

Security of Distributed Digital Criminal Dossiers

Martijn Warnier¹, Frances Brazier¹ and Anja Oskamp²

¹Intelligent Interactive Distributed Systems, Faculty of Sciences, VU University Amsterdam, the Netherlands

²Computer Law Institute, Faculty of Law, VU University Amsterdam, the Netherlands

Email: {warnier, frances}@cs.vu.nl a.oskamp@rechten.vu.nl

Abstract—Securely managing shared information in distributed environments across multiple organisations is a challenge. Distributed information management systems must be able to support individual organisations’ information policies whilst securing global consistency and completeness. This paper proposes a multi-agent approach to a distributed multi-organisational system design based on this principle, focusing on the example of the distributed digital criminal dossier used in the Courts of Amsterdam and Rotterdam, compiled and managed by the Public Prosecution. Security requirements are identified and a distributed multi-agent architecture proposed.

Index Terms—Security, Distributed Systems, Digital Criminal Dossiers, Legal Domain

I. INTRODUCTION

Managing shared information securely and efficiently across multiple independent organizations is a challenge - the challenge this paper addresses. In general, individual organizations manage their own information locally on their own systems, according to their own information policies. If information is to be shared, however, additional policies are needed to manage global correctness and consistency.

One example of an environment in which information needs to be shared between multiple semi-independent organisations is the environment in which the Public Prosecution compiles and manages distributed digital criminal dossiers. Such dossiers are currently being used in the Courts of Amsterdam and Rotterdam in a number of pilot studies. Earlier work [1] explored the potential of a centralised system for digital criminal dossier management. This paper explores the potential of a *distributed digital criminal dossiers*, supported by a multi-agent system architecture [2], to improve consistency, completeness, integrity and security of the information in such dossiers.

The distributed architecture allows *physically* distributed information sources, such as Municipals and the prison systems, to remain responsible for the integrity of their own information content, each monitored by one or more of their own software agents. The Public Prosecutor

This paper updates and extends “Secure Distributed Dossier Management in the Legal Domain,” by M. Warnier, F.M.T. Brazier, M. Apistola, and A. Oskamp, which appeared in the Proceedings of The Second International Conference on Availability, Reliability and Security (ARES’07), Vienna, Austria, April 2007. © 2007 IEEE.

This work has been funded by the NWO TOKEN program, grant number 634.000.431.

has a centralized role and is responsible for providing the infrastructure that enables other organizations to securely add information and securely access information in criminal dossiers. Together these organizations form a semi-open environment: an environment in which organizations have their own control over their own information. In this environment Dutch Law, however dictates the exchange of this information with other organizations. This paper discusses some of the details involved in the use of such dossiers, focusing on security issues. See [3], [4] for more details on enforcing consistency and completeness and on implementation details.

All legal and procedural details discussed in this paper are interpreted in the context of Dutch law, but can be extended to other legislation.

The paper first explores security requirements in the semi-open distributed environment associated with the compilation of criminal dossiers. Section III introduces distributed digital dossiers, Section IV discusses the associated security architecture in light of the security requirements. Domain specific legal requirements are discussed in Section V and the paper ends with a discussion and conclusions.

II. SECURITY REQUIREMENTS

Security is essential in the environment of distributed digital criminal dossiers addressed in this paper. Security requirements (1) hold for all comparable distributed computer systems, but also requirements (2) that hold in this specific context. Nine requirements, related directly to relatively standard security requirements for distributed clinical information systems [5], have been identified:

- A. *Access Control*: each individual dossier and each referring record (see Section IV) contained therein must be marked with an access control list naming the people or groups of people who may read and alter data. The system must prevent anyone not on the access control list from accessing the dossier in any way.
- B. *Dossier creation*: a dossier is always created by the Public Prosecutor.
- C. *Control*: separate records in the dossier are the responsibility of individuals/organizations on the access control list. This control can be transferred to other persons/organizations when required.

