

The meaning of logs

Citation for published version (APA):

Etalle, S., Massacci, F., & Yautsiukhin, A. (2007). The meaning of logs. In C. Lambrinoudakis, G. Pernul, & A. Min Tjoa (Eds.), *Trust, Privacy and Security in Digital Business (4th International Conference, TrustBus 2007, Regensburg, Germany, September 4-6, 2007. Proceedings)* (pp. 145-154). (Lecture Notes in Computer Science; Vol. 4657). Springer. https://doi.org/10.1007/978-3-540-74409-2_17

DOI:

[10.1007/978-3-540-74409-2_17](https://doi.org/10.1007/978-3-540-74409-2_17)

Document status and date:

Published: 01/01/2007

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

The Meaning of Logs^{*}

Sandro Etalle^{1,2}, Fabio Massacci¹, and Artsiom Yautsiukhin¹

¹ University of Trento, DIT

² University of Twente, The Netherlands

sandro.etalles@utwente.nl, {evtiukhi,massacci}@dit.unitn.it

Abstract. While logging events is becoming increasingly common in computing, in communication and in collaborative environments, log systems need to satisfy increasingly challenging (if not conflicting) requirements. In this paper we propose a high-level framework for modeling log systems, and reasoning about them. This framework allows one to give a high-level representation of a log system and to check whether it satisfies given audit and privacy properties which in turn can be expressed in standard logic. In particular, the framework can be used for comparing and assessing log systems. We validate our proposal by formalizing a number of standard log properties and by using it to review a number of existing systems. Despite the growing pervasiveness of log systems, we believe this is the first framework of this sort.

1 Introduction

In the past few years we have witnessed a struggle between two competing forces: privacy protection and fight against cyber-crime. Privacy protection has called for new regulations [15,5], new technological solutions [2,4] and re-thinking of business interactions [8]. On the other hand, efforts in countering cyber-crime, have led to increasingly invasive laws [13] and new auditing techniques [24,3,1].

Such clash is most evident in the realm of auditing in general, and in the regulations on how logs should be taken, maintained and deleted in particular. A folklore pun well describes the problem as follows: if logs mention private information they are forbidden and if they do not - they are useless. For instance, an important privacy requirement for log systems is the compliance with the maximal retention period (the time after which a company has to delete user's data) which in some cases must be determined on a need basis [2,4,12] (e.g. service providers have to delete logged data when they do not need it any longer to offer their services). On the other hand, logs have to be kept for audit purpose or for computer forensics. This problem goes beyond privacy in databases: Internet Service Providers (ISPs) have similar regulations [10,27]. A recent amendment to EU Directive N 2002/58/EC [13] requires service providers (i.e. ISPs, e-mail services, communication providers) to store their logs for not less than 6 months to help law enforcement agencies. Consequently, sensitive information about a user may be in the system after the user's own account has been deleted.

^{*} This work is partly supported by the project EU-IST-IP-SERENITY.

We notice that even though logs are ubiquitous in computing and telecommunication security and there is a significant amount of work on *analyzing logs*¹, we find relatively few papers on *design and analysis of log systems* [28,22,14] and on what security properties a log system may or should exhibit [17]. This is somehow striking in comparison with the large body of work on security properties for e.g. security protocols or security models for access control.

In this paper we define a formal framework for modeling and analyzing log systems, which allows one to provide a high level specification of a log system, thereby allowing her to check whether it has the expected properties (e.g., if it meets given privacy or audit requirements). In particular, our framework can be used to compare different log systems with each other.

To validate our proposal, we include a survey of the requirements that are applicable to log systems, and we show how to represent them formally. In addition, we have considered a number of log systems taken from the literature and we show how they compare to each other when modeled in our framework.

2 Log Requirements

First we need to specify some notation: here we talk about (real world) *events* and call *trace* a sequence of events. In turn, a trace may be logged in a *log*; by *recovering a trace* we indicate the action of associating to a given log the trace(s) of events that could have generated it. To be useful, logs often have to meet various requirements. Here we list the most common of them (collected from various papers in the literature: ISO17799 [16], CC [17], [6]); later, we will be able to give a precise formalization of these properties.

