

Safekeeping Digital Evidence with Secure Logging Protocols: State of the Art and Challenges

Rafael Accorsi
Department of Telematics
University of Freiburg, Germany
raccorsi@acm.org

Abstract

While log data are being increasingly used as digital evidence in court, the extent to which existing secure logging protocols used to collect log data fulfill the legal requirements for admissible evidence remain largely unclear. This paper elucidates a subset of the necessary secure requirements for digital evidence and extensively surveys the state of the art secure logging protocols, thereby demonstrating that none of the current protocols completely fulfills the elucidated requirements for admissible evidence. In analyzing the shortcoming of logging protocols, the paper also elaborates on the related research challenges.

Keywords

Digital evidence; Preservation; Secure logging protocols; Legal and security requirements.

1. Introduction

Log data record the “dynamics” of a system, i.e. the *events* that change a system’s state. Each event includes information about when it occurs, which subject triggers it, and which objects it involves. Hence, rather than a static snapshot of a system’s state, log data allow the reconstruction of complex event chains, thereby being a central source of digital evidence for legal disputes.

To be admitted as evidence, log data must fulfill several (legal) requirements regarding the *architecture* and the *cryptographic operations* used to collect and protect log data. While the former are well-studied [25], the latter remain largely unclear. In consequence, system engineers are often confused by the parlance used to define the legal requirements for admissibility. Moreover, since requirements are not completely clear, there is a mismatch between what has to be offered and what a particular logging protocol in fact fulfills. Such a mismatch has a fatal implication: more often than not, digital evidence gained from log data loses on probative force or it is not admissible in the first place [12].

Two questions are of foremost relevance here: (1) how legal requirements translate into security requirements common to system engineers and (2) which and how security requirements are satisfied by the existing secure logging protocols prescribing such cryptographic operations.

This paper addresses these questions, focusing on the latter. First, in Section 2 it elucidates what security requirements are necessary – yet not sufficient – for logging protocols to render

log data as admissible digital evidence, expressing them as usual protection goals of computer security. This is not meant to be a complete treatise on admissibility for digital evidence, though. What is admissible in a court of law varies greatly among the legislations (see, e.g., [39]). Hence we consider only a common, in fact basic tenet underlying these legislations. Second, based on a threat model for logging protocols introduced in Section 3 and on a comprehensive survey of state of the art secure logging protocols given in Section 4, in Section 5 the paper reports on what secure logging protocols have to offer for admissible evidence.

This is the first comprehensive analysis of existing secure logging protocols for digital evidence safekeeping. It focus on their operation, revealing that state of the art protocols fail to lay a basis for admissible digital evidence. We observe two issues: first, existing protocols focus either on the transmission or on the storage of log data, whereas both phases must be equally covered. Second, even protocols designed for the respective phases contain vulnerabilities that lead to erroneous evidence, so that log data thought as being authentic are in fact manipulated. Based on these observations, in Section 6 the paper elaborates on a number of research issues to advance secure logging protocols to fully satisfy the evidentiary standards.

Logging architecture: We assume the architecture proposed in [28]. It is based on principals taking the roles:

- *Device*: senses the events and generates the corresponding log messages. In some protocols, the device applies cryptographic operations to protect log messages, e.g. by signing them for entity authentication.
- *Relay (optional)*: receives log messages and forwards them to a relay or collector. Relays modify the headers of the corresponding datagrams, not messages' contents.
- *Collector*: receives log messages and generates the corresponding log file (also called “audit trail”). The collector may apply cryptographic operations to protect recorded entries against, e.g., illegal reading and deletion.

Two additional roles may exist: \mathcal{V} is a semi-trusted verifier and \mathcal{T} is a trusted verifier. These principals assert the authenticity of and administrate the secrets used to protect the audit trail.

Log files are either *sequential* or *circular*. The former grow indefinitely as entries are appended. For the latter, there is a bound – e.g. the total number of log entries or file size – that whenever reached, forces the oldest entries to be overwritten with the new ones. In the context of digital evidence, log files are taken as being sequential.

2. From legal to security requirements

If evidence gained from log data is to influence the outcome of a legal dispute, it must be *admissible* [22]. Admissibility is a legal concept that prescribes requirements to judge the acceptance of any kind of evidence – be it digital or not. Kenneally surveys several ways in which log data can be eventually admitted as evidence [24]. Underlying these postures, the basic evidentiary tenet governing admissibility determinations is that there are guarantees of “*trustworthiness*” associated with the evidence, otherwise it can be challenged [12]. (The term “trust” is very unsettled. We define it as the measure of belief in the soundness of a piece of digital evidence.)

To-date there is no consensus on how to characterize trustworthiness as *security* requirements. Clearly, collected log data must be a faithful representation of the events communicated by the devices. The collection itself must guarantee completeness and authenticity. The provenance of data must be traceable and guaranteed. Moreover, once stored, it should not be possible to alter log data in a way that cannot be detected by a verifier. But are these the sole requirements to be guaranteed by secure logging protocols? If not, what are the others?

We translate Cohen’s “trustworthiness” requirement into security requirements that guarantee the *authenticity* of audit trails, thereby distinguishing between the requirements on the transmission of log messages and those on their storage. In doing so, we explicitly focus on two of the most relevant aspects of digital evidence safekeeping. Further work aims at extending this survey correspondingly, as reported in Section 7.