- D. *Notification*: defendants must be informed of the content of the dossier as required by Dutch law. During the trial period, i.e., when a dossier is finalized and sent to the Court, all parties (public prosecution, lawyers and court) need to be informed of updates to the dossier. Defendants have the right to challenge the correctness of the information contained in the dossier during trial.
- E. *Persistence*: during the trial period, no party may have the ability to delete (parts of) the dossier, unless this is mandated by the Dutch law because the time for enforcement has passed (extinguishment).
- F. *Attribution*: all changes to (records of) the digital dossier must be marked with the subject's (users/organizations) identity as well as date and time. An audit trail must also be kept of all deletions¹.
- G. *Information Flow*: no information may be copied from one dossier to another unless this is allowed by the access control policy.
- H. *Aggregation control*: aggregation of information is only possible for persons/organizations (such as the Public Prosecution) who have the explicit right to do so.
- I. *Trusted computing base*: computer systems that handle digital dossiers need maintained in a trusted environment, handled by trusted system administrators.

In addition, in the specific context of the digital dossier the following additional requirements hold:

- J. *Secure transfer*: the (physically) distributed organizations may only exchange information over secure communication channels that guarantee confidentiality and integrity of the transferred data.
- K. *Compartmentalization of information*: organizations may only access those parts of the dossier for which they are responsible and/or those parts where the need to access has been identified.
- L. *Consistency*: the data in the dossier must be (internally) consistent.
- M. *Completeness*: each dossier must be complete when it is sent to the Court. The Court forwards the dossier to the relevant judge(s) and lawyer(s).
- N. *Backups*: periodically backups of each dossier must be made. These backup copies are secured against unauthorized access in a similar fashion as the original dossiers.

Consistency and completeness are especially challenging requirements: they must be guaranteed. When an authorized organization, e.g. the Council for Child Welfare, adds a record to a specific digital dossier, the system needs to check whether the information is *consistent* with all other information in the dossier, for example whether personal information, such as name, address, age and sex of a subject, is consistent across records/documents.

¹For our purpose attribution and auditability can be regarded as similar requirements.

Completeness requirements include general completeness requirements, and offense specific completeness requirements. General completeness requirements specify that, for example, a subject's personal administrative information, the offense for which he/she is charged, and the official report filed by the Police, must be included *in each distributed digital criminal dossier*. In addition, offense specific completeness requirements hold, for example, if the offense for which a subject is being charged is a drunken driving charge then an alcohol test by an authorized lab must be included in the dossier. Completeness of the dossier must be guaranteed before it is transferred to the Court.

The next sections propose a design for an agent based support system for the distributed digital criminal dossiers that fulfills the above mentioned requirements.

III. DISTRIBUTED DIGITAL DOSSIERS

The intrinsic nature of criminal dossiers with physically distributed sources of information distributed over different organizations is the motivation for proposing a *distributed digital criminal dossier*. Figure 1 below shows some of these (distributed) sources².

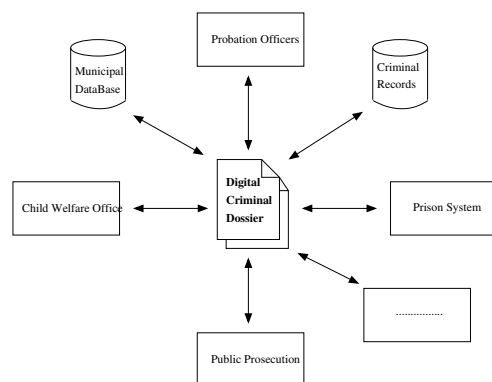


Figure 1. Some of the Information Sources used to compile a Digital Criminal Dossier

This section describes the design, organization, abstract implementation, life cycles models of distributed criminal dossiers and how such dossiers can be managed by agents.

An example based on juvenile repeat offenders [3] is used in this paper to illustrate the types of information included in a distributed digital criminal dossier and their sources, the focus is on the information exchange between Public Prosecution and the Council for Child Welfare³. This scenario has been chosen because repeat

²In this context, the Police has a special kind of role. It delivers information needed for compiling dossiers, but that does not, for example, changes information in dossiers created by the Public Prosecution, as other organizations such as the Municipals, will do in our approach. Information exchange between Police and other organizations does thus not occur via distributed digital criminal dossiers.

³In the Dutch context the Council for Child Welfare has the task to investigate all crimes of minors. In addition to the criminal offences of the minor, the family situation and other relevant social factors are taken into account. This results in a motivated advice for suitable punishment (if applicable) of the suspect.

offenders and especially juvenile offenders represent an interesting and socially important subject. It also clearly illustrates how multiple physically distributed organizations can work together in an efficient way in the semi-open environment responsible for criminal dossiers.

