a vector of values for these features. Depending on the system, there can be many different features vectors consisting of these four features. Since an event can cause multiple entries in multiple log files, we define the set of feature vectors that are generated as the *evidence set* of that event.

More formally, let Σ be the set of all possible events that can happen in the system and are of interest to the SIEM. When some event $\sigma \in \Sigma$ happens, log entries are generated. The *evidence set* $E(\sigma)$ of event σ is the set of all subsets of feature vectors that are thereby generated by σ. It is technically necessary, that the evidence set is closed under subsets. Intuitively, it can be interpreted as the fact that partial evidence is also evidence of the event.

It is obvious that the evidence sets of different events may overlap. To be able to detect an event, it is necessary to calculate the *characteristic evidence set* $CE(\sigma)$[10] of an event σ, which is defined as the set that contains only feature vectors that are caused by σ and *not* by any other event $\sigma' \in \Sigma$. Formally, the set of characteristic evidence of an event σ with respect to a set of other events Σ' is defined as follows:

$$CE(\sigma, \Sigma') = E(\sigma) \setminus \bigcup_{\sigma' \in \Sigma'} E(\sigma')$$

The set of characteristic evidence of an event is also called *characteristic fingerprint* of that event.

As one can see in the formula above, a characteristic fingerprint is defined for a specific reference set Σ'. All feature vectors that are caused by $\sigma' \in \Sigma'$ are not in $CE(\sigma, \Sigma')$. One can say, that $CE(\sigma, \Sigma')$ is the evidence set $E(\sigma)$ minus all other evidence sets of events in Σ'. Let $|CE(\sigma, \Sigma')|$ and $|\Sigma'|$ be sufficiently large, then a match of the feature vector with the log files of a system is a clear indication that σ happened and not σ'. The size of CE is an indication for the discriminative power of the evidence. The larger the set, the higher is the probability that the event may be detected no matter what the reference set Σ' looks like. However, it is also possible, that $CE(\sigma, \Sigma')$ is empty, i.e., one cannot detect reliably the occurrence of σ.

Characteristic Evidence without Personal Data We now evaluate the impact of pseudonymization on the existence of characteristic evidence that is needed for event detection. The evaluation setting is based on the DINGfest architecture as described by Latzo and Freiling [21] [20]. We calculated evidence and characteristic evidence sets for 45 different events (see also Table 2 that typically appear in Linux server environments as one typically finds them in small and medium-sized enterprises. In our threat model, we consider an adversary with root privileges that were either gained via a privilege escalation attack or by having them anyway (i.e., a malicious insider). Basically, it is not possible to determine the intention of an administrator's input. Hence, most events can also be used maliciously, e.g., for information retrieval, covering traces, etc. So, basically all events might be interesting during a forensic analysis or event detection.

The higher the number of feature vectors in a characteristic fingerprint, the better the quality of that fingerprint. This is intuitive since a feature vector in a fingerprint is basically an indicator of an event. In the evaluation, we compare the size of characteristic evidence sets with and without taking personal data into account. More concretely, we consider the following two feature sets:

- $F_1 = \{source,\ type\_id,\ path,\ misc\}$
- $F_2 = \{source,\ type\_id,\ misc\}$

First, F_1 is a feature set that has turned out to be reasonable for our events. However, F_1 includes the *path* feature that may contain personal data. Features of F_2 do not contain personal data. For calculating the fingerprints, an event was executed 40 times (trainings set). Furthermore, the reference set Σ' for a characteristic fingerprint of σ are always all other events. Table 2 compares sizes of the characteristic evidence set using the two feature sets including the decrease rate when using F_2 instead of F_1. As one can see, omitting the path as a feature has a huge impact on the size of characteristic evidence. On the average, using F_2 reduces a characteristic fingerprint by half of its feature vectors. For three events, there is no characteristic fingerprint, anymore. Figure 4 shows the absolute number of feature vectores of the characteristic evidence sets using F_1 and F_2. "Big"events, that come originally with big characteristic fingerprints are in general more affected than smaller events.