- *Completeness*: All events in a trace of events can be recovered from its log.
- *Partial Completeness*: All events in a trace of events matching a given property (relevant events) can be recovered from its log.
- *Past Independence*: In a trace, older events have no influence on the log and recovery of newer events.
- *Future Independence*: In a trace, newer events have no influence on the log entries of older events, nor on their recovery.
- *Context Independence*: The conjunction of past and future independence.
- *Chaining*: Valid logs become invalid if an intermediate record is altered.
- *Exactness*: The recovery of a log of a trace is unambiguous: given a log there is a unique trace of events which could have generated it.

These notions allow one to characterize the precision and completeness of the audit.

Events in a trace usually have attributes (e.g. date, user name, address); the following properties concern whether a given log system allows or not to recover a certain attributes. This is particularly important for privacy protection.

- *Complete Anonymity (w.r.t. attribute A)*: The recovery of an event does not give any information on the value of its attribute A.

¹ See the RAID conference series, for example.

- *Ambiguity (w.r.t. attribute A)*: The recovery of an event does not allow one to establish the value of its attribute A .
- *Linkability (w.r.t. attribute A)*: It is possible to determine whether two recovered events had the same value for attribute A (notice that the system could still be ambiguous w.r.t. A).
- *Positive/Negative Monotonicity*: Newer events do not introduce/reduce anonymity in older events.

These notions allow one to characterize the extent to which a log system protects private information. Linkability is common because it allows precise auditing even if some information is hidden.

An example of a log system which is not past independent is e.g. the log system in Linux, which records a user's name together with the assigned pseudonym². An example of a system which does not satisfy future independence is one in which log entries are destroyed after a given retention time. Such system is not complete either. Positive monotonicity is important when we do not want to lose information we logged. Negative monotonicity is important from the privacy perspective.

3 A Formal Model of Logs

To introduce our framework we give the definition of the world model, which is the environment where logging takes place. Here and in the sequel, given a set X we denote by 2^X its powerset and by X^* the set of sequences of elements from X .

Definition 1. *A World Model is a tuple $\langle E, T, AD, \{AF_i\}_{i \in I} \rangle$; where: E is a set of real world events; $T \subseteq 2^{E^*}$ is a set of valid traces; AD is a general attribute domain which includes all possible dimensions (e.g. strings, real, data, etc.); $\{AF_i\}_{i \in I}$ is a set of attribute functions, which given a sequence of events return the corresponding sequence of attribute values: $E^* \mapsto (2^{AD})^*$ (e.g. $user()$, $date()$).*

Now we can define a log system which records events from the world model.

Definition 2. *Let $WM = \langle E, T, AD, \{AF_i\}_{i \in I} \rangle$ be a world model, then a Log System for WM is a tuple $\langle R, L, Log(), Rec() \rangle$; where: R is a set of records; $L \subseteq 2^{R^*}$ is a set of valid logs in the system. $Log: E^* \mapsto R^*$ is a function mapping a trace of events into the corresponding log. $Rec: R^* \mapsto 2^{E^*}$ is the function which given a log returns the corresponding set of traces (of events).*

In other words, the recovery function $Rec()$ maps a log into the set of traces of events that could have originated the log. Considering that some information might be lost during the logging process (e.g., in the case of anonymous systems), it can well be the case that the $Rec(l)$ contains more than one trace. We denote events by e and records by r . A trace is represented by $t = \langle e_1, e_2 \dots e_n \rangle$. Similarly, a log is denoted by $l = \langle r_1, r_2 \dots r_n \rangle$. In the sequel, $x \circ x'$ means that sequence x' is appended to (after) sequence x preserving elements order.

² In Linux pseudonyms are used for convenience, and not to preserve users' privacy.