A. Design and Organization

Each individual digital dossier is -by Dutch law- created by the Public Prosecution once it decides, on the basis of information available, to prosecute a defendant. A newly created dossier consists of records and meta data. The meta data contains information such as the access control list for the dossier, logging information on who altered or accessed information at what time, and when the last backup of the dossier was made. This meta data part of the dossier is stored centrally by the Public Prosecution.

Individual records are the responsibility of different organizations. Administrative information, for example, is managed and maintained by the defendant's local authorities⁴ and information on a juvenile's family situation is provided by the Council for Child Welfare. Distributing both the *data* (information) and the *responsibility* for the data ensures that information in the digital dossier is kept as up-to-date as possible. Changes in data are flagged by the relevant organizations and transmitted to the Public Prosecution for synchronization of the complete (distributed) dossier.

Thus the basic dossier itself is stored by the Public Prosecution while relevant records are maintained and stored by the responsible organizations and then synchronized with the digital dossier by the Public Prosecution.

B. Dossier Implementation

Dossiers are implemented with a light weight 'skeleton' framework based on XML. Fields in this XML document contain keywords and references to other (possibly, but not necessarily, XML-based) documents. The distributed nature of the dossier emerges through the references which can be both to local and remote systems.

The central part of the dossier, as created by the Public Prosecution, specifies which information is to be included, such as: status (mandatory, optional), dependencies, data maintainer (which organization), access rights, etc. This meta data depends on the charged crime. This paper assumes that standard XML templates exist for each type of crime, and that these templates are used by the Public Prosecution to structure the meta-data in a dossier, see [3] for a proposal to use automatic clustering techniques to construct (part of) this information automatically. Figure 2 depicts an example of part of a dossier with both meta-data and data provided by the Public Prosecution [4]. It is based on a real dossier from an actual case, but all

⁴In the Dutch context Municipals are responsible for keeping administrative records of its citizens. Other countries use other organizational units to maintain these records, or non at all, as is, for example, the case in the United States.

data has been anonymized [3]. The case of a juvenile repeat offender has been chosen to illustrate the use of distributed digital dossiers, as a number of physically distributed organizations are involved.

Meta-data includes the dossier number, creation date, type of offense and access control lists. The example depicted in Figure 2 specifies that in this case the dossier has been opened by the Public Prosecution of Amsterdam on the first of July 2007. The offense for which the subject is being charged is shoplifting, and the subject is a (known) juvenile repeat offender. This dossier further specifies that four named employees of the Public Prosecution have been assigned permission to read this document (identified in this example by numbers 1234, 2345, 3456, 4567) and that only two of these four employees have permission to edit parts of the dossier, i.e., write permission, see Section IV for more details on the access control mechanism and other security issues.

```

<Dossier>
  <MetaData>
    <Ref>PubProsAmsterdam-00001</Ref>
    <CreationDate>1-7-2007</CreationDate>
    <Offense>
      <Main>Shoplifting</Main>
      <Category>JuvenileRepeatOffender</Category>
    </Offense>
    <Access>
      <Read>1234,2345,3456,4567</Read>
      <Write>1234,4567</Write>
    </Access>
    ...
  </MetaData>
  <Records>
    <MandatoryInfo>
      <PersonalInfo>ref:MuniDatAmsterdam-123456
    </PersonalInfo>
      <ReportCCW>ref:CCW-234567</ReportCCW>
      ...
    </MandatoryInfo>
    <OptionalInfo>
      <StatementSubject>ref:PP-00001statement.pdf
    </StatementSubject>
      <Media>ref:PP-00001-movie1.avi,
              ref:PP-00001-movie2.avi
    </Media>
      ...
    </OptionalInfo>
  </Records>
</Dossier>

```

Figure 2. Example of the centralized 'skeleton' Digital Criminal Dossier [4]

The template distinguishes between mandatory and optional information for a specific crime. Such information represents the knowledge used by the Public Prosecution to check completeness of a dossier - and provides the structure needed for automated completeness checks [3].