Table 2. The events used for the evaluation

Class	*Name*	*Description*	$\lvert CE(\sigma, \Sigma') \rvert$ F_1	F_2	*Loss Factor*
CLI	ls	Lists files	1	1	0.0
	cp	Copies file	4	1	0.75
	mv	Moves file	2	1	0.5
	cat	Cats file	0	0	0
	vmstat	Virtual memory statistics	6	1	0.833
	netstat	Network statistics	15	1	0.933
	tar	Creates compressed tar archive	5	4	0.2
	rm	Removes file	1	1	0.0
	shred	Shreds file	2	1	0.5
	curl	Downloads file	1	0	1
CLI Root	tailShadow	Reads /etc/shadow	7	2	0.714
	catCredentials	Reads Wordpress config file	4	2	0.5
	vimHosts	Opens /etc/hosts in Vim	220	3	0.986
	rmSudo	Removes file with sudo	2	2	0.0
	shredSudo	Shreds file with sudo	9	3	0.667
Web	wordpressLogin	Wordpress Login	63	10	0.841
	wordpressSearch	Wordpress Search	3	0	1
	wordpressOpen	Opens Wordpress website	0	0	0
Service	sshLogin	SSH login (server side)	2219	466	0.79
	apacheStop	Stops apache web server	1712	15	0.991
	mysqlWp	Login into Wordpress DB via command line	47	1	0.979
Kernel Modules	lsmod	Lists loaded kernel modules	251	1	0.996
	insmod	Loads kernel module	10	3	0.7
	rmmod	Unloads kernel module	12	3	0.75
Docker	dockerHelloWorld	Starts docker hello world example	28	3	0.893
	dockerUbuntuLog	Starts docker ubuntu and show log	23	5	0.783
	dockerImages	Lists all docker images	1	1	0.0
	dockerPs	Lists all running dockers	0	0	0
	dockerPSA	Lists all dockers container	0	0	0
	dockerUbuntuSleep	Starts docker in background	2	2	0.0
	dockerRm	Removes all docker containers	0	0	0
	dockerNginx	Runs nginx docker and curl it	65	8	0.877
	dockerUbuntuBash	Attaches bash of container	0	0	0
	dockerPrune	Removes unused container	1	1	0.0
	dockerPruneVolumes	Removes unused objects and volumes	1	1	0.0
	dockerRmImages	Removes all images	2	2	0.0
	dockerUbuntuBashCp	Attaches container and runs cp	0	0	0
	dockerUbuntuBashMv	Attaches container and runs mv	18	1	0.944
	dockerUbuntuBashRm	Attaches container and runs rm	3	1	0.667
	dockerUbuntuBashCat	Attaches container and runs cat	24	0	1
Nextcloud	nextcloudStatus	Shows Nextcloud status	3	2	0.333
	nextcloudAppList	Lists Nextcloud apps	44	2	0.955
	nextcloudUserList	Lists Nextcloud user	3	2	0.333
	nextcloudUserAdd	Adds new Nextcloud user	103	16	0.845
	nextcloudGroupList	List Nextcloud groups	5	2	0.6
Average					0.508

We have shown that it is possible to calculate characteristic evidence for all events even if features that contain personal data are not used. However, the fingerprints that we generated had a lower quality, i.e., the size of the characteristic evidence set was reduced by an average of about 50%. By extending the feature set and the set of traces acquired from the SIEM, we conjecture that fingerprints can also be calculated for this action even if data is pseudonymized.

In the following we want to compare the matching results using characteristic fingerprints with F_1 and F_2. For matching, we calculate a score that indicates

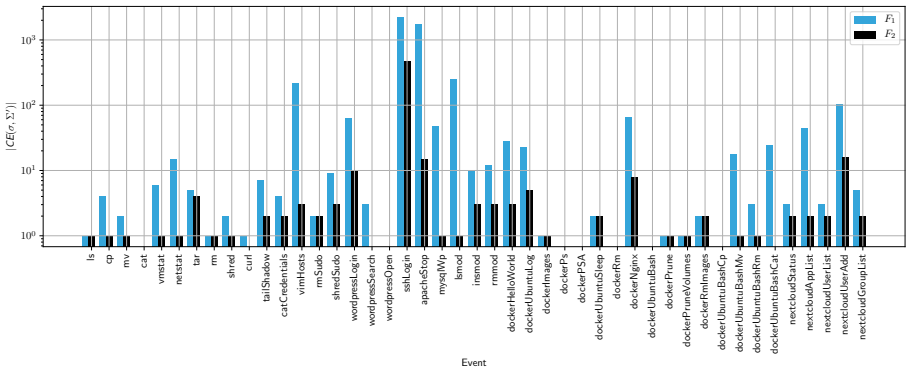

Fig. 4. Comparison of powers of characteristic evidence sets with personal data (blue) and without personal data (black).