Example. Let us describe a log system using pseudonyms (as in [18]). Consider a hospital-based database containing medical and personal data of patients. The hospital keeps track of all accesses to the database both to prevent data linkage (privacy) and for accountability purposes. To define the World Model, we introduce the following domains: *Time* is a set of positive integer values which denote time; *Operator* is a set of users (represented by strings) who have access to patient data; *Patient* is a set of all possible patients of the hospital (represented by strings); *Status* is the set: {successful, failed} used to denote whether an action was carried out successfully or not. The general attribute-domain is $AD = Time \cup Operator \cup Patient \cup Status$. Finally, attribute functions are defined and named according to the domains above $AF = \{Time(), Operator(), Patient(), Status(), Data()\}$. In the world model, there are six types of events (here, $\tau \in Time$; $o \in Operator$; $p \in Patient$; $s \in Status$):

$$E = \{ \begin{array}{ll} \textit{login}(\tau, o, s) & \text{(Operator) } o \text{ logged-in at time } \tau; \\ \textit{logoff}(\tau, o, s) & o \text{ logged-off at time } \tau; \\ \textit{add}(\tau, o, p, s) & o \text{ added the record of } p \text{ to the system at time } \tau; \\ \textit{read}(\tau, o, p, s) & o \text{ read the record of } p \text{ at time } \tau; \\ \textit{update}(\tau, o, p, s) & o \text{ updated the record of } p \text{ at time } \tau; \\ \textit{delete}(\tau, o, p, s) & o \text{ deleted } p \text{ from the system at time } \tau \end{array} \}$$

Having defined the set of possible events, a valid trace is any ordered (in time) sequence of such events. $T = \{t \in E^* \mid \forall e_i \in t \wedge \forall e_j \in t. i < j \implies Time(e_i) < Time(e_j)\}$. We can now move on to the definition of the log system. Let us first define some additional domains: *Patient_id* is a set of all possible identifiers (strings) of all patients; *Record_id* is a set of integers which unambiguously point to a log record. Note that the *Patient* domain from the world model differs from *Patient_id*, as the real names of patients are substituted with pseudonyms. We underline the identifier to refer to the pseudonym, so \underline{p} is the pseudonym of patient p . We also underline the records to distinguish between records and events. The log system has four types of records (here: $j \in Record_id$; $\tau \in Time$; $o \in Operator$; $\underline{p} \in Patient_id$; $s \in Status$):

$$R = \{ \begin{array}{ll} \underline{\textit{add}}(j, \tau, o, \underline{p}, s) & o \text{ added the record of } \underline{p} \text{ to the system at time } \tau; \\ \underline{\textit{read}}(j, \tau, o, \underline{p}, s) & o \text{ read record of } \underline{p} \text{ at time } \tau; \\ \underline{\textit{update}}(j, \tau, o, \underline{p}, s) & o \text{ changed record of } \underline{p} \text{ at time } \tau; \\ \underline{\textit{delete}}(j, \tau, o, \underline{p}, s) & o \text{ deleted } \underline{p} \text{ from the system at time } \tau \end{array} \}$$

Record identifiers (j) are assigned incrementally. We can now define the *Log()* function:

$$Log(t) = \begin{cases} \underline{\textit{add}}(j, \tau, o, \underline{p}, s) \circ Log(t') & \text{if } t = e \circ t' \text{ and } e = \textit{add}(\tau, o, p, s); \\ \underline{\textit{read}}(j, \tau, o, \underline{p}, s) \circ Log(t') & \text{if } t = e \circ t' \text{ and } e = \textit{read}(\tau, o, p, s); \\ \underline{\textit{update}}(j, \tau, o, \underline{p}, s) \circ Log(t') & \text{if } t = e \circ t' \text{ and } e = \textit{update}(\tau, o, p, s); \\ \underline{\textit{delete}}(j, \tau, o, \underline{p}, s) \circ Log(t') & \text{if } t = e \circ t' \text{ and } e = \textit{delete}(\tau, o, p, s); \\ Log(t') & \text{if } t = e \circ t' \text{ and none of the above applies;} \\ \varepsilon & \text{otherwise.} \end{cases}$$

The mapping between a patient and his pseudonym is done with a special binding table to which access is restricted. We assume that *login* and *logoff* events are not logged. For the recovery function, let M be the set of bijective mappings $Patient_id \mapsto Patient$; given $m \in M$ we define R_m as follows:

$$\begin{aligned} R_m(l \circ \underline{add}(j, \tau, o, \underline{p}, s)) &= R_m(l) \circ add(j, \tau, o, m(\underline{p}), s) \\ R_m(l \circ \underline{read}(j, \tau, o, \underline{p}, s)) &= R_m(l) \circ read(j, \tau, o, m(\underline{p}), s) \\ R_m(l \circ \underline{update}(j, \tau, o, \underline{p}, s)) &= R_m(l) \circ update(j, \tau, o, m(\underline{p}), s) \\ R_m(l \circ \underline{delete}(j, \tau, o, \underline{p}, s)) &= R_m(l) \circ delete(j, \tau, o, m(\underline{p}), s) \end{aligned}$$

(where $R_m(\varepsilon) = \varepsilon$); the recovery function is: $Rec(l) = \{t \mid t = R_m(l) \text{ for some } m \in M\}$.