Certain information must always be included in a dossier, such as the subject's administrative data, and the original police report concerning an incident. Other information is mandatory for a specific kind of offense and/or type of offender. In this particular example, independent of the crime for which the subject has been

charged, the subject is a juvenile: according to Dutch law, a report by the Council for Child Welfare is mandatory. Other information, however, is optional. In this example, footage of the incident on the basis of which the subject has been charged, recorded by a video surveillance camera, is included. Note that a court case can only commence if all mandatory information is included in a dossier. The dossier also contains a number of references, indicated by the `ref:` keyword, to specific documents and their source. In this example references local to the Public Prosecution (PP) are the subject's original statement (`ref:PP-00001statement.pdf`) and the video evidence (`ref:PP-00001-movie1.avi` and `ref:PP-00001-movie2.avi`). References to XML documents to be provided by other organizations, in this example, include the (national) Council for Child Welfare (`ref:CCW-234567`) and the Municipal database of Amsterdam (`ref:MuniDatAmsterdam-123456`).

A distributed digital criminal dossier thus consists of a number of documents in a networked structure distributed over physically distributed locations with the Public Prosecution as the central coordinator. Note that digital criminal dossiers always form the root of this network, i.e., digital criminal dossiers are not included in (referenced from) other dossiers, unless the referring dossier is also a digital criminal dossier. This ensures that digital criminal dossiers themselves are always controlled by the Public Prosecution, an important security requirement.

C. Life Cycle model

When the Public Prosecutor decides that a case is ready for Court the dossier is 'frozen': the dossier is finalized and forwarded by the Court to the presiding judge and the defendant's lawyer. From this point on the criminal dossier is no longer distributed and other organizations are no longer responsible for 'their' records. Note, that consequently trials are based on information available at this point in time.

As is currently the case, the Public Prosecutor decides which information is included in the frozen version of the dossier, and which not, based on its judgment of its relevance. Additional information can, from this moment on, only be added by one of the parties involved (prosecutors, judges and defense lawyers) by a special procedure that ensures that the relevant additions to the dossier are distributed to all concerned parties.

Once a case has been tried, a dossier can be 'defrosted', i.e., made distributed again, re-'frozen' when needed for a trial, etc. This process can be repeated numerous times (re-trials, appeals etc.) until a dossier is finally closed.

'Freezing' and 'defrosting' of distributed digital dossiers can be implemented in various ways. Two different life cycle models are discussed here: the straightforward (but inefficient) *naive* life cycle model and the efficient (though somewhat more involved) *semi-freezing* life cycle model.

- *The naive life cycle model*, is the conceptually most straightforward model. A dossier is 'frozen'

(static and centralized) and 'defrosted' (dynamic and distributed) as required. Note that a technical solution for defrosting a dossier is non-trivial, as it requires identification of the appropriate organization for each record in the dossier and complete new resynchronization of information contained in the dossier's records. Figure 3 illustrates the naive model graphically.




Figure 3. The naive life cycle model

A *distributed digital criminal dossier* is shown on the far left. This dossier is frozen, resulting in a static file with the current state of the data. Next, as indicated by the solid black arrows copies of this static dossier are distributed to the Public Prosecution, the Court and the defendant's lawyer. If a new trial is needed, e.g. due to a miss-trial or an appeal, the dossier is distributed again (defrosted), as shown on the far right of the picture. Each organization is again responsible for 'its' records in the dossier.

- *The semi-freezing life cycle model*, is a less drastic solution. Freezing entails making a local central copy of the dossier as in the above case. The difference is that the distributed version of the dossier still exists. If and when a dossier is defrosted, a new version is instantly available (again). This model is technically preferable, the only difficulty that can arise is that during trial, additional information may have been added to the dossier (by the Court or the defense). This information needs to be distributed to the relevant parties. The Public Prosecution is responsible for distributing this information to the responsible parties. The semi-freezing model is schematically displayed in Figure 4.




Figure 4. The semi-freezing life cycle model

As in the naive model, the *distributed digital crimi-*

nal dossier is shown on the far left. Once the dossier is frozen and copied (solid arrows), the dossier remains distributed (indicated by the dashed arrows). If and when a case is directed to a new Court, a completely new dossier can be acquired on the basis of the information as known to the distributed organizations responsible for the records. New information that surfaced during the trial, however, needs to be incorporated into the rest of the distributed dossier.

Both models have technical and conceptual advantages. As the semi-freezing model ensures a maximum of both control and responsibility for all parties involved this model is preferred. As such it is used in the remainder of this paper.