what proportion of feature vectors in event traces are matched by a characteristic fingerprint. Fig. 5 (F_1) and Fig. 6 (F_2) show the corresponding matching matrices. The values there are average values of 10 traces of the event (test set). It stands out that the matching matrices are quite similar. In Fig. 6 there is much less matched noise. The characteristic fingerprints in Fig 6 are much smaller than in Fig. 5, though. For that events for which we could calculate a characteristic fingerprint, the matching results are similar good. This is also confirmed by the *Receiver Operating Characteristic* (ROC) curve in Fig. 7. There, the true positive rate (sensitivity) and the false positive rate are plotted against each other with different thresholds. It shows, that the sensitivity with F_2 is only a little smaller with about 78% versus the sensitivity of F_1 with about 84%.

In this section we showed that it is possible to calculate characteristic fingerprints with a reduced feature set that does not contain personal data. While the size of the characteristic fingerprints decreased when using that feature set, the matching results were very similar. It showed, that when there is a characteristic fingerprint, matching usually also works. So, to improve this approach, future work should focus in extracting more features from logs that help to increase the size of characteristic fingerprints.

5.2 Legal Evaluation

As we have seen above, to generate fingerprints of high quality, data must be obtained by processing previously collected data during the data acquisition which clearly relates to a personal user's actions and thus is considered as personal data. Hence, the general principles mentioned in Sec. 2.2 relating to processing of personal data and especially a lawful processing are important. Processing shall be lawful only if one of the Art. 6 GDPR included reasons applies. Applicable and best suited for the SIEM case is Art. 6 (1) lit. f, when processing is necessary for the purposes of the legitimate interests. This clause is different to

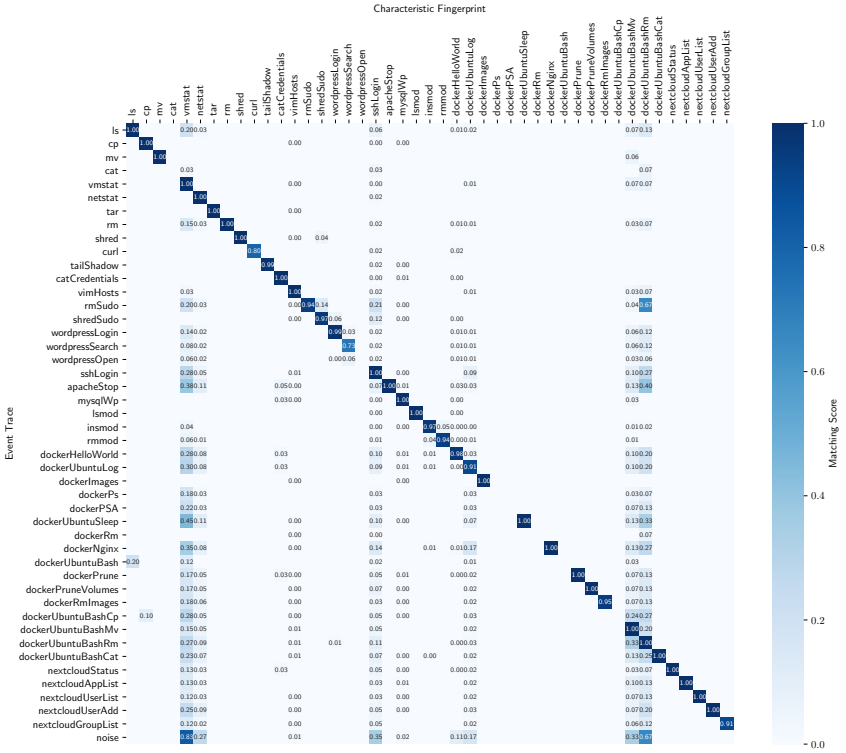

Fig. 5. Matching matrix using F_1. The events listed on the y-axis are the ground truth, the events on the x-axis correspond to the characteristic fingerprints [20]

the other lawful bases as it is not centered around a particular purpose and it is not necessary that the individual has specifically agreed to (consent). Legitimate interests are more flexible and could in principle apply to any type of processing for any reasonable purpose. Art. 6 (1) lit. f states:

"1.Processing shall be lawful only if and to the extent that at least one of the following applies: (f) processing is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, [...]".