Notice that the recovery function maps a log into a set of traces. Consider the following list of events: *Login(8:58 21/10/2006,³Edward Green,successful)* *Add(10:30 21/10/2006,Edward Green,Mackle Daniels,successful)* *Login(12:00 21/10/2006,Suzi Wallach,successful)* *Changed(12:21 21/10/2006,Suzi Wallach,Paul Anderson,failed)* *Changed(12:22 21/10/2006,Suzi Wallach,Mackle Daniels,successful)* Then the corresponding log is:

Record ID	Cause	Time	Operator	Patient	Status
1	Add	10:30 21/10/2006	Edward Green	102	successful
2	Update	12:21 21/10/2006	Suzi Wallach	101	failed
3	Update	12:22 21/10/2006	Suzi Wallach	102	successful

As one can see the log file itself (without knowledge of the bijection mapping) does not disclose any information about the patients of the hospital other than the fact that records 1 and 3 concern the same patient. If an operator who has no access to the private data tries to recover the log he obtains six possible traces: one for each pseudonym-user assignment.

4 Properties

The formal log system allows us to give a precise definition of the informal properties stated in Section 2, providing us with a basis for assessing and comparing different log systems.

Having a formal definition of these properties is very important to make them precise, which is a less trivial task than it may seem at first. If one argues that a log system where i) everything is logged but ii) old records are deleted is complete then we need to change both the informal and the formal definitions. Let $WM = \langle E, T, AD, \{AF_i\}_{i \in I} \rangle$ be a world model, and $\langle R, L, Log(), Rec() \rangle$ be a log system for WM:

³ Time is stored as an integer value, but for the sake of simplicity it is represented as usual.

Definition 3 (Properties)

- *Trace Completeness*: $\forall t \in T \ t \in \text{Rec}(\text{Log}(t))$.
- *Partial Trace Completeness* (w.r.t. a property P . Here we simply indicate by $P(t)$ the subsequence of t consisting of all and only events satisfying property P). $\forall t \in T \ . \ P(t) \in \text{Rec}(\text{Log}(t))$.
- *Future Independence*: $\forall t, t' \in T \ . \ t' \in \text{Rec}(\text{Log}(t)) \iff \forall t_1 \in T \ \exists t'_1 \in T \ . \ t' \circ t'_1 \in \text{Rec}(\text{Log}(t \circ t_1))$
- *Past Independence*: $\forall t, t' \in T \ . \ t' \in \text{Rec}(\text{Log}(t)) \iff \forall t_1 \in T \ \exists t'_1 \in T \ . \ t'_1 \circ t' \in \text{Rec}(\text{Log}(t_1 \circ t))$.
- *Context Independence*: conjunction of future and past independence.
- *Chaining*: $l \circ \langle r \rangle \circ l' \in L \implies \forall r' \neq r \ \ l \circ \langle r' \rangle \circ l' \notin L$.
- *Exactness*: $\forall t \in T \ \{t\} = \text{Rec}(\text{Log}(t))$

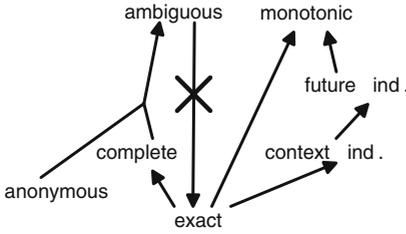
To express most privacy-related properties we need to be able to make the correspondence between single events and single log entries. In particular, if e is an event in a trace t and $t' \in \text{Rec}(\text{Log}(t))$ we have to be able to tell which event in t' corresponds to the original e . We denote this event by $t' \downarrow e$. In most cases, the correspondence function \downarrow is realized quite simply by assigning consecutive numbers to events and log entries.