D. Managing Distributed Digital Dossiers with Agents

This section describes the functional design of a multi agent system for distributed digital dossier management. Distributed multi-agent systems provide a promising paradigm for large scale distributed autonomous systems [2]. Agents are pro-active, autonomous, possibly mobile systems capable of interacting with other agents and services, and adapting to their environment [6]. The main reasons for choosing the agent paradigm is that it provides a conceptually clear model for autonomous systems that supports modularity, security and scalability. Specific tasks can be implemented by dedicated agents, allowing for a clear separation of concerns and straightforward integration of new functionality as new agents, when needed.

From a technical perspective, one or more computer hosts that are maintained by the same organization together form a *location*. A dedicated middleware layer, the *agent platform* ensures that all hosts with a location can be viewed as one logical unit. The middleware ensures that all agents can uniquely be identified (using a lookup service), that agents on different locations can communicate with each other and that, if required, agents can migrate between locations⁵. Examples of such agent systems include AgentScape [7], JADE [8] and SeMoA [9].

Each organization has its own hosts with its own security policies. This allows *local control* and responsibility of data, and access to data with each organization, while at the same time it supports the use of *global security policies* that can guarantee a minimal set of (global) security requirements across organizations. Section IV describes several of the dedicated agents that are used to manage distributed digital dossiers.

Note that, in our model all interaction with the digital dossier is facilitated by means of agents. Agents are the single points of entry to dossiers. For example, only authorized agents can alter records in a dossier within a organization. The next section describes the specifics of the security architecture.

⁵Not all agent systems allow migration of agents.

IV. SECURITY ARCHITECTURE

The security requirements identified in Section II provide the outline for the presentation of the security architecture.

A. Access Control

The semi-open nature of the digital criminal dossier environment makes access control a particular challenge. Access control regulations in such systems generally tend to make *all* data more difficult to access, including the less sensitive information that can be of interest to a large public. This phenomenon is known as ‘label creep’ in the literature [10].

Our proposed solution handles this problem by means of a two-tier access model. On the first level role based access control [11] is used as an access control mechanism for access to the distributed digital criminal dossier. Each ‘role’, such as judge, lawyer, Public Prosecutor, or clerk, has certain rights with regard to the dossier. This typically depends on specific security policies, for example, a clerk may only add information to the dossier, not delete or modify anything, while a Public Prosecutor may change existing information in the dossier and even create new dossiers.

Additionally, the second access level uses access control lists [12] to limit the access of individuals *after* their role has been determined. Each dossier contains, in its meta data, information on specific individuals: it specifies per individual who may change, read, delete or add information to a dossier. For example, judge A may have permission to read a specific dossier (as it’s his/her case), while judge B may not (even though judge B has the role of ‘judge’). Thus in order to change records in the digital dossier (via an agent) a user not only has to be assigned a specific role, he/she also has to be on the access control list of a specific criminal dossier. Figure 2 shows such an access control list where subjects identified by numbers 1234, 2345, 3456 and 4567 have read access for a dossier and only subjects 1234 and 4567 have write access.

The distinction between roles and individuals is crucial in a dynamic environment (such as is the case associated with the digital dossier). Security policies based on roles can be regarded as static (or at least ‘long lived’) and are typically globally valid (at all possible locations), while individual access control lists are typically dynamic (or ‘short lived’). Individual policies typically only apply per dossier, or even shorter, for example, when a criminal dossier is handed over to another clerk, prosecutor or judge. The combination of static and dynamic access control rules should also limit the aforementioned ‘label creep’ phenomena.

In addition, each user also has a public/private key pair and a corresponding digital certificate, as specified in the X509 standard [13]. The certificates are organized in a standard PKI infrastructure [12], which is run at the Public Prosecution, and are used for signatures on individual

records, to enforce integrity of a digital criminal dossier in its totality.

B. Dossier Creation

The above-mentioned access control model is used to authenticate (the role of) Public Prosecutor(s). A dedicated agent then enforces the policy that only Public Prosecutors can create dossiers.

C. Control

The above mentioned access control lists are maintained per document, where a document can (1) either be a digital criminal dossier maintained by the Public Prosecutor (as seen in Figure 2), (2) an administrative dossier maintained by one of the other organizations (for example, `ref:MuniDatAmsterdam` in Figure 2) or (3) another kind of document (for example `ref:PP-00001-movie1.avi` in Figure 2). Each document thus has its own access control list. This means, in particular, that someone can have access to the digital criminal dossier, stored at the Public Prosecution, without having access to all of the *referring records* the criminal dossier contains. This ensures that local organizations keep control over access to their information. Note however that organizations are not completely free to determine who can access a specific document. They do have to act in accordance with Dutch law. For example, the Council for Child Welfare must inform the Public Prosecutor, that is responsible for a case involving a minor, of the minor's social environment. The Council for Child Welfare can chose to only allow this specific Public Prosecutor access to the document, without, for example, allowing legal clerks to view the information it contains.