Legitimate interest is balanced with personal data protection Since legitimate interests can apply in a wide range of circumstances, it is mandatory that the controlling party puts its legitimate interests and the necessity of processing the personal data to the interests, rights and freedoms of the individual in balance. To provide a balance-test, the key elements of the legitimate interests provision is contained in a so-called three-part test. Whereas this test is not ex-

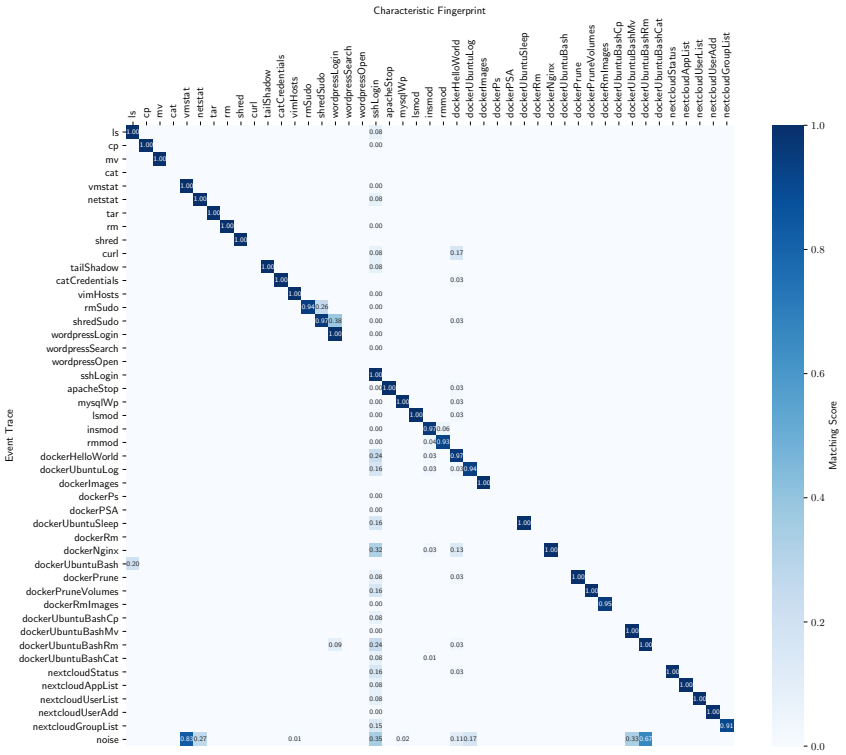

Fig. 6. Matching matrix using F_2. The events listed on the y-axis are the ground truth, the events on the x-axis correspond to the characteristic fingerprints.

plicitly named in the GDPR, the legitimate interests provision does incorporate three key elements:

- Purpose test: there must be a legitimate interest behind the processing.
- Necessity test: the processing must be necessary for that purpose.
- Balancing test: the legitimate interest must be balanced with the individuals interests, rights or freedoms.

This concept of a three-part test for legitimate interests has been confirmed by the Court of Justice of the European Union in the Rigas case (C-13/16, 4 May 2017) in the context of the Data Protection Directive 95/46/EC, which contained a very similar provision. This means, the controller must be able to meet all three requirements of the test prior to commencing the processing of personal data.

Firstly, **purpose** is clearly given as the whole purpose of the SIEM architecture as given in Sec. 4 is to detect unlawful use of information systems and their data, it is important to make clear that the European Parliament has already considered the legitimate interest of processing personal data necessary for the

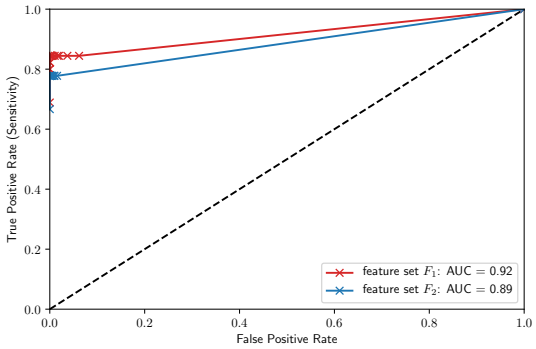

Fig. 7. ROC curve with F_1 in red and F_2 in blue. The differences are quite small.

purposes of preventing fraud. This is explicitly backed by recital 47: "*[...] The processing of personal data strictly necessary for the purposes of preventing fraud also constitutes a legitimate interest of the data controller concerned. [...]*".

Secondly, with regards to the condition relating to the **necessity** of processing personal data, it is important that derogations and limitations in relation to the protection of personal data must apply only in so far as is strictly necessary. In that regard, communication of features which do not contain personal data, does not make it possible to identify a person with enough precision in order to be able to bring an action against him. Accordingly, for that purpose, it is necessary for the SIEM system to obtain also the possibility of full identification of that person, i.e. allow to de-anonymize and retain authenticity proofs in order to construct substantial and reliable evidence of an unlawful use of the system against that person.