2.1. Transmission phase

Log *messages* are in transit between the device and the collector, where the following requirements are set:

- *origin authentication*: the collector must ensure that log messages were sent by authorized devices.
- *message confidentiality*: the log messages must remain confidential during the transmission.
- *message integrity*: the message cannot be modified during the transmission.
- *message uniqueness*: log messages are logged just once.
- *reliable delivery*: log messages sent by the device must reach the collector. This is not related to denial of service (DoS) attacks, but with the transport protocol employed.

These requirements must also hold in cases where the collector communicates with \mathcal{V} and \mathcal{T} .

2.2. Storage phase

In the storage phase, the collector appends log *entries* to the audit trail. The following requirements are set:

- *entry accountability*: log entries must include information with regard to the subject (i.e. device and collector) that appends the entry to the audit trail.
- *entry integrity*: audit trails cannot be changed once recorded. Specifically, we define integrity as “accuracy” (log entries were not modified), “completeness” (log entries were not deleted), and “compactness” (log entries were not illegally appended).
- *entry confidentiality*: entries are not stored in clear-text.

Tamper evidence is needed while enforcing some of these requirements. If an attacker succeeds in violating the integrity, e.g. by altering the contents of a log entry, an auditor must be able to detect this tampering.

These security requirements are necessary, yet not sufficient. Complementary, non-architectural measures protecting, e.g., the access to log files and the isolation of the logging processes to avoid covert-channels and side-channels are necessary as well. However, they traditionally fall outside the scope of secure logging protocols.




Figure 1. Classes of secure logging protocols

3. Attacker models for logging protocols

An *attacker model*, i.e. a model of the capabilities of attackers, is necessary when analyzing the threats to the audit trails and the protection offered by secure logging protocols.

One of the shortcomings in analyzing digital evidence and their admissibility is the absence of well-accepted attacker models for audit trails. Nevertheless, [2], [30], and [38] give informal descriptions of attacker capabilities. Following the distinction for security requirements above, we separate the attacker’s capabilities for the transmission and storage phases. Analyses of logging protocols must must consider both phases.

3.1. Attacker model for the transmission phase

This phase considers an attacker residing in the network and attacking log messages between the device and the collector over a relay, and those between the collector and \mathcal{V} and \mathcal{T} .

The attacker model for this setting is that of Dolev-Yao [14]. The attacker controls the communication channel in the sense that he can insert, eavesdrop, delay, schedule, modify, and completely block transmitted messages. However, the attacker can only gain access to the contents of an encrypted message if he possesses the corresponding keys; and he can only produce meaningful encrypted messages if the attacker possesses the necessary submessages. The attacker is also computationally bounded, being subject to the same complexity constraints for cryptographic computation as legitimate principals.

3.2. Attacker model for the storage phase

This phase considers a penetrator inside the collector attacking directly log data at rest in the audit trail. Such an attacker can read, (over)write, (re)move and delete (fields of the) log entries. In doing so, such the attacker may generate message items from the items he already possesses. However, as in the case of log data in transit, the attacker can only obtain the contents of an encrypted log entry if he possesses the corresponding decryption key. It also assumes that the attacker is computationally bounded with regard to cryptography.

Both attacker models assume that attackers do not possess the secrets used in the logging protocol, i.e. the cryptographic keys are initially *not* compromised. If an attacker gains access to these secrets, no security properties can be guaranteed.

4. Secure logging protocols: State of the art

We group the existing logging protocols in four classes according to the techniques they build upon, as in Fig. 1. The “Syslog” class consists of protocols based on the syslog; the “Schneier/Kelsey” class consists of protocols extending that of [38]; the “Ma/Tsudik” class includes the novel logging techniques in [30]; the “Encrypted Search” class contains protocols to search encrypted log data.

4.1. Syslog class

The syslog [28] standard in several *nix distributions, was *not* developed to provide for secure logging. (It transmits and stores messages in clear-text and uses the unreliable UDP to deliver log messages.) This led to a number of extensions of the traditional syslog focusing on the transport of data from the device to the collector, thus neglecting the storage phase. The most relevant extensions of the syslog are:

Syslog-ng: syslog’s new generation [41] provides a secure, reliable log service, being backward compatible with syslog. The improvements against the syslog are: use of TCP as transport protocol; support to IPv6 protocol; and encrypted transmission with the TLS protocol.

Syslog-sign: It adds origin authentication, message integrity, replay resistance, and detection of missing messages to syslog [23]. For this, it uses cryptographically signed messages that contain the signatures of previously sent syslog messages. Fig. 2 depicts the generation of “signature blocks.” (The necessary cryptographic keys are generated in the initialization phase.) The syslog-sign does not provide confidentiality during the transmission of log data, nor accounts for encrypted storage of log entries. Since signature blocks are supposed to be deleted after the authentication, tamper evidence is only partially fulfilled. Implementations of syslog-sign do not exist.

Reliable syslog: While the syslog-ng and syslog-sign are based on the RFC 3164, the reliable syslog is based on the RFC 3195 [32] (as opposed to RFC 3164) and aims to implement reliable delivery for syslog messages. It uses TCP as transport protocol and provides for device authentication, mechanisms to protect the integrity of log messages, and protection against replay attacks. The San Diego Supercomputer Center implemented the reliable syslog [36].