Similarly, organizations or individuals may transfer their access rights to others as long as this is allowed by Dutch law. For example, a Public Prosecutor can transfer a criminal dossier, together with the access rights to the dossier, to another Public Prosecutor.

D. Notification

Dedicated agents are deployed to automatically inform all involved parties of the progress of a case.

Users are also notified when certain 'exceptional' behavior occurs within the system. Examples include inconsistencies found between records in a dossier (point *L* below) or multiple failures to comply to the access control mechanism (point *A* above). In such instances dedicated clerks are notified who deal with the problem.

E. Persistence

Information in a criminal dossier changes during the compilation of a dossier. To ensure that no information is accidentally or intentionally removed, *all* versions of a dossier are securely stored by a dedicated agent. This storage system also serves as a backup store. The

information is encrypted and chained signatures, i.e., for each version, are added to ensure both confidentiality and integrity of the criminal dossiers. An outside trusted third party is responsible for maintaining this backup site. See also points *F* and *N* below.

Another dedicated agent guards the lifelines of data entries in the digital dossier. For example, information obtained from the Council for Child Welfare concerning juvenile suspects may, by Dutch law, only be kept for a maximum period of five years and should also be destroyed when the subject turns 18. Note that simply removing a reference in the central part of the dossier at the Public Prosecution is enough to obtain this result.

F. Attribution

Criminal dossiers can only be accessed through agents. Each user of the system, for example a public prosecutor or a judge, uses its own agent to access a criminal dossier (or part thereof). Each individual also has his/her own X509 public/private key-pair. Whenever an agent performs an action, such as reading or altering a file, this is logged by a signature performed by the agent, *on behalf of its owner*. This ensures that actions can be attributed to specific users.

A dedicated logging agent is responsible for logging all information per dossier (the signatures that attribute who changed what, when, etc.). This information is safely stored (preferably offline and encrypted) and needs to be integrity preserving (using signatures). The same trusted third party that ensures persistence of the criminal dossiers (point *E* above) is responsible for storing this information securely.

G. Information Flow

Access control (point *A* above) together with control (point *C*) are used to limit illicit information flow. However, totally disallowing illicit information flow is in practice not possible. See the discussion in Section VI for more on this point.

H. Aggregation Control

Effectively limiting illicit aggregation control is probably as difficult as effectively stopping illicit information flow (see above). In principal this is handled via access control, restricting access to only a handful of dossiers makes aggregating information impossible. However, an organizations typically *want* to aggregate their information, for example to do efficiency studies or scientific research, this is not an option. By limiting the number of persons that have access to (almost) all dossiers unwanted aggregation can be kept to a minimum. See the discussion in Section VI for more on this subject.

I. Trusted Computing Base

All computer systems running the agent platforms must be trusted: form a trusted computing base. This is in

particular true for the centralized system run by the Public Prosecution. This computer system stores all skeleton files of digital criminal dossiers and runs both the PKI infrastructure and the lookup service needed to find agents and services in the distributed system.

Viewed in its totality as a distributed system, damage is only severe if the Public Prosecution's computer system is compromised, see Section VI for a discussion of this issue. Dedicated computers (not used for day-to-day computing), dedicated users, including system administrators, and audit procedures are means to insure trust in the system.

J. Secure Transfer

The agent platform provides secure transfer (based on SSL tunnels) between agent platform locations for free. Other security requirements are also handled by the agent platform [14]: integrity of agents and their data and secure communication between platform locations run by organizations such as the Public Prosecution and the Council for Child Welfare.

K. Compartmentalization of Information

Compartmentalization of information is again related to the access control system. As access to criminal dossiers is restricted to a few dossiers per employee this property can be guaranteed.

L. Consistency

Consistency is checked whenever information from an outside source (such as from the Council for Child Welfare) is entered in the digital criminal dossier. A dedicated consistency agent (per dossier) checks if all administrative data from the outside source matches the data in the digital dossier. If this is not the case a (human) agent needs to decide how to act further. The consistency agent marks such an event in the meta data of the digital dossier. It is also possible for the agent to make an 'educated guess' related to (simple) consistency issues, but a human user needs to confirm this. See [3] for more details on this complicated issue.