Thirdly, it is necessary to make a **balancing test** to justify any impact on individuals. During the test the controller takes into account "the interests or fundamental rights and freedoms of the data subject which require the protection of personal data", and makes sure they dont override his interests. In recital 75 speaks of the risks of the rights and freedoms of natural personas: "*The risk to the rights and freedoms of natural persons, of varying likelihood and severity, may result from personal data processing which could lead to physical, material or non-material damage, in particular: where the processing may give rise to discrimination, identity theft or fraud, financial loss, damage to the reputation, loss of confidentiality of personal data protected by professional secrecy, unauthorized reversal of pseudonymisation, or any other significant economic or social disadvantage; where data subjects might be deprived of their rights and freedoms or prevented from exercising control over their personal data; [...]; where personal aspects are evaluated, in particular analyzing or predicting aspects concerning performance at work, [...] ; or where processing involves a large amount of personal data and affects a large number of data subjects.* "

Since the data acquisition module collects data from all monitored computing resources in the company, one can assume a great danger for personal data of employees and customers. Also, the analysis of personal data in the data stream by fingerprinting and pattern recognition and especially the merging of data is in general - interfering with the privacy rights of a natural person. And finally, is the reporting module and the included long-term storage of analyzed incidents as well as the reporting to the authorities itself a potential risk for personal data. Since the complete monitoring of the users without cause is not compliant with the GDPR, especially since the user does not have any possibility to intervene, fundamental rights would be violated, if the controller is not implementing appropriate technical and organizational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate, for example the pseudonymisation and encryption of personal data.

Mechanisms for GDPR-compliance in DINGfest architecture DINGfest's GDPR architecture counters the above-mentioned problems by implementing special steps as part of their work flow for a GDPR-compliant SIEM system:

First of all, by a continuous pseudonymisation through obfuscation during the data acquisition. The suspension will only be carried out under certain conditions determined by controller and, in particular, in case of suspicion of a criminal offence. Since the public key is always provided by a trusted third party (TTP) and policies provide the organizational background before a special field gets encrypted, the balance between the rights of the controller and the user should be met. All technical steps in which personal data is processed are accompanied by a special pseudonymization method through obfuscation. If data analysis has then found indications for an (possible) incident, the data protection officer has to approve this case as an "incident case" within the data analysis module. Only then the TTP receives a key identifier for the data packet - not the packet's contents, not even in pseudonymous form -, in order to find the appropriate key. The critical point is the policy (see the "Policy" defined by the Data protection office (DPO) in Fig. 2), which is being consulted by the TTP before sending the decryption key to the data analysis module. This organizational measure ensures a level of security appropriate to the risk, which is to reveal private data to the controller. By providing a log of every request to de-pseudonymize data fields the architecture enables to comply with transparency requirements, like the right of access.

Only after passing this safety measure the DINGfest architecture allows to reveal data to the controller (the data analyst) using the private key B, to de-pseudonymize all necessary fields that contain pseudonymized information within the data package. Only if the analyst decides to include this in the report the resulting data package is then transferred to the incident reporting module, during transfer and storage it gets again encrypted under key A. Again, the access to the encrypted long-term storage of the data within the IoC vault, including the use of data for further analysis and incident reports, only applies for cases that deserve an attention because of potential unlawful behavior. This is

clearly a legitimate interest which is not overridden by recital 47 of the GDPR. Furthermore, it depends on the individual use case, which data is to be excluded from the report and what data must be removed or stays pseudonymized. This extra step, in which the analyst balances the rights of both parties, is the very essence of the balance test, and here the DINGfest GDPR architecture implements this check to balance the legitimate interest with the individuals interests. As said, if the data analysis module decides to keep pseudonymization of fields then only pseudonymized data of this incident is transferred to the incident reporting module. When the report is generated, the data is again encrypted with the public key of the recipient in most cases an authority. This way, the reported data can only be opened by the correct recipient.

DINGfest's GDPR architecture reduces the risks for personal data By using the methods of pseudonymisation and encryption of personal data, it is to be concluded that scope, context and purposes of processing as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller has implemented appropriate technical and organizational measures to ensure a level of security appropriate to the risk.