Definition 4 (Privacy Properties). *Let AF be an attribute function.*

- *Complete Anonymity* (w.r.t. AF): $\forall t \in T \ \forall t' \in \text{Rec}(\text{Log}(t)) \ \forall e_1, e_2 \in t', \ AF(e_1) = AF(e_2)$.
- *Ambiguity* (w.r.t. AF): $\forall t \in T \ \forall e \in t \ |AF(\text{Rec}(\text{Log}(t)) \downarrow e)| > 1$.
- *Linkability* (w.r.t. AF): $\forall t \in T \ \forall e_i, e_j \in t \ . \ AF(\text{Rec}(\text{Log}(t)) \downarrow e_i) = AF(\text{Rec}(\text{Log}(t)) \downarrow e_j) \iff AF(e_i) = AF(e_j)$
- *Positive Monotonicity* (w.r.t. AF): $\forall t, t' \in T \ \forall e \in t \ AF(\text{Rec}(\text{Log}(t)) \downarrow e) \subseteq AF(\text{Rec}(\text{Log}(t \circ t')) \downarrow e)$
- *Negative Monotonicity* (w.r.t. AF): $\forall t, t' \in T \ \forall e \in t \ AF(\text{Rec}(\text{Log}(t)) \downarrow e) \supseteq AF(\text{Rec}(\text{Log}(t \circ t')) \downarrow e)$

We have now the formal machinery to relate some of these properties to each other and these relations are shown in Figure 1 (see [11] for the proof).

Example. Consider again the system shown in the Section 3. The system is not complete since exist events (e.g., $t' = \text{Login}(t, o, s)$) that are not logged corresponding logs; it is partially complete w.r.t. the property P which is true for all events except for *login* and *logout*. The system in our example is context independent because the recovery of a record does not depend on other records; it is not chained since changes in the log are not noticeable by the system since by definition of L any sequence of records from R is valid; it is not exact because pseudonyms are mapped back (by the recovery function) to any person belonging to the set of patients; for the same reason, it is completely anonymous. Notice however that the system is still linkable: it allows us to see if two events pertain to the same patient (though the presence of the pseudonym does not allow to see which patient it is) as every user has only one identifier and visa versa. It is monotonic since it is context independent.



Theorem 1

1. Every exact system is complete.
2. Every ambiguous system is not exact.
3. Every exact system is context independent.
4. Every exact system is (positively and negatively) monotonic.
5. Every future independent system is (positively and negatively) monotonic.
6. Every complete and anonymous system is ambiguous.

Fig. 1. Relationships among properties

Other examples. We now use the properties just defined to assess and compare some existing logging systems.

Pseudonyms based systems [18,19,20]. These systems use pseudonyms to hide user identities to regular log users while allowing special authorized parties (who have access to the pseudonymization function) carry out precise auditing. The basic idea is to substitute private information with an arbitrary string (the pseudonym). The correspondence between pseudonyms and user identifiers is stored in some binding database with restricted access (here we should mention that even the use of pseudonyms does not guarantee complete protection from the use of statistical methods to reconstruct the behavior of user [9,20]).

Linux logs. Sometimes pseudonyms are used for convenience rather than for privacy reasons. The Linux log system is an example of such pseudonymization ante-literam: here, user identities are partially hidden using group and user identifiers which can be considered as pseudonyms (e.g. user "grayyoga" has pseudonym 1001:100). The binding between user and her pseudonym is stored in the */etc/passwd* file. All kernel audit information contains no references to the username but only to the user id⁴. On the other hand, when a new user is added in the system his identity and pseudonym are stored in a log file. This means that even if access to the */etc/passwd* file is denied it is possible to recover a user identity by consulting the log file.

Buschkes-Kesdogan system [7]. Buschkes and Kesdogan proposes a log system which uses group pseudonyms (e.g. for access right management). Before logging into a server a user receives from a Trusted Third Party (TTP) a credential containing a Group Reference Pseudonym (GRP). When she connects to the server, she reveals her credential and the server authenticates the user as a member of a group according to the GRP. The main peculiarity of the log system is that a pseudonym id corresponds to a *set* of user identifiers.

