

Data Degradation: Making Private Data Less Sensitive Over Time

Nicolas AnCIAUX^{*}, Luc BouGANIM^{*}, Harold van Heerde^{***}, Philippe Pucheral^{**}, Peter M.G. Apers^{***}

^{*} INRIA Rocquencourt
Le Chesnay, France
<Fname.Lname>@inria.fr

^{**} PRISM Laboratory
University of Versailles, France
<Fname.Lname>@prism.uvsq.fr

^{***} CTIT
University of Twente, The Netherlands
{heerdehw,apers}@ewi.utwente.nl

ABSTRACT

Trail disclosure is the leakage of privacy sensitive data, resulting from negligence, attack or abusive scrutinization or usage of personal digital trails. To prevent trail disclosure, data degradation is proposed as an alternative to the limited retention principle. Data degradation is based on the assumption that long lasting purposes can often be satisfied with a less accurate, and therefore less sensitive, version of the data. Data will be progressively degraded such that it still serves application purposes, while decreasing accuracy and thus privacy sensitivity.

Categories and Subject Descriptors

H.2.7 [Database Administration]: Security, integrity, and protection

General Terms

Algorithms, Management, Security, Human Factors

Keywords

Privacy, limited retention, data degradation

1. INTRODUCTION

Personal digital trails are difficult to protect. As any data, they are exposed to accidental disclosures resulting from negligence or piracy. The personal details of 25 million UK citizens have been recently lost inadvertently and some of the data published by AOL about Web search queries of 657,000 Americans have been deanonymized. Even the most defended servers (including those of Pentagon, FBI and NASA) are successfully attacked. But more, personal digital trails are often weakly protected by obscure and loose privacy policies which are presumed accepted when exercising a given service. This fosters ill-intentioned scrutinization and abusive usages justified by business interests, governmental pressures and inquisitiveness among people. No one is sheltered because common events, like applying for a job or a credit, can suddenly make anybody's digital trail of utmost interest for someone else. Companies like Intelius or ChoicePoint make scrutinization their business while others like ReputationDefender provide a lucrative service to destroy the sensitive part of personal digital trails subject to scrutinization.

We define *trail disclosure* as the leakage of data pertaining to a personal digital trail and resulting from negligence, attack or abusive scrutinization or usage. By definition, its occurrence assumes that all security mechanisms have been bypassed or that the ac-

cess control policy has been defined too weakly. Promoted by most legislation protecting personal data, the *limited data retention* principle consists of attaching a lifetime to a data compliant with its acquisition purpose, after which it must be withdrawn from the system. The shorter the retention period is, the smaller the total amount of data needlessly exposed, reducing the impact of trail disclosure [3]. Beyond the protection of personal data, the limited data retention principle is also a cornerstone of the ISO/IEC 27002:2005 recommendation for protecting enterprise information systems.

However, the limited data retention principle is difficult to put in practice. In some countries, minimal retention periods can be fixed for law enforcement or legal processes purposes (e.g., banking information in UK cannot be destroyed before 7 years). In this case, the same retention limit is used for privacy preservation purposes, for which such limits are usually too large. For the large amount of data not covered by law, the retention limit is supposed to reflect the best compromise between privacy preservation and application purposes reach. In practice, the same data item is likely to serve different purposes, leading to selecting the largest retention limit compatible with all purposes, as suggested in [3]. Moreover, the purposes exposed in most privacy policies are fuzzy enough to defend very long retention limits (years or decades), denaturing the initial principle. As a consequence, retention limits are seen by civil rights organizations as a deceitful justification for long term storage of personal data by companies.

We propose a new approach which opens up an alternative to reason about and implement limited data retention. It is based on the assumption that long lasting purposes can often be satisfied with a less accurate, and therefore less sensitive, version of the data. The objective of the proposed approach is to progressively *degrade* the data after a given time period such that (1) the intermediate states are informative enough to serve application purposes and (2) the accurate state cannot be recovered by anyone after this period, not even by the server. To the best of our knowledge, this approach is the first attempt to implement the essence of the limited data retention principle, which is limiting the retention of any information to the period strictly necessary to accomplish the purpose for which it has been collected. Hence, if the same information is collected to serve different purposes, degraded states of this information and their respective retention limits are defined according to each application purpose.

By degrading attributes forming a quasi-identifier, *k*-anonymisation [4] shares some similarities with our data degradation model. However, both models pursue different objectives. *k*-Anonymisation transforms the database content such that the data of a single individual cannot be distinguished from the one of *k*-1 other individuals. Thus, no other purposes than statistics computa-

Copyright is held by the author/owner(s).

CIKM'08, October 26–30, 2008, Napa Valley, California, USA.

ACM 978-1-59593-991-3/08/10.

tions can be satisfied. By contrast, our data degradation model applies to attributes describing a recorded event while keeping the identity of the user intact. Hence, user-oriented purposes are preserved.

The threat model considered by data degradation is the same as for limited data retention. It does the two following assumptions on the recording system (i.e., the DBMS): (1) the system implements without malice all security policies which have been defined (including retention control); (2) the system cannot prevent all forms of attacks, negligence or weakly defined policies which could expose, at any given time, the personal digital trail of a victim to an adversary.

