Sensitive data transaction in Hippocratic Multi-Agent Systems

Ludivine Crépin[1,4], Yves Demazeau[1], Olivier Boissier[3], and François Jacquenet[4]

[1] Laboratoire d'Informatique de Grenoble - CNRS
{Ludivine.Crepin, Yves Demazeau}@imag.fr
[2] Centre G2I
Ecole Nationale Supérieure des Mines de Saint-Etienne
Olivier.Boissier@emse.fr
[3] Laboratoire Hubert Curien (UMR CNRS 5516)
Université Jean Monnet, Saint-Etienne, France
Francois.Jacquenet@univ-st-etienne.fr

Abstract. The current evolution of Information Technology leads to the increase of automatic data processing over multiple information systems. The data we deal with concerns sensitive information about users or groups of users. A typical problem in this context concerns the disclosure of confidential identity data. To tackle this difficulty, we consider in this paper the context of Hippocratic Multi-Agent Systems (HiMAS), a model designed for the privacy management. In this context, we propose a common content language combining meta-policies and application context data on one hand and on the other hand an interaction protocol for the exchange of sensitive data. Based on this proposal, agents providing sensitive data are able to check the compliance of the consumers to the HiMAS principles. The protocol that we propose is validated on a distributed calendar management application.

Keywords: Privacy, Sensitive Data Transaction, Confidentiality, Multi-Agent Systems, Interaction Protocol

1 Introduction

With the use of multiagent technologies, the sensitive data transmission problem in Multi-Agent Systems (MAS) is all the more present since users delegate their sensitive data to an autonomous agent (the interaction is an essential feature of Multi-Agent Systems). Spread of sensitive data over the Internet using autonomous entities becomes an important risk that requires to be considered nevertheless this problem has not received enough attention by the researchers in the domain until now.

We have proposed in [1] the model of Hippocratic Multi-Agent Systems (HiMAS) that takes into account this data sensitivity regarding moral issues and not legal aspects. This model defines the concept of *private sphere* for an agent

or a user to structure and to represent the data involved in the management of privacy, and the nine principles that should govern the functioning of a HiMAS so that privacy is preserved in the Multi-Agent Systems. In order to engineer agents societies according to this conceptual framework we focus in this article on a precise objective of the design of such a system: sensitive data protection during sensitive data transaction. Such a transaction represents a sensitive data transaction between two agents. To tackle this problem, we propose a sensitive data transaction protocol inspired by [2, 3], with an associated content language in the HiMAS context. This protocol is our first step for the implementation of a HiMAS. To illustrate this protocol, we have chosen the distributed calendar management application presented in [4].

The next section briefly presents the model of Hippocratic Multi-Agent Systems in order to draw the global context in which we place our present work. Section 3 focuses on the definition of the content language and the associated semantics used in the protocol that we propose in section 4. We present an application of our sensitive data transaction protocol in section 5. Finally we talk about related work in section 6 and conclude with some perspectives on the future work.

2 Foundations: Hippocratic Multi-Agent Systems (HiMAS)

As introduced in the previous section, the HiMAS model proposed in [1] is composed of two main components: the private sphere representation and some hippocratic principles that we present in the following sections. The reader interested in more information about this model and the private sphere, may refer to [1].

2.1 Private sphere, consumer and provider

The private sphere contains information that an agent considers as sensitive, represented by sensitive data, and all the associated management rules. For instance, in the context of calendar management [4], sensitive data is the user's slots of time or meetings that are delegated to an agent. The agent's private sphere represents all this kind of data and all the rules defining the conditions of its disclosure, its use or its sharing for example.

To define the private sphere dimensions, we are inspired by many researches in social science [1]. The first one focus on the **ownership rights** of sensitive data. They are only assigned to agents concerned by this data [5]. Moreover the private