It is worth noting that the original syslog protocol has been composed with other protocols to provide for security properties. Most notable is the deployment of syslog over a TLS protocol session (recently accredited RFC 5425) and the tunneling of syslog messages over SSH. Two observations are due: first, they cannot be seen as “pure” logging protocols any longer; second, they still consider only the transmission phase, as the extensions described above.

4.2. Schneier-Kelsey class

Building on the work of Bellare and Yee [7], Schneier and Kelsey [38] devise a logging protocol for the storage phase. They employ *hash chains* [26] and *evolving cryptographic systems* [18] as cryptographic techniques to protect of audit trails. A hash chain is a successive application of a cryptographic hash function to a string. Evolving keys is an encryption technique

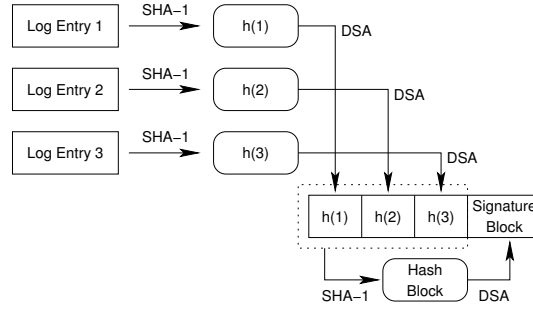


Figure 2. Creation of signature blocks in syslog-sign

$$L_i := \begin{array}{|c|c|c|c|} \hline W_i & \{E_i\}_{K_i} & Y_i & Z_i \\ \hline \end{array}$$

Figure 3. Format of a log entry in Schneier and Kelsey [38]

in which (usually symmetric) keys evolve over time, possibly as a function of the previous keys. The timespan in which the key is kept the same is called an *epoch*. The idea is to limit the damage of compromise: attackers who obtain a key can only decrypt entries stored at that epoch.

Based on these primitives, the authors provide algorithms to create/close an audit trail, append entries to it, and authenticate its entries. (Authentication checks audit trails for tampering.)

Fig. 3 depicts an entry in Schneier and Kelsey. The field W_i is an authorization mask that controls the access to the contents of the entry L_i . (Only subjects authorized in W_i gain access to the contents E_i of L_i .) $\{E_i\}_{K_i}$ stands for the i th event E_i encrypted with the unique, symmetric *evolving* key K_i .

The key K_i is generated by hashing the authorization mask W_i and the entry authentication key A_i . After generating K_i and computing the value Z_i (see below), a new entry authentication key A_{i+1} for the next entry is generated, which is then used to generate K_{i+1} and provide for key evolution. Thus, each entry corresponds to an epoch. Besides assuming the entry authentication keys A are kept secret, the authors assume that A_i and K_i are irretrievably deleted after use.

The field Y_i is the i th link of the hash chain. To intertwine entries of the audit trail, the hash chain link Y_i is generated by hashing the fields W_i , $\{E_i\}_{K_i}$, and the link Y_{i-1} of the previous entry L_{i-1} . In doing so, a verifier can assert the absence of tampering by checking the chain for *broken links*. Finally, the field Z_i stands for the message authentication code (MAC) – computed on the basis of A_i – to authenticate the entry. (Z_i is not included in the generation of the link Y_i .)

While Schneier and Kelsey provide no implementation of their protocol, Chong et al. embed the protocol in a tamper resistant hardware [10] and Accorsi and Hohl focus on the delegation of logging tasks using trusted computing [3].

Informal threat analyses of the Schneier and Kelsey protocol revealed attacks in which modifications of the audit trail could go undetected, thereby violating entry integrity and tamper evidence. This led to the following extensions of [38]:

Stathopoulos et al.: In [40], the authors apply the logging protocol in the public communication networks. Their extension aims to prevent internal attacks, in which a fraudster could reconstruct parts of the audit trail in way that tamper could not be detected. (This attack is possible if the attacker knows some A_j .) They introduce a “regulatory authority” (RA) to assure that the collector follows the protocol accordingly.

To this end, the authors employ an approach similar to the signature blocks used in the syslog-sign (see Section A.2): periodically, RA receives a signature block from the collector, storing these blocks for further analysis. If suspicious arises, the signature blocks is recomputed and compared with those stored at the RA . The authors also informally analyze the security guarantees provided by their logging protocol, demonstrating that the protocol is robust against such attacks.

The BBox: For Alles et al. [5] and Oppliger and Ritz [34], trustworthy digital evidence requires a kind of “black box” logging, similar to flight recorders used in airplanes. With this picture in mind, Accorsi devises the **BBox**: a digital black box [2, Chap. 3]. (A previous version of his work appears in [1], [37].)

By means of an informal threat analysis, the author demonstrates that truncation attacks are possible. The attacker may remove the last n entries of the audit trail in a way that goes undetected at the authentication. (Truncation goes undetected because it does not lead to a broken hash chain.) To fix this vulnerability, as well as the one introduced by Stathopoulos, Accorsi modifies the protocol of Schneier and Kelsey and applies trusted computing modules and novel cryptographic approaches to digital evidence [31]. In doing so, he focuses not only on the storage and authenticated, keyword-based retrieval of audit trails, but also defines security protocols for the transmission of log data between the devices, the **BBox**, and external auditors.