In a trail disclosure resulting from a *piracy attack*, the adversary is a hacker or a dishonest employee breaking into the server with the objective to get access to a set of user's digital trails (e.g., for a lucrative purpose) and the victims are the targeted users (potentially all users of the system). In a trail disclosure resulting from a weak policy declaration, the adversary is anyone (e.g., an insurance or credit company, an employer, a governmental agency) having a particular interest to scrutinize the digital trail of an identified victim (e.g., a future client or employee, a suspect citizen). In a trail disclosure resulting from a negligence, the adversary is anyone getting access to the disclosed digital trails (e.g., could be internet user in the AOL disclosure scandal) and the victim can be anyone having a singular personal digital trail (e.g., user 4417749 identified as Thelma Arnold, a 62-year-old Georgian widow).

Data degradation however, as any data retention model, cannot defeat trail disclosures performed by an adversary spying the database system from its creation. To be effective, such attack must be repeated with a frequency smaller than the duration of the shortest degradation step. Such continuous attacks are much more easily detectable thanks to Intrusion Detection Systems and Auditing Systems. Besides, data degradation is complementary to any access control techniques.

An important question is whether data degradation can be reasonably implemented in a DBMS. Even guaranteeing that data cannot be recovered after a regular delete is not easy [1]. Data degradation is a more complex process which includes physical data deletion but impacts more thoroughly the data storage, indexing, logging and locking mechanisms to deal with data traversing a sequence of states. As retention limits become shorter, the number of degradation steps increases and the performance problem arises. For more details we refer to [2].

2. MOTIVATING SCENARIO

To illustrate our approach, we sketch here a concrete scenario where data degradation can be applied. It takes place in an organization which wants to provide new services to its employees, enabled by the analysis of employees' web activity. Each internet access is captured as one tuple in a personal *trail* containing four attributes: *ID* is the employee's identifier, *URL* is the resource locator of the visited page, *TIME* and *DUR* are respectively the date and the duration of the visit.

Such personal trails are sensitive and may violate privacy if disclosed. Indeed, a full trail gives complete information about web activities (what, when and how long). Past timetables can be reconstituted simply from attributes ID and TIME. Moreover, the full trail may give sensitive information about any given employee such as the delay before starting a task (reactivity), the

time taken to accomplish a task (productivity), the time spent on useless or extra work activities, etc. Notice that we consider "an organization" in this scenario, but the privacy invasion is even stronger when tracking people within the private sphere.

Examples of services enabled in such a context are:

S1: Synchronous co-browsing. Employees can see the current web trail of colleagues. For example, people making a web survey to prepare a meeting may use this feature to search in parallel.

S2: Asynchronous co-browsing. Each employee can retrieve the past web trails of other colleagues. This would enable people to suspend/resume investigations [5]

S3: Passive co-browsing. Employees benefit from implicit feedback of colleagues, computed from their past trails. In particular, each link in visited web pages is complemented with statistics generated by colleagues having followed the link (e.g., name along with duration and frequency).

S4: Profiling employees. HR department often requires employees' domains of expertise, e.g., to identify interesting trainings to be proposed or to find the correct person to which delegating the organization of an event (e.g., a seminar). Domains of expertise are obtained from web trails (e.g., MySQL experts most probably visit mysql.com often and extensively).

To fulfil its purpose, service *S1* requires accurate recent histories (e.g., 1 day) to enable parallel investigations. *S2* requires longer histories (e.g., 1 week) but with degraded attributes *TIME* and *DUR*: the date in days can replace the exact time in seconds; the exact duration can be degraded to a value from the domain *Binary_dur* = {*short*; *long*}. *S3* requires even longer trails (e.g., 1 month) and can be fulfilled with a *TIME* value expressed in weeks and a *URL* value degraded to its static part (i.e., the path without parameters, hence, no content generated based on users' inputs). *S4* requires long term trails (e.g., 1 year) with *TIME* in months and *URL* reduced to the domain subpart (e.g., "MySQL.com").

3. RESEARCH AGENDA

Data degradation strictly implements the limited retention principle, increasing privacy with respect to trail disclosures, without touching the ability for applications to reach its purposes. We already devised a simple and limited version of the degradation model [2]. Current work aims at a better understanding of the impact of degradation on traditional database technology. In the next steps, we investigate necessary extensions of the model in terms of flexibility and usability, where users are better involved into the degradation process (e.g., personalized degradation policies given optins purposes), and where the content of the data influences the degradation.

4. REFERENCES

- [1] Stahlberg, P., Miklau, G., Levine, B.N. Threats to privacy in the forensic analysis of database systems. In SIGMOD, 2007
- [2] N. L. G. Ancaiaux, L. Bouganim, H. J. W. van Heerde, P. Pucheral, and P. M. G. Apers. InstantDB: Enforcing Timely Degradation of Sensitive Data. In ICDE, April 2008
- [3] Agrawal, R., Kiernan, J., Srikant, R., Xu, Y. Hippocratic databases. In VLDB, 2002
- [4] Sweeney, L. K-anonymity: A model for protecting privacy. Int. Journal on Uncertainty Fuzziness and Knowledge-based Systems, 10, 5 (Oct. 2002)
- [5] Morris, M.R., Horvitz, E. SearchTogether: An Interface for Collaborative Web Search. Proceedings of ACM UIST 2007.