Collaborative Fraud Detection in Outsourcing Scenarios: Issues of and Solutions for Privacy and Confidentiality

Ulrich Flegel[1] and Florian Kerschbaum[1] and Richard Wacker[2]

[1] SAP Research Center Karlsruhe, Germany
{ulrich.flegel|florian.kerschbaum}{@}sap.com
[2] Center for Applied Jurisprudence, University of Karlsruhe, Germany
wacker{@}ira.uka.de

Abstract. In this paper we investigate the privacy dimension of collaborative fraud detection envisioned for outsourcing scenarios. Firstly, we investigate the privacy requirements derived from privacy law and present the resulting judicial argument for pseudonymizing audit data generated for the purpose of fraud detection. Second, we summarize the requirements for such pseudonymization derived from the requirements of the misuse detection approach for fraud detection. Third, we describe our approach for pseudonymization of audit data and two approaches for hiding timestamps in audit data.

1 Introduction

Fraud often spans different organizations, and in the face of ubiquitous outsourcing, new opportunities for fraud will be created. Detecting fraud requires a complete picture of the executed business processes, obviously necessitating collaboration of the involved organizations for detecting fraud. A main obstacle to collaborative fraud detection are data confidentiality and privacy where parties are reluctant to share their possibly sensitive data. We briefly analyze the statutory situation, and the effects of some fraud-detection specific privacy-enhancing technologies in Sect. 2.1 and in Sect. 2.2 we summarize the techniqual constraints to be respected to be able to effectively detect fraud incidents on pseudonymized data. In Sect. 3 we present three of our solutions for collaborative privacy-enhanced fraud detection and conclude with a brief outlook in Sect. 4.

2 Requirements

2.1 Dimension of Privacy Law

Setting the scene for a discussion of fraud detection with respect to privacy law, this technology can be briefly described as a technique to analyze the behavior of users in a technical environment focusing on harmful behavioral patterns. As an asynchronous process any Fraud Detection System (FDS) requires the storing of activities that occur

Dagstuhl Seminar Proceedings 08302
Countering Insider Threats
http://drops.dagstuhl.de/opus/volltexte/2008/1794

on a system or use data that is held for other purposes. This information, is from the legal point of view and as long as the users of a system happen to be natural persons, personal data. As a result the application of FDS in most real world scenarios has to adapt to the privacy law, which leads to the question, what requirements have to be fulfilled by the system provider, to ensure that the application of his FDS is legitimate from the legal perspective. Afterwards it is possible to think of technical solutions that might be helpful to meet these requirements and to improve the data privacy standards of the whole system. As one example implementing the EU Privacy Directive 95/46/EG we focus on the German privacy law, which is one of the most strict implementations of the directive, such that results that hold for this law, will with high probability also hold for less stringent implementations.

Main exigencies are the guarantee of transparency and purpose-binding. Firstly the user has to be informed about the data that is collected and processed, about the objectives the controller (typically the one who uses an FDS) and the identity of the controller, if that is not obviously cognizable by the purpose. Secondly the controller is bound to that purpose. In the case that the data is collected for other reasons (for example transactional data) the legislator restricts the circumstances under which a change of the objectives is possible. If the data is collected just for the purpose of system security any processing for other reasons is prohibited. Any technical change, which serves both or one of these requirements of the system, improves the legal compliance.

One important issue from the legal point of view is that the analyzing process runs on the original data that contains personal data of the persons concerned. The pure existence and accessibility of this data is a potential risk that the data might be used in an illegal way. A solution that might limit or even eliminate this risk could significantly improve the legal compliance of FDS. A solution that is at hand consists of three important characteristics Sect. 3.1.