Assuming a PKI, each log message is encrypted and signed for the transmission to the collector. The **BBox** ensures the origin and integrity of the log message, while the replay attacks are addressed by means of sequence numbers and timestamps. As for the storage, the author removes the field Z_i of an entry, requiring the hash chain links to be signed by the **BBox**. The authentication of the audit trail detects truncation and other kinds of tampering.

Accorsi (to a certain extent formally) demonstrates the security properties provided by the **BBox** and successfully carries out a series of tests with the resultant implementation.

The Logcrypt: Holt devises Logcrypt, a logging protocol that provides strong cryptographic assurances that data stored by a collector cannot “be modified after the compromise without detection” [20]. The most significant improvement is the ability to use a public key cryptography in computing the field Z . Using the symmetric techniques proposed by [38], any entity \mathcal{V} who wishes to verify the a log must possess the secret A used as parameter of the MAC function. This secret also gives the entity the ability to falsify the log entries, which would render evidence gathered this way utterly wrong. Public key – specifically, Identity-Based Encryption, see Section 4-4 – allows signatures to be created with one key and verified with another. Thus, the major aspect of Logcrypt is the replacement of MAC with signatures, in a way similar to Accorsi [2].

Other improvements described in [20] include a method of securing multiple, concurrently maintained log files using a single initial value, and a method of aggregating multiple log entries to reduce latency and computational load. Holt provides algorithms to this end and reports on performance analysis obtained with a prototypical implementation.

4.3. The Ma and Tsudik class

Ma and Tsudik [30] propose a logging protocol for the storage phase that employs a novel authentication technique called “Forward-Secure Sequential Aggregate” (FssAgg) [29].

In an FssAgg, forward-security, append only signatures generated by the same signer are sequentially combined into a single aggregate signature. Successful verification of an aggregate signature is equivalent to that of each component signature, whereas failed verification of an aggregate signature implies that at least one component signature is invalid.

Carried over to the logging protocol, these properties allow Ma and Tsudik to build a secure logging protocol that provides *forward security* as Schneier and Kelsey [38], *stream security*, so that audit trail reordering is impossible, and *integrity*, i.e. insertion of new entries as well as modification and deletion of existing entries renders the final aggregated invalid.

At a given timepoint an authenticated log file consists of two parts: audit trail $[L_1, \dots, L_{i-1}]$ and two FssAgg authentication MACs $\mu_{\mathcal{T},i-1}$ and $\mu_{\mathcal{V},i-1}$. In addition to that, the collector possesses two random symmetric keys A_i and B_i , which are created during the initialization phase and updated – by means of a cryptographic hash function – for each new entry. The initial keys A_1 and B_1 are communicated during the initialization phase to a semi-trusted verifier \mathcal{V} and the trusted verifier \mathcal{T} , respectively.

When the i th event occurs, the collector \mathcal{U} creates and appends the entry L_i to the audit trail, updating the authentication tags $\mu_{\mathcal{V}}$ and $\mu_{\mathcal{T}}$ as follows:

- 1) \mathcal{U} generates a MAC for the verifier \mathcal{V} $mac_{A_i}(L_i)$. It computes $\mu_{\mathcal{V},i}$ such that $\mu_{\mathcal{V},i} = \mathcal{H}(\mu_{\mathcal{V},i-1} || mac_{A_i}(L_i))$, where \mathcal{H} is an one-way, collision-resistant hash function acting as an aggregator.
- 2) \mathcal{U} updates the second FssAgg MAC (for \mathcal{T}) in the same way using the value B_i : $\mu_{\mathcal{T},i} = \mathcal{H}(\mu_{\mathcal{T},i-1} || mac_{B_i}(L_i))$.
- 3) \mathcal{U} evolves both keys A_i and B_i , applying a hash function to obtain the values A_{i+1} and B_{i+1} . \mathcal{U} irretrievably deletes the prior keys.

The authentication tag $\mu_{\mathcal{V}}$ can be “unrolled” as:

$$\mu_{\mathcal{V},i} = \mathcal{H}(\mathcal{H}(\dots \mathcal{H}(\mu_{\mathcal{V},i} || mac_{A_1}(L_1)) || \dots || mac_{A_i}(L_i))),$$

where the same holds true for $\mu_{\mathcal{T},i}$ with the symmetric key B_i . That is, the final aggregator signatures are a function of *all* the previous aggregators, thereby ensuring unforgeability.

Besides presenting the design of the logging protocol, the authors also report on its implementation, analyzing the performance of the implementation in different settings.

4.4. Encrypted search class

Some settings assume a collector that records the entries in clear-text and, eventually, stores the resultant audit trail (or backups thereof) in marginally trusted sites. This makes it necessary

$$R_i := \boxed{E_{K_i}(m_i) \quad H(R_{i-1}) \quad c_{w_a}, c_{w_b}, c_{w_c}}$$

Figure 4. Format of a log entry in Waters et al. [43]

to encrypt the *whole* audit trail. The downside of this measure is that encryption makes it extremely difficult to search for and extract entries from the audit trail, which is essential for digital evidence generation. In fact, a naïve approach to do so would require decrypting every entry separately. This brings along a number of disadvantages, such as unintended access to and possible disclosure of classified data, and huge computational cost associated with entry description.