6 Conclusion

In this paper we presented an architecture for a GDPR compliant SIEM system, as implemented in the DINGfest prototype SIEM system. We first identified central questions that must be answered for the development of such a platform. The questions affected the necessary conditions for a GDPR compliant data processing, techniques for the incident recognition on pseudonymized data as well as a lawful de-pseudonymization techniques in the case of occurred incidents. We then answered these questions with the help of our architectural design and evaluated them both from a technical and legal perspective. Using this evaluation, we have shown that it is possible to comply with the legal requirements for pseudonymization, while at the same time keeping detectability. Altogether we presented a base architecture for a GDPR compliant SIEM system with this work. Although it was developed based on our underlying system DINGfest, it may also serve as a draft for other security systems that have to be adapted to GDPR specifications.

In the context of this work, it was revealed that the performance of the recognition mechanisms used can be impaired using pseudonymization. This is one challenge that could be addressed in future work. Beyond that we defined the fundamental boundaries of a GDPR conform architecture with this work. However, various details were not considered. An example for this is to transfer our architecture to already established SIEM systems. Each system is tailored to its infrastructure and thus, it is necessary to define, which of the collected data sets needs to be protected. This applies both to data that is collected during initial data acquisition and to data that is prepared for a report. To support this process, it would be helpful to develop a central repository that defines the data

points relevant to data protection for frequently used data sources. Furthermore, it will also be necessary to develop the needed details for the data protection policy within SIEM systems in future works. It would be conceivable to develop a generally applicable basic policy and specific implementations of this policy adapted to individual systems.

Regarding the legal probative value DINGfest using malleable signatures allows to balance integrity protection for evidence and GDPR-compliant removal or pseudonymization of the gathered data. To achieve this the data acquisition module emits malleable signed data –instead of simply digitally signed data– and hence any subsequent modification due to GDPR-compliant processing does not inhibit the verification of the integrity and origin of the remaining data. With a scheme that is accountable and private and supports mergeability, previously obfuscated parts of an entry can be subsequently de-obfuscated and the signature still verifies and provide means of evidence.

Acknowledgement

This research was supported by the Federal Ministry of Education and Research, Germany, as part of the BMBF DINGfest project (https://dingfest.ur.de/). The research of H.C.Pöhls was carried out in the project SEMIOTICS funded by EUs H2020 grant no. 780315.

References

1. Ateniese, G., Chou, D.H., de Medeiros, B., Tsudik, G.: Sanitizable Signatures. In: Proc. of European Symposium on Research in Computer Security (ESORICS 2005). LNCS, vol. 3679, pp. 159–177. Springer (2005)
2. Bilzhause, A., Pöhls, H.C., Samelin, K.: Position Paper: The Past, Present, and Future of Sanitizable and Redactable Signatures. In: Proc. of International Conference on Availability, Reliability and Security (ARES 2017). pp. 87:1–87:9. ACM (Sept 2017)
3. Biskup, J., Flegel, U.: Transaction-based pseudonyms in audit data for privacy respecting intrusion detection. In: International Workshop on Recent Advances in Intrusion Detection. pp. 28–48. Springer (2000)
4. Böhm, F., Menges, F., Pernul, G.: Graph-based visual analytics for cyber threat intelligence. Cybersecurity **1**(1), 16 (Dec 2018)
5. Brzuska, C., Pöhls, H.C., Samelin, K.: Non-Interactive Public Accountability for Sanitizable Signatures. In: Revised Selected Papers of European PKI Workshop: Research and Applications (EuroPKI 2012). LNCS, vol. 7868, pp. 178–193. Springer (2012)
6. Burkhart, M., Strasser, M., Many, D., Dimitropoulos, X.: Sepia: Privacy-preserving aggregation of multi-domain network events and statistics. Network **1**(101101) (2010)
7. Büschkes, R., Kesdogan, D.: Privacy enhanced intrusion detection. Multilateral security in communications, information security pp. 187–204 (1999)