1. The data that is stored or analyzed is preliminarily pseudonymized. To the FDS-Administrators it contains no personal data because linking the data to the "persons concerned" causes a disproportional effort.
2. The mapping of the original user-identifying attributes to the pseudonyms used in the FDS is held by a trusted third party - the data privacy law establishes a data privacy official which is widely privileged and obliged to discretion. This person can be qualified as trusted.
3. The exposure of the pseudonyms is limited to cases in which the analysis of the data raises a manifest suspicion against the "person concerned"
   (a) In cases where the behavioral pattern that leads to this suspicion was preliminary known, this process can be automated. That shall be defined as *technical purpose-binding*.
   (b) When a new kind of harmful behavior is recognized the exposure of a pseudonym affords the collaboration of the data privacy official and the FDS. That shall be called *organizational purpose binding*.

By judicial argumentation we found that this solution at least supports and strengthens the compliance of an FDS.

2.2 Technical Dimension

Considering requirements of law and competitive enterprises we find the following conflicting goals for collaborative fraud detection:

1. *detection effectiveness*,
2. *privacy* of honest individuals,
3. organizational business-related *confidentiality* requirements, and
4. *efficiency*.

For this contribution we focus on *lossless (information) reductions* that work by *splitting the information* contained in structured data objects into *open data* and *private (covered, masked, blinded) data* before forwarding it to other organizations [1].

The open data of a lossless reduction is sufficient for detecting fraud, possibly in conjunction with exploiting certain properties of the private data, or in conjunction with some *support data* that must be additionally generated depending on the specific application. For this chapter we focus on the misuse detection approach, where known fraud schemes are modeled and matched to the current business process execution event stream. If a fraud incident is detected, the private data may be disclosed for the legal purpose of effectively handling the fraud case. The respective open data together with the private data allows for the reconstruction of the original information, subject to the detected fraud. The data with the disclosed information can be used to hold perpetrators accountable. The following requirements are crucial for the misuse detection approach:

R1: certain data fields (except for timestamps) need to be compared to certain other data fields for equal content, or equal prefix content
R2: certain data fields (except for timestamps) need to be compared to values outside of the open data, e.g. constant values, entries of a database
R3: distances of alarm timestamps need to be computed and compared to values outside of the open data, i.e. a constant value
R4: the order of alarm timestamps needs to be determined

We have then refined these requirements to the technical level of pseudonymizing data fields in audit data for fraud detection. In summary the resulting requirements are that (1) the generated pseudonyms respect the syntax of the audit data, (2) pseudonyms must be equal for a given data value in different data fields, if R1 requires so, (3) the data reduction (e.g. hash) function is collision resistant, and (4) pseudonyms for constant values can be computed during detection to be compared to pseudonyms in the audit data, if R2 requires so. Additionally, if R1 does not require that the pseudonyms for a given data value in different data fields needs to be compared, the pseudonyms for the data value shall be different, in order to reduce inferences with respect to the value.

3 Selected Solutions

3.1 Pseudonymizing Audit Data

The above result translates to our approach for pseudonymizing audit data, such that it can still be analyzed for occurrences of pattern, e.g. characterizing specific known fraud scenarios; while while balancing the conflicting requirements for accountability and anonymity [2]. In our approach audit data is pseudonymized immediately after it has been generated, such that users appear under pseudonyms in the audit data. The pseudonymized audit data maintains the degree of linkability required for fraud detection. The pseudonymization process also produces additional *private data* that allows for the recovery of the original data, subject to specified conditions. The fraud detection component merely analyzes the pseudonymized audit data with respect to fraud (suspicions). Only if a (*threshold*) *alert* occurs, i.e., a fraud suspicion has been detected, the private data can be used to disclose the original data. After data disclosure accountability can be established in order to further investigate anbd mitigate the current incident.