⁴ See <http://www.die.net/doc/linux/man/man8/auditctl.8.html>

IDA system [25]. The IDA log system uses pseudonyms as well, but instead of substituting the private information with a pseudonym encrypts data. The main advantage is that for re-identification of logs no binding database is required: only the decryption key is needed for full recovery.

Waters et al. [28]. Our last example is to the system of Waters et al., where log entries are encrypted as a whole. This solution eliminates the need to store the correspondence between users and their pseudonyms. The disadvantage of this approach is that – in general – searching in encrypted logs is difficult. To overcome this problem the system stores a separate list of keywords for each log entry encrypted with a unique key. This allows the system to carry out limited search actions while offering good data protection. The authors also use a hash function to preserve the order and integrity of the log.

We can now compare these systems along our classification.

	Context Independence	Complete Anonymity	Ambiguity	Monotonicity	Linkability	Exactness ⁵	Chaining
Linux logs	–	–	–	✓	✓	–	–
Lundin–Johnson [18]	✓	✓	✓	✓	✓	–	–
Buschkes–Kesdogan [7]	✓	–	✓	✓	–	–	–
IDA [25]	✓	✓	✓	✓	✓	–	–
Waters et al. [28]	✓	✓	✓	✓	–	–	✓

5 Related Works and Conclusion

There are a few works focusing on audit log properties. Billable and Yee in [6] introduce forward integrity property and propose a system enforcing it. The well-known Common Criteria security standard [17] reports some requirements which we have referred to as properties. The audit requirements specify what should be stored in the logs and how to use the logs collected.

Most log formalizations have been developed for monitoring purposes [23,21,26]. Roger and Goubault-Larrecq [23] investigate linear time logic for log auditing and propose another logic consisting of Wolper-style linear-time formulae which make auditing more efficient. Spanoudakis et al. [26] propose a formal description of compliance checking for web-service based systems. Mansouri-Samani and Sloman [21] present GEM (generalized event monitoring language) which is used for monitoring networks and distributed systems. B. Waters et al. [28] provide a formal description of a searchable temper-resistant log model.

⁵ Note, that all these systems except [7] become exact if the operator has access to the inverse mapping from pseudonyms to user names. Also note, that in contrast to the system shown in our example, the system in [18] is complete.

In this paper we propose an abstract framework for formalizing and reasoning about log systems. Our framework allows one to model a concrete log system and to check whether it satisfies certain properties; in particular it allows one to check whether the system meets various requirements such as the one we have collected from the literature [6,17,25,7].

The practical motivation for realizing this framework is given by the need compare and assess precisely different log systems against the properties they have to satisfy, which in our system can be expressed in a precise way using a simple logic formalism. To validate our framework, we have encoded in it a number of different systems ([18,7,25,28]). To the best of our knowledge, this is the first framework of this kind.

References

1. Agrawal, R., Bayardo, R.J., Faloutsos, C., Kiernan, J., Rantau, R., Srikant, R.: Auditing compliance with a hippocratic database. In: Nascimento, M.A., Özsu, M.T., Kossmann, D., Miller, R.J., Blakeley, J.A., Schiefer, K.B. (eds.) Proceedings of the 30th International Conference on Very Large Data Bases (VLDB'04), pp. 516–527. Morgan Kaufmann, San Francisco (2004)
2. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic databases. In: Bernstein, P.A., Ioannidis, Y.E., Ramakrishnan, R., Papadias, D. (eds.): Proceedings of the 28th International Conference on Very Large Data Bases (VLDB'02). Morgan Kaufmann (2002)
3. Allen, J., Christie, A., Fithen, W., McHugh, J., Pickel, J., Stoner, E.: State of the practice of intrusion detection technologies. Technical Report CMU/SEI-99-TR-028, Carnegie Mellon, SEI (January 2000)
4. Ashley, P., Hada, S., Karjoth, G., Powers, C., Schunter, M.: Enterprise privacy authorization language (EPAL 1.1). Technical report, IBM (October 2003)
5. Baumer, D.L., Earp, J.B., Poindexter, J.: Internet privacy law: a comparison between the united states and the european union. *Computers & Security* 23(5), 400–412 (2004)
6. Bellare, M., Yee, B.: Forward integrity for secure audit logs. Technical report, University of California at San Diego (1997)
7. Buschkes, R., Kesdogan, D.: Privacy enhanced intrusion detection. In: Mueller, G., Rannenber, K. (eds.) Proceedings of the Conference on Multilateral Security for Global Communication, pp. 187–207. Addison-Wesley-Longman (1999)
8. Cranor, L., Langheinrich, M., Marchiori, M., Presler-Marshall, M., Reagle, J.: The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C, 1.0 edn. (April 2002)
9. Domingo-Ferrer, J., Torra, V.: Disclosure risk assessment in statistical microdata protection via advanced record linkage. *Statistics and Computing* 13, 343–354 (2003)
10. Electronic Privacy Information Center. Data retention (23/05/2007) Available at <http://www.epic.org/privacy/intl/data%5Fretention.html>
11. Etalle, S., Massacci, F., Yautsiukhin, A.: The meaning of logs. Technical Report TR-CTIT-07-24, Centre for Telematics and Information Technology, University of Twente (2007)
12. EU. Directive 95/46/EC of the european parliament and of the council (1995)