8. Coppolino, L., D'Antonio, S., Mazzeo, G., Romano, L., Sgaglione, L.: How to Protect Public Administration from Cybersecurity Threats: The COMPACT Project. In: 2018 32nd International Conference on Advanced Information Networking and Applications Workshops (WAINA). pp. 573–578 (May 2018). https://doi.org/10.1109/WAINA.2018.00147
9. Deutscher Bundestag: Gesetz zur Erhöhung der Sicherheit informationstechnischer Systeme (2015), https://www.bmi.bund.de/SharedDocs/Downloads/DE/Gesetzestexte/it-sicherheitsgesetz.pdf
10. Dewald, A.: Characteristic evidence, counter evidence and reconstruction problems in forensic computing. it - Information Technology **57**(6), 339–346 (2015)
11. European Commission: NIS Directive 2016/1148 (EU) of the European Parliament and of the Council (2016), http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016L1148
12. European Parliament and the Council of the European Union: Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC. Official Journal **OJ L 257 of 28.8.2014**, 73–114 (Jul 2014)
13. European Parliament and the Council of the European Union: 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). Official Journal **OJ L 119 of 4.5.2016**, 1–88 (May 2016)
14. Gartner Inc.: Security information and event management (siem) (2018), https://www.gartner.com/it-glossary/security-information-and-event-management-siem
15. Goldstein, M., Asanger, S., Reif, M., Hutchison, A.: Enhancing security event management systems with unsupervised anomaly detection. In: ICPRAM. pp. 530–538 (2013)
16. Höhne, F., Pöhls, H.C., Samelin, K.: Rechtsfolgen editierbarer Signaturen. Datenschutz und Datensicherheit - DuD **36**(7), 485–491 (Jun 2012), http://dx.doi.org/10.1007/s11623-012-0165-8
17. Jensen, M.: Challenges of privacy protection in big data analytics. In: 2013 IEEE International Congress on Big Data. pp. 235–238. IEEE (2013)
18. Johnson, R., Molnar, D., Song, D., Wagner, D.: Homomorphic signature schemes. In: Proc. of the RSA Security Conference - Cryptographers Track. pp. 244–262. Springer (Feb 2002)
19. Lanzi, A., Balzarotti, D., Kruegel, C., Christodorescu, M., Kirda, E.: Accessminer: using system-centric models for malware protection. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS 2010, Chicago, Illinois, USA, October 4-8, 2010. pp. 399–412. ACM (2010)
20. Latzo, T.: Efficient fingerprint matching for forensic event reconstruction (2020), under submission
21. Latzo, T., Freiling, F.: Characterizing the limitations of forensic event reconstruction based on log files. In: 2019 IEEE Trustcom/BigDataSE. IEEE (2019)
22. López, J., Oppliger, R., Pernul, G.: Why have public key infrastructures failed so far? Internet Research **15**(5), 544–556 (2005)
23. Menges, F., Böhm, F., Vielberth, M., Puchta, A., Taubmann, B., Rakotondravony, N., Latzo, T.: Introducing dingfest: An architecture for next generation siem sys-

tems. In: Langweg, H., Meier, M., Witt, B.C., Reinhardt, D. (eds.) SICHERHEIT 2018. pp. 257–260. Gesellschaft für Informatik e.V, Bonn (2018)
24. Miller, D., Harris, S., Harper, A., VanDyke, S., Blask, C.: Security information and event management (SIEM) implementation. Network pro library, McGraw-Hill, New York, NY (2011)
25. Miloslavskaya, N., Tolstoy, A.: New siem system for the internet of things. In: World Conference on Information Systems and Technologies. pp. 317–327. Springer (2019)
26. Mokalled, H., Catelli, R., Casola, V., Debertol, D., Meda, E., Zunino, R.: The applicability of a siem solution: Requirements and evaluation. In: 2019 IEEE 28th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE). pp. 132–137. IEEE (2019)
27. Nespoli, P., Gómez Mármol, F.: e-health wireless ids with siem integration. In: Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC18), Barcelona, Spain. pp. 15–18 (2018)
28. Park, H.A., Lee, D.H., Lim, J., Cho, S.H.: Ppids: privacy preserving intrusion detection system. In: Pacific-Asia Workshop on Intelligence and Security Informatics. pp. 269–274. Springer (2007)
29. Parliament of the United Kingdom: Freedom of Information Act 2000. URL http://www.legislation.gov.uk/ukpga/2000/36/pdfs/ukpga_20000036_en.pdf (November 2000)
30. Pöhls, H.C.: Increasing the Legal Probative Value of Cryptographically Private Malleable Signatures. Ph.D. thesis, University of Passau (2018)
31. Pöhls, H.C., Höhne, F.: The Role of Data Integrity in EU Digital Signature Legislation - Achieving Statutory Trust for Sanitizable Signature Schemes. In: Meadows, C., Fernandez-Gago, C. (eds.) 7th International Workshop, STM 2011, Copenhagen, Denmark, June 27-28, 2011, Revised Selected Papers. Lecture Notes in Computer Science (LNCS), vol. 7170, pp. 175–192. Springer Berlin Heidelberg (2011), http://dx.doi.org/10.1007/978-3-642-29963-6_13
32. Pöhls, H.C., Höhne, F.: The Role of Data Integrity in EU Digital Signature Legislation - Achieving Statutory Trust for Sanitizable Signature Schemes. In: Revised Selected Papers from the 7th International Workshop on Security and Trust Management (STM 2011). LNCS, vol. 7170, pp. 175–192. Springer (2011), http://dx.doi.org/10.1007/978-3-642-29963-6_13
33. Pöhls, H.C., Samelin, K.: Accountable Redactable Signatures. In: Proc. of International Conference on Availability, Reliability and Security (ARES 2015). pp. 60 – 69. IEEE (Aug 2015)
34. Pöhls, H.C., Samelin, K., Posegga, J., de Meer, H.: Transparent Mergeable Redactable Signatures with Signer Commitment and Applications (MIP-1206). Tech. Rep. MIP-1206, Faculty of Computer Science and Mathematics (FIM), University of Passau (Aug 2012)
35. Putz, B., Menges, F., Pernul, G.: A secure and auditable logging infrastructure based on a permissioned blockchain. Computers and Security p. 101602 (2019)
36. Rieck, K., Holz, T., Willems, C., Düssel, P., Laskov, P.: Learning and classification of malware behavior. In: International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment. pp. 108–125. Springer (2008)
37. Schlette, D., Böhm, F., Caselli, M., Pernul, G.: Measuring and visualizing cyber threat intelligence quality. International Journal of Information Security (2020). https://doi.org/10.1007/s10207-020-00490-y, https://doi.org/10.1007/s10207-020-00490-y
38. Sgaglione, L., Mazzeo, G.: A GDPR-Compliant Approach to Real-Time Processing of Sensitive Data. In: Intelligent Interactive Multimedia Systems and Services. pp. 43–52. Springer International Publishing, Cham (2019)

39. Sobirey, M., Fischer-Hübner, S., Rannenberg, K.: Pseudonymous audit for privacy enhanced intrusion detection. In: Information Security in Research and Business, pp. 151–163. Springer (1997)
40. Stahlberg, P., Miklau, G., Levine, B.N.: Threats to privacy in the forensic analysis of database systems. In: Proceedings of the 2007 ACM SIGMOD international conference on Management of data. pp. 91–102. ACM (2007)
41. Steinfeld, R., Bull, L., Zheng, Y.: Content extraction signatures. In: Proc. of International Conference on Information Security and Cryptology (ICISC 2001). vol. 2288, pp. 163–205. Springer (2002)
42. The Apache Software Foundation: Apache http server project (2019), https://httpd.apache.org/
43. The National Archives: Redaction toolkit – editing exempt information from paper and electronic documents prior to release. URL http://www.nationalarchives.gov.uk/documents/information-management/redaction_toolkit.pdf [last accessed: Nov. 2019] (Jul 2011)
44. United Kingdom Ministry of Justice: Lord Chancellor's Code of Practice on the management of records issued under section 46 of the Freedom of Information Act 2000. URL http://www.nationalarchives.gov.uk/documents/foi-section-46-code-of-practice.pdf [last accessed: Sep. 2019] (Jul 2009)
45. van Geelkerken, F.W.J., Pöhls, H.C., Fischer-Hübner, S.: The legal status of malleable- and functional signatures in light of Regulation (EU) No 910/2014. In: Proc. of the 3rd International Academic Conference of Young Scientists on Law & Psychology 2015 (LPS 2015). pp. 404–410. L'viv Polytechnic Publishing House (Nov 2015), https://drive.google.com/file/d/0B-Yu3Ni9z3PXM2lBajhCXzhoWk0/view
46. Vielberth, M., Menges, F., Pernul, G.: Human-as-a-security-sensor for harvesting threat intelligence. Cybersecurity **2**(23) (2019)
47. Vielberth, M., Pernul, G.: A security information and event management pattern. In: 12th Latin American Conference on Pattern Languages of Programs (SugarLoafPLoP 2018) (2018)
48. Wang, R.Y., Strong, D.M.: Beyond accuracy : What data quality means to data consumers. Journal of Management Information Systems **12**(4), 5–34 (1996), http://w3.cyu.edu.tw/ccwei/PAPER/ERP/dataquality%28JMIS%29.pdf
49. Williams, A.T., Nicolett, M.: Improve it security with vulnerability management. Technical Report - Gartner Inc. (2005)