This class of logging protocols is concerned with the secure storage of encrypted audit trails and efficient retrieval of entries from these trails. (See [13] for techniques tailored for encrypted search only.) We classify the existing protocols into those based on *Identity-Based Encryption* (IBE), i.e. a special kind of PKI where *any* string can be used as a public key [9], and those approaches based on *Bloom Filters*, i.e. probabilistic data structures to test whether an element is a member of a set [8].¹

The approach of Waters et al. They devise protocols for the search of encrypted audit trails using IBE [43]. In IBE, public keys can be arbitrary strings; private keys are derived from public keys through a system-wide master secret known (only!) to a trusted authority. In Waters et al., a designated trusted party \mathcal{T} – the so-called *key escrow agent* – constructs keyword search capabilities which allows (less trusted) parties \mathcal{V} – the so-called *investigators* – to search for and decrypt entries matching a given keyword. The escrow agent can either carry out the retrieval himself or delegate this task to the investigators.

The server generating the audit trail encrypts the entries with the public keys corresponding to the keywords that are derived from those entries. The key escrow agent holding the IBE master secret can construct a search capability for a given keyword as the private key corresponding to the given keyword, whereas additional security properties of the Boneh and Franklin’s scheme hold.²

The audit trail consists of a set of entries $\{R_0, \dots, R_n\}$ built up as depicted in Fig 4: $E_{K_i}(m_i)$ stands for the encryption of the i th event m under a key K_i . $H(R_{i-1})$ is the hash of the previous entry, thereby generating a hash chain. $c_{w_a}, c_{w_b}, c_{w_c}$ represent the information about the keywords w_a, w_b, w_c that index the contents m_i of the entry R_i .

The relation between the keywords w_a, \dots, w_c and the encryption key K_i is of foremost importance for entry retrieval. To construct a searchable audit trail, the key escrow server first extracts keywords of the entry. Next, it encrypts the entry m_i using the key K_i derived from the keywords, so that for each keyword w , the server computes the IBE c_i of the string (**flag**| K) using the w as the public key the IBE setup P known by the server, where **flag** is a constant

¹While false positives are possible in Bloom Filters, false negatives are not.

²Attackers cannot tell which public key was used to create a cipher-text.

bit-string with length l .

Retrieval consists of the search and the decryption of entries. Assuming a search for the keyword W , the escrow agent constructs the corresponding capability d_w as the IBE private key for the string w . For each log entry and for each c the investigator attempts to IBE-decrypt c using the private key d_w . If the prefix of the result matches **flag**, then the investigator extracts K as the remainder of the result. If none of the results contain **flag**, the log entry does not match, moving to the next entry. Whenever one of the results match, the capability holder computes K to decrypt $E_K(m)$. Here, the cryptographic properties of the underlying IBE scheme prevent some investigator with capability d_w from deriving a capability $d_{w'}$ to access entries containing the keyword w' .

Ohtaki's approach

Similar to Waters et al. [43], Ohtaki proposes an approach based on Bloom Filters for the partial disclosure of audit trails [33]. However, in contrast to [43], Ohtaki allows for Boolean searches consisting of AND and OR operators.

The use of Bloom filters is for efficiency: they reduce the storage space for the encoded index. A Bloom filter is a probabilistic data structure to test whether an element is a member of a set [8]. Filters are represented by an array of m bits. There are also r independent hash functions h_1, \dots, h_r that take all inputs and produce outputs that are distributed uniformly over a range $[1 \dots m]$. An arbitrary set $S = \{e_1, \dots, e_n\}$ is represented with the m -bit array as follows:

- *Filter construction:* Initially, every array position is set to 0. To add an element e , independent hash functions of the element are calculated and array bits at position $h_1(e), \dots, h_r(e)$ are set to 1. This procedure is repeated to each $e_n \in S$.
- *Deciding set membership:* To decide whether $x \in S$, the bits at positions $h_1(x), \dots, h_r(x)$ are checked. If *any* selected bit is 0, x is *not* a member of S ; if all the checked bits are 1, then x is a member of the set S . There is some probability of false positives: x appears to be in S but is actually not a member of it. This happens when all bits corresponding to x were set by other elements in S . The false positive rate can be regulated through careful selection of h , r and m .

The architecture proposed by Ohtaki foresees principals taking one of the following roles: the *administrator* manages the access to the log file and the *searcher* (or *auditor*) retrieves log entries for a keyword. Assuming a PKI, the administrator generates a pair of keys K_P and K_S , publishes the public key K_P , and sets the parameters h , r and m .

For every incoming log entry L^i with unique record identifier ID^i , the administrator generates two kinds of data: first the encrypted data V^i for partial disclosure; second, the Bloom filter BF^i used for the search and retrieval of entries. In detail, the encryption key K^i used to encrypt L^i is different for each entry, so that entry-wise decryption is possible. To this end, the symmetric key K^i is derived from the record identifier ID^i and the administrator's private key K^i , for each entry L .³

³Precisely, for some entry L^i , the corresponding key is obtained as follows: $K^i = K_S(ID^i)$, where K_S stands for the public key encryption with K_S . The cipher-text V^i is generated by encrypting L^i with K^i .

For the Bloom filter generation, the data used for matching is built by inserting *all* possible combinations of keywords and Boolean operators AND and OR, whereas a normalization step removes equivalent index patterns. Each normalized pattern w is converted into a “coded word” T_w , which is inserted to the Bloom Filter BF^i after the concatenation with the corresponding record identifier ID^i , denoted $ID^i||T_w$.