M. Completeness

A dedicated agent checks the completeness of each criminal dossier. A dossier should always include the required minimal information such as personal information, criminal charge and warrants. Additional completeness checks are performed on a per case basis. For example, a dossier concerning a drunken driver case should include a rapport that details the factual information of the alcohol blood level at the time of the offense, again see [3].

N. Backups

A special purpose backup agent is deployed that facilitate secure backups of the digital dossier (combined with the persistence database mentioned in point *E*). This agent's functionality is combined with the logging agent (from point *F*), as these agents share a lot of functionality.

Note that the architecture described is inherently modular, new security requirements can be added by new agents that can be used whenever required.

V. DOMAIN SPECIFIC LEGAL REQUIREMENTS

A distributed digital criminal dossier in combination with a multi-agent system to access and secure digital dossiers has major benefits. However, a number of domain specific legal issues in relation to the use of agent technology in a criminal trial, in a Dutch legal setting, remain.

From a legal perspective it is important for all parties to know to whom an agent belongs, as it acts on behalf of this person/organization [15]. Identification is also important when dealing with liability and compliance with agreements. Legally, software agent themselves are not responsible for their actions. The owner/user is always responsible [16]: there must be a link between the agent and its provider or user, see Section IV-F.

Relatively strict privacy regulations apply. As a guideline, Dutch law forbids the exchange of administrative data, unless there is a valid legal reason to do so. In all cases in which, for example, the Council for Child Welfare is legally allowed to release personal data for the digital dossier, it needs to inform the offending minor of its action. The Public Prosecution is then expected to carefully handle the received reports to guarantee the privacy of the minor. In such instances the Council for Child Welfare is also required by law to add a record to the (local) dossier of the offending minor that states, when and to whom what data was supplied.

Another criterion that applies in certain cases is the so called 'protection of others criterion' [17]. In the example of data exchange between the Public Prosecution and the Council for Child Welfare this means that the Public Prosecution can only give information to the Council for Child Welfare when this is necessary for a good execution of the tasks of the Public Prosecution and insofar a weighty public interest is involved. Furthermore, giving information to the Council for Child Welfare must serve a purpose as stated in legislation. In the case of possible child abuse or other domestic violence in which minors are victims, the purpose is prevention of criminal acts or supporting victims. In that case information is given to the Council for Child Welfare. For the preparation of dossiers the Minister of Justice gives copies of reports in personal dossiers to the Director of the Council for Child Welfare. Additional research for each type of offense is needed to obtain a more complete overview of all legal requirements and their applicability.

VI. DISCUSSION AND CONCLUSIONS

All main security requirements for an information management system for criminal dossiers have been addressed in this paper. Distributed digital criminal dossiers used together with a dedicated multi agent system solve most of these. The remainder of this section discusses some of the open questions and issues identified above.

In addition to the individual security properties described above, the system as a whole also has some high level security properties. The main security attribute is that each organization can use its own refined security policy. Such security policies state rules for all security related issues. Some of these are global policies, that is, they hold for all organizations involved. The role-based access control mechanism described above is one example. The use of one shared trusted lookup service to find individual agents in the distributed system another. Other rules in the security policy are however local: each local organization controls *who* has access to its files, but also which specific backup procedures are used, which privacy policy regarding its data holds, etc.

As long as the computer system of the Public Prosecution has not been compromised, the digital criminal dossiers are not at risk. The computer systems of the Public Prosecution, however, need to be trusted completely. The lookup service, public key infrastructure and global authentication mechanism are all hosted by the Public Prosecution. If other systems are compromised the digital dossier is not necessarily effected. Automatic consistency checking, performed each time part of a dossier is altered, detects modifications. If these modifications are unwarranted and detected by both the dedicated agent and the user, once informed, the original data is restored from the secure backup service.

As stated earlier, completely preventing illicit information flow remains an issue that is not solved at this moment. Although it is not possible to move (parts of) criminal dossiers to other dossiers as these are guarded by access control lists, it is always possible for someone to move the *information* contained in a criminal dossier to another dossier 'by hand'. More research on this subject is needed.