13. EU. Directive 2006/24/EC of the European Parliament and of the Council. Official Journal of the European Union, 105/54 (2006)
14. Hansen, J.V.: Audit considerations in distributed processing systems. *Communications of the ACM* 26(8), 562–569 (1983)
15. HIPAA.: Health insurance reform: Security standards; final rule. *Federal Register* 68(34), 8333–8381 (2003)
16. ISO/IEC. Information technology–Security techniques–Evaluation criteria for IT security (November 2001)
17. ISO/IEC. Common Criteria for Information Technology Security Evaluation. Common Criteria Project Sponsoring Organisations, 2.2 edn. (January 2004)
18. Lundin, E., Jonsson, E.: Privacy vs. intrusion detection analysis. In: *Proceedings of the 2nd International Symposium on Recent Advances in Intrusion Detection* (1999)
19. Lundin, E., Jonsson, E.: Anomaly-based intrusion detection: privacy concerns and other problems. *Computer Networks* 34(4), 623–640 (2000)
20. Malin, B., Sweeney, L.: How (Not) to Protect Genomic Data Privacy in a Distributed Network: Using Trail Re-identification to Evaluate and Design Anonymity Protection Systems. *Journal of Biomedical Informatics* 37(3), 179–192 (2004)
21. Mansouri-Samani, M., Sloman, M.: GEM: a generalized event monitoring language for distributed systems. *Distributed Systems Engineering Journal* 4, 96–108 (1995)
22. Ohtaki, Y., Kamada, M., Kurosawa, K.: A scheme for partial disclosure of transaction log. *IRICE Transaction on Fundamentals of Electronics Communications and Computer Sciences* E88(1), 222–229 (2005)
23. Roger, M., Goubault-Larrecq, J.: Log auditing through model checking. In: Young, D.C. (ed.) *Proceedings of the 2001 IEEE Computer Society Security Foundations Workshop*, pp. 220–236. IEEE Computer Society Press, Los Alamitos (2001)
24. Schneier, B., Kelsey, J.: Secure audit logs to support computer forensics. *ACM Transactions on Information and System Security* 2(2), 159–176 (1999)
25. Sobirey, M., Fischer-Hoebner, S., Rannenber, K.: Pseudonymous audit for privacy enhanced intrusion detection. In: Yngstroem, L., Carlsen, J. (eds.) *Proceedings of the IFIP TC11 13 international conference on Information Security (SEC '97) on Information security in research and business*, London, UK, 1997, pp. 151–163. Chapman & Hall, Ltd (1997)
26. Spanoudakis, G., Mahbub, K.: Non intrusive monitoring of service based systems. *International Journal of Cooperative Information Systems* 15(3), 325–358 (2006)
27. Statewatch. UK-EU call for mandatory data retention of all telecommunications (2005), Available at <http://www.statewatch.org/news/2005/jul/05eu-data-retention.htm>
28. Waters, B.R., Balfanz, D., Durfee, G., Smetters, D.K.: Building an encrypted and searchable audit log. In: *Proceedings of the 11th Annual Symposium on Network and Distributed System Security*, San Diego (2004)