For the private data, the approach leverages threshold schemes for cryptographic secret sharing: The fraud suspicions for intrusion detection are modeled as thresholds of secret sharing schemes. We have constructively shown that determining appropriate thresholds can be achieved by statically analyzing the given fraud models [3]. The private data contains the encrypted identifying data that is replaced by the pseudonyms, and it contains shares of the respective decryption keys. As a result, the disclosure of the encrypted identifying data is enforced cryptographically, such that decryption is possible, if and only if the pseudonyms are involved in a sufficient suspicion of fraud (*technical purpose binding*), i.e., the number of shares associated with the pseudonyms exceeds the threshold in the model of the misuse suspicion. Note that it may be necessary to provide the ability to recover the decryption keys independently of a priori defined models of misuse suspicion in order to investigate misuse that has not (yet) been modeled. In that case, the grounds for decryption must be scrutinized by one or more trusted parties (*organizational purpose binding*).

3.2 Synchronization on Hidden Time Stamps

An important indicator for correlation is time and we introduced a scheme where a third party can compare timestamps, but only if they are within a certain distance threshold. If the distance of the two timestamps exceeds the threshold, the third party cannot make any conclusion about the timestamps. We achieve this by pseudonymizing the timestamp using a message authentication code, but giving local distance information. The pseudonyms are aligned to so-called grid points and if they match one can compute the distance from the local distance information [4].

A problem with this or any approach that allows similar computations is that an attacker can gain additional information given a set of timestamps. We showed that it is unavoidable given such a mechanism. A second problem with the timestamp pseudonymization approach is that it requires synchronized clocks. Synchronized clocks can be a very strong assumption in distributed systems and as a consequence logical clocks have been developed, that replace absolute time with causality information.

Vector clocks provide the most causality information of logical clocks. In a vector clock system each process keeps an estimate of the other processes' clocks. Every time a message is sent these estimates are updated. During such an exchange vector clocks may reveal information about communication with other parties. In order to prevent this leakage, we developed privacy-preserving vector clocks that encrypt the clock information [5]. All operations of vector clocks are replaced with simple secure computation protocols on the encrypted data and allow a third party to compare vector clock timestamps without gaining any additional knowledge.

4 Outlook

Not only for timestamp comparisons, but for most operations in collaborative fraud detection, privacy-preserving alternatives have been defined. Practical systems using data pseudonymization and hiding (as in the timestamp pseudonymization approach) are currently emerging, but it is becoming obvious that the security provided by the combination of privacy-preserving techniques is limited. An avenue of future research is therefore to either model formally the provided privacy and security of combined approaches and hopefully provide some useful limitations. A number of attacks can also possibly be found on the pseudonymized data, as has been done for a large collection of anonymized data.

Another avenue of research is to strengthen the security of the approaches. This usually comes at the expense of further performance, but some modern security technologies can help. We currently investigate an outsourcing scenario, where an enterprise outsources non-core services as well as the ability for fraud detection. In this scenarion the enterprise as well as the service providers wish to keep their data confidential. Also the third party fraud detection provider considers his domain knowledge a competitive asset, which he wishes to keep secret. We developed a searchable encryption scheme that protects both the privacy of the fraud detection queries as well as the privacy of the databases.

References

1. Ulrich Flegel and Joachim Biskup. Requirements of information reductions for cooperating intrusion detection agents. In Günter Müller, editor, *ETRICS*, volume 3995 of *Lecture Notes in Computer Science*, pages 466–480. Springer, 2006.
2. Joachim Biskup and Ulrich Flegel. Threshold-based identity recovery for privacy enhanced applications. In Sushil Jajodia and Pierangela Samarati, editors, *CCS*, pages 71–79. ACM, November 2000.
3. Ulrich Flegel. *Privacy-Respecting Intrusion Detection*, volume 35 of *Advances in Information Security*. Springer, New York, 2007.
4. Florian Kerschbaum and Julien Vayssière. Privacy-preserving logical vector clocks using secure computation techniques. In *ICPADS*, pages 1–8. IEEE Computer Society, 2007.
5. Florian Kerschbaum. Distance-preserving pseudonymization for timestamps and spatial data. In Peng Ning and Ting Yu, editors, *WPES*, pages 68–71. ACM, 2007.