Preventing illicit aggregation control is another problematic issue, as organizations typically want to aggregate their information. One possibility is to only allow aggregation of anonymized criminal dossiers, thereby disallowing queries to individuals. But anonymizing information in such a way that it is possible to draw some high level conclusions -wanted by the organizations involved- and disallow querying specific individuals is very hard, if not impossible [18].

This paper studies the security issues related to the use of distributed digital criminal dossiers. Other requirements, such as reliability, scalability and performance, have not been studied. These issues are addressed in a prototype implementation of the system proposed in this paper which is currently under development. The agent platform Agentscape [7] is being used to realize this

system.

ACKNOWLEDGMENT

The authors thank Martin Apistola for comments and contributions to early drafts of this paper. This research is supported by the NLnet Foundation, <http://www.nlnet.nl>, and is conducted as part of the Agent-based Criminal Court Electronic Support Systems (ACCESS) project, <http://www.iids.org/access>, initiated by the VU University Amsterdam together with the Courts of Amsterdam and Rotterdam, and financed by the Dutch Court for Jurisdiction (Dutch: Raad voor de Rechtspraak) and by the NWO TOKEN program.

REFERENCES

- [1] L. Storchi, "Het controleren van het digitale dossier op volledigheid met behulp van software agenten," Master's thesis, Vrije Universiteit Amsterdam, IIDS group, sept 2005.
- [2] M. Luck, P. McBurney, and C. Preist, *Agent Technology: Enabling Next Generation Computing (A Roadmap for Agent Based Computing)*. AgentLink, 2003.
- [3] M. Warnier, F. M. T. Brazier, M. Apistola, and A. Oskamp, "Towards automatic identification of completeness and consistency in digital dossiers," in *Proceedings of the Eleventh International Conference on Artificial Intelligence and Law (ICAIL'07)*. ACM Press, 2007, pp. 177–182.
- [4] —, "Distributed Digital Data: Keeping files consistent, timely and small," in *Proceedings of the eGovernment Interoperability Campus 2007 Conference (eGovINTEROP'07)*, 2007, to appear.
- [5] R. J. Anderson, "Clinical System Security – Interim Guidelines," *British Medical Journal*, vol. 312, pp. 109–111, 1996.
- [6] M. Wooldridge and N. R. Jennings, "Intelligent Agents: Theory and Practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995.
- [7] B. J. Overeinder and F. M. T. Brazier, "Scalable middleware environment for agent-based internet applications," in *Proceedings of the Workshop on State-of-the-Art in Scientific Computing (PARA'04)*, ser. Lecture Notes in Computer Science, vol. 3732. Copenhagen, Denmark: Springer, June 2004, pp. 675–679.
- [8] F. Bellifemine, A. Poggi, and G. Rimassa, "JADE—A FIPA-compliant agent framework," *Proceedings of PAAM*, vol. 99, pp. 97–108, 1999.
- [9] V. Roth and M. Jalali-Sohi, "Concepts and architecture of a security-centric mobile agent server." in "Proc. of the Fifth International Symposium on Autonomous Decentralized Systems (ISADS 2001)". IEEE Computer Society, 2001, pp. 435–442.
- [10] A. Sabelfeld and A. C. Myers, "Language-Based Information-Flow Security," *IEEE Journal on selected areas in communications*, vol. 21, no. 1, 2003.
- [11] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-Based Access Control Models," *Computer*, vol. 29, no. 2, pp. 38–47, 1996.
- [12] C. Kaufman, R. Perlman, and M. Speciner, *Network Security, PRIVATE Communication in a PUBLIC World*, 2nd ed. Prentice Hall, 2002.
- [13] C. Adams and S. Farrell, "RFC2510: Internet X. 509 Public Key Infrastructure Certificate Management Protocols," *Internet RFCs*, 1999.

- [14] G. van 't Noordende, F. M. T. Brazier, and A. S. Tanenbaum, "Security in a mobile agent system," in *Proceedings of the First IEEE Symposium on Multi-Agent Security and Survivability*, Philadelphia, 2004.
- [15] L. B. B. W. Schermer, M. Durinck, "Juridische aspecten van autonome systemen," ECP.NL, Tech. Rep., 2005.
- [16] B. W. Schermer, "Handreiking voor gedragsregels autonome systemen," ECP.NL, Tech. Rep., 2006.
- [17] M. R. Bruning, *Over sommige kinderen moet je praten*. Universiteit Leiden, 2006, oratie.
- [18] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*. John Wiley and sons, Inc., 2001.