To obtain the contents of entries, the auditor receives the whole encrypted audit trail and communicates the keywords of interest to the administrator. Let W denote the normalized keyword pattern; the administrator converts W into the a coded keyword T_W , using to this end his private key K_S . With the encoded keyword T_W and the encrypted audit trail, the auditor queries the audit trail for matching entries. Here, the Bloom filter of each entry is verified against the concatenation of record identifier and coded keyword. That is, the auditor computes r hash functions of $ID^i||T_W$. If any of BF^i is 0, the entry does not match the search criterion. If all the bits are 1, the entry matches the keyword and, for the set of entries fulfilling the search criterion, the administrator must generate the keys allowing the auditor to decrypt them.

Ohtaki analyzes the probability of false positives and its relationship to the parameters of the Bloom filter, demonstrating that the probability of false positives already in simple Bloom filters amounts to less than 0,01%

5. What logging protocols offer

Table 1 relates the security requirements for digital evidence of Section 2 and the logging protocols presented in Section 4. Each of the protocols is deployed for a particular trust model. The extensions of the Syslog, for example, assume a trusted collector, whereas the transportation medium is malicious. Hence, digital evidence generated by, e.g. syslog-ng, can be challenged on the basis that it does not protect stored log data. From this perspective, except for the **BBox**, all other protocols fail to address the whole spectrum of security requirements.

Problematic is the fact that even focusing on the strengths of each class, one finds vulnerabilities that lead to the violation of security requirements. The truncation and the replacement attacks upon the Schneier and Kelsey [38] protocol exemplify how subtle the design of logging protocols can be. Similarly, the **BBox** architecture [2] – omitted due to space constraints – builds on a number of assumptions that rely on strong tamper resistance on the side of the devices.

Secure logging protocols in the Encrypted Search class do not address the information flows or entropy that might happen when an external auditor is allowed to search through the whole encrypted file for particular keywords. Moreover, in possession of the cipher-text, an attacker can infer other keywords and conduct chosen cipher-text attacks, thereby compromising the confidentiality of other entries.

Finally, the Ma/Tsudik class is based aggregated signatures for which correctness proofs are still missing. Experience with hash chains in [38] and more generally in the development distributed protocols (see [11]) shows that such subtle mistakes are rule rather than the exception.

Taking stock, existing logging protocols enforce a number of essential security requirements. Still, none of the protocols analyzed above convincingly satisfy the evidentiary standards characterizing the admissibility requirements proposed in Section 2.

Table 1
Secure logging protocols and fulfilled security requirements

Secure logging protocol	Security Requirements							
	Transmission phase					Storage phase		
	confidentiality	or. authentication	integrity	uniqueness	rel. delivery	accountability	integrity	confidentiality
syslog	no	no	no	no	no	no	no	no
syslog-ng	yes	no	yes	no	yes	no	no	no
syslog-sign	no	yes	yes	yes	no	no	no	no
reliable syslog	yes	yes	yes	yes	yes	no	no	no
Schneier/Kelsey	no	no	no	no	no	yes	no	yes
Stathopoulos et al.	no	no	no	no	no	no	no	yes
BBox	yes	yes	yes	yes	yes	yes	yes	yes
Logcrypt	no	no	no	no	no	yes	yes	yes
Waters et al.	no	no	no	no	no	no	yes	yes
Ohtaki	no	no	no	no	no	yes	yes	yes
Ma/Tsudik	no	no	no	no	no	yes	yes	yes

6. Open research issues

Apart from the legal issues concerning the use of log data as digital evidence in the court, we see various technical research issues towards the provision of digital evidence, e.g.:

Obtaining reliable signals: Secure logging protocols assure that events, once sensed by the devices, are transmitted to and stored in a way that they remain authentic. These protocols neglect the *veracity* of events though, so that devices may be tricked into communicating wrong events (or signals) to the collector. Clearly, evidence generated on this basis is bogus. While devising techniques for reliable signalling is a challenge that arises in several settings, such as usage control [35] and DRM [16], convincing solutions are still not available.

Advanced adversary models: There is a visible gap between the secure logging protocols used to generate evidence and the attacker models they aim to protect against. In general, such a gap leads to threat analysis rendering (logging) protocols as “secure”, whereas they in fact allow manipulation that cannot be detected by verifiers [19]. Hence, the development and corresponding evolution of adequate, commonly (i.e. both technically and legally) accepted attacker models is of foremost relevance in the realm of admissible digital evidence.

Formal verification of logging protocols: In contrast to authentication and access control protocols, which have been thoroughly analyzed and proved secure, formal approaches to asserting the security properties of logging protocols do not exist. Techniques to do so can focus both on the communication steps between the principals (see, e.g., Drieslma et al. [15]) and on the cryptographic primitives (see, e.g., Backes et al. [6]), building on existing approaches in these settings.

Standard evidence formats: Once obtained from audit trails, there is to-date no formal format for digital evidence or an *evidence bag* for digital evidence. In consequence, concepts, such as chain of custody and immutability [39], cannot be taken for granted. While work has been done towards digital evidence bags, e.g. [21] and [42], we are not aware of standards for it.

Consolidation of log data: While digital evidence may be found in a single audit trail, in practice this is the exception rather than the rule. That is, auditors must consolidate data from remote, disparate sources possibly collected using different formats. This boils down to the problem of event correlation, which is known to be as challenging as relevant for sound digital evidence [17]. As in the case of formal verification, there is lack of work towards policy or suspicion-driven correlation techniques.

Evidence mining: Automated approaches to evidence mining focus on the correlation of events in the log file to decide whether security policies have been violated during the execution. These approaches are not tailored to detect subtle, illegal information flows happening between the containers, such as covert/side-channels [27]. Dedicated correlation techniques to this end are strongly needed, yet not available [4], [44].

7. Conclusion

While log data are being increasingly used as digital evidence in judicial disputes, we have shown that protocols used to collect and store log data fail to satisfy some of the security

and hence legal requirements for admissible evidence. In consequence, more often than not, digital evidence based on log data can be successfully challenged in the court, leading to inadmissibility or loss on probative force. Intensive research into logging protocols is needed to advance digital evidence, in particular into the elucidation of further, fine-grained security requirements and into the verification of protocols with regard to these requirements. As further work, we intend to address these issues and investigate the trustworthiness of log data and the intricacies delegation causes to secure logging protocols, two relevant topics neglected in our analysis due to space constraints. Moreover, in the course of our analysis realized that a finer-grained view of admissibility – systematically characterizing what is admissibility across different legislations – is needed to analyze protocols and improve the expressiveness of our classification. This is also subject of further work.

Acknowledgements

I would like to thank Dr. Erin Kenneally and the anonymous referees for the helpful comments on the previous versions of this paper, and Thomas Stocker for helping me in the typography.

References

- [1] R. Accorsi, “On the relationship of privacy and secure remote logging in dynamic systems,” in *Security and Privacy in Dynamic Environments*, ser. IFIP Conference Proceedings, S. Fischer-Hübner, K. Rannenberg, L. Yngström, and S. Lindskog, Eds. Springer-Verlag, 2006, vol. 201, pp. 329–339.
- [2] —, “Automated counterexample-driven audits of authentic system records,” Ph.D. dissertation, University of Freiburg, 2008.
- [3] R. Accorsi and A. Hohl, “Delegating secure logging in pervasive computing systems,” in *Security in Pervasive Computing*, ser. Lecture Notes in Computer Science, J. Clark, R. Paige, F. Pollack, and P. Brooke, Eds. Springer-Verlag, 2006, vol. 3934, pp. 58–72.
- [4] R. Accorsi and C. Wonnemann, “Detective information flow analysis for business processes (extended abstract),” in *Business Processes, Services Computing and Intelligent Service Management*, ser. Lecture Notes in Informatics, W. Abramowicz, L. Macaszek, R. Kowalczyk, and A. Speck, Eds. Bonner Köllen Verlag, 2009, vol. 147, pp. 223–224.
- [5] M. Alles, A. Kogan, and M. Vasarhelyi, “Black box logging and tertiary monitoring of continuous assurance systems,” *Information Systems Control Journal*, vol. 1, 2003.
- [6] M. Backes, M. Berg, and D. Unruh, “A formal language for cryptographic pseudocode,” in *Logic for Programming, Artificial Intelligence, and Reasoning*, ser. Lecture Notes in Computer Science, 2008, vol. 5330, pp. 353–376.
- [7] M. Bellare and B. Yee, “Forward integrity for secure audit logs,” University of California, San Diego, Dept. of Computer Science & Engineering, Tech. Rep., 1997.
- [8] B. Bloom, “Space/time trade-offs in hash coding with allowable errors,” *Comm. ACM*, vol. 13, no. 7, pp. 422–426, 1970.

- [9] D. Boneh and M. Franklin, "Identity based encryption from the Weil pairing," *SIAM Journal of Computing*, vol. 32, no. 3, pp. 586–615, 2003.
- [10] C. Chong, Z. Peng, and P. Hartel, "Secure audit logging with tamper-resistant hardware," in *Security and Privacy in the Age of Uncertainty*, ser. IFIP Conference Proceedings, D. Gritzalis, S. D. C. di Vimercati, P. Samarati, and S. Katsikas, Eds., vol. 250. Kluwer, 2003, pp. 73–84.
- [11] J. Clark and J. Jacob, "A survey of authentication protocol literature: Version 1.0," 1997, available at <http://www.cs.york.ac.uk/~jac/papers/drareview.ps.gz>.
- [12] F. Cohen, *Challenges to Digital Evidence*. ASP Press, 2008.
- [13] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, "Searchable symmetric encryption: Improved definitions and efficient constructions," in *ACM CCS*, A. Juels, R. Wright, and S. D. C. di Vimercati, Eds. ACM, 2006, pp. 79–88.
- [14] D. Dolev and A. Yao, "On the security of public key protocols," *IEEE Transactions on Information Theory*, vol. 2, no. 29, pp. 198–208, 1983.
- [15] P. H. Drielsma, S. Mödersheim, L. Viganò, and D. A. Basin, "Formalizing and analyzing sender invariance," in *Workshop on Formal Aspects in Security and Trust*, ser. Lecture Notes in Computer Science, T. Dimitrakos, F. Martinelli, P. Ryan, and S. Schneider, Eds., 2006, vol. 4691, pp. 80–95.
- [16] J. Erickson, "Fair use, DRM, and trusted computing," *Comm. ACM*, vol. 46, no. 4, pp. 34–39, 2003.
- [17] K. Franke and S. Srihari, "Computational forensics: An overview," in *2nd Workshop on Computational Forensics*, ser. Lecture Notes in Computer Science, K. Franke and S. Srihari, Eds. Springer-Verlag, 2008, vol. 5158, pp. 1–10.
- [18] M. Franklin, "A survey of key evolving cryptosystems," *International Journal of Security and Networks*, vol. 1, no. 1/2, pp. 46–53, 2006.
- [19] V. Gliglor, "On the evolution of adversary models in security protocols," in *Workshop on Security Protocols*, ser. Lecture Notes in Computer Science, B. Christianson, B. Crispo, J. Malcom, and M. Roe, Eds. Springer-Verlag, 2007, vol. 4631, pp. 276–283.
- [20] J. Holt, "Logcrypt: Forward security and public verification for secure audit logs," in *Australasian Symposium on Grid Computing and e-Research*, ser. CRIPT, R. Buyya, T. Ma, R. Safavi-Naini, C. Steketee, and W. Susilo, Eds., vol. 54. Australian Computer Society, 2006, pp. 203–211.
- [21] C. Hosmer, "Digital evidence bag," *Comm. ACM*, vol. 49, no. 2, pp. 69–70, 2006.
- [22] Kahn Consulting, "Computer security log files as evidence," http://www.kahnconsultinginc.com/images/pdfs/KCI_ArcSight_ESM_Evaluation.pdf, August 2006.
- [23] J. Kelsey and J. Callas, "Signed syslog messages," IETF Internet Draft, 2005, <http://www.ietf.org/internet-drafts/draft-ietf-syslog-sign-16.txt>.

- [24] E. Kenneally, "Digital logs – Proof matters," *Digital Investigation*, vol. 1, no. 2, pp. 94–101, 2004.
- [25] K. Kent and M. Souppaya, *Guide to Computer Security Log Management*, National Institute of Standards and Technology (NIST), September 2006.
- [26] L. Lamport, "Password authentication with insecure communication," *Comm. ACM*, vol. 24, no. 11, pp. 770–772, 1981.
- [27] B. Lampson, "A note on the confinement problem," *Comm. ACM*, vol. 16, no. 10, pp. 613–615, October 1973.
- [28] C. Lonvick, "RFC 3164: The BSD syslog protocol," Request for Comments, 2001, <http://www.ietf.org/rfc/rfc3164.txt>.
- [29] D. Ma, "Practical forward secure sequential aggregate signatures," in *Proceedings of the ACM Symposium on Information, Computer and Communications Security*, M. Abe and V. D. Gligor, Eds. ACM, 2008, pp. 341–352.
- [30] D. Ma and G. Tsudik, "A new approach to secure logging," *ACM Trans. on Storage*, vol. 5, no. 1, pp. 1–21, 2009.
- [31] U. Maurer, "New approaches to digital evidence," *Proceedings of the IEEE*, vol. 92, no. 6, pp. 933–947, June 2004.
- [32] D. New and M. Rose, "RFC 3195: Reliable delivery for syslog," Request for Comments, 2001, <http://www.ietf.org/rfc/rfc3195.txt>.
- [33] Y. Ohtaki, "Partial disclosure of searchable encrypted data with support for boolean queries," in *Workshop on Advances in Policy Enforcement*, S. Jakoubi, S. Tjoa, and E. Weippl, Eds. IEEE, 2008, pp. 1083–1090.
- [34] R. Oppliger and R. Ritz, "Digital evidence: Dream and reality," *IEEE Security & Privacy*, vol. 1, no. 5, pp. 44–48, 2003.
- [35] A. Pretschner, M. Hilty, C. Schaefer, F. Schütz, and T. Walter, "Usage control enforcement: Present and future," *IEEE Security & Privacy*, vol. 6, no. 4, pp. 44–53, 2008.
- [36] "Reliable syslog," <http://security.sdsc.edu/software/sdsc-syslog/>.
- [37] S. Sackmann, J. Strüker, and R. Accorsi, "Personalization in privacy-aware highly dynamic systems," *Comm. ACM*, vol. 49, no. 9, pp. 32–38, 2006.
- [38] B. Schneier and J. Kelsey, "Security audit logs to support computer forensics," *ACM TISSEC*, vol. 2, no. 2, pp. 159–176, 1999.
- [39] F. Smith and E. Kenneally, "Electronic evidence and digital forensics testimony in court," in *Handbook of Digital and Multimedia Forensic Evidence*, J. Barbara, Ed. Humana Press Inc., 2008, pp. 103–132.

- [40] V. Stathopoulos, P. Kotzanikolaou, and E. Magkos, "A framework for secure and verifiable logging in public communication networks," in *Critical Information Infrastructures Security*, ser. Lecture Notes in Computer Science, J. Lopez, Ed. Springer-Verlag, 2006, vol. 4347, pp. 273–284.
- [41] "Syslog-ng web site," http://www.balabit.com/products/syslog_ng/.
- [42] The Common Digital Evidence Storage Format Working Group, "Standardizing digital evidence storage," *Comm. ACM*, vol. 49, no. 2, pp. 67–68, 2006.
- [43] B. Waters, D. Balfanz, G. Durfee, and D. Smetters, "Building an encrypted and searchable audit log," in *Network and Distributed System Security Symposium*, 2004.
- [44] C. Wonnemann, R. Accorsi, and G. Müller, "On Information Flow Forensics in Business Application Scenarios," in *Workshop on Security, Trust, and Privacy Business Application Scenarios*, IEEE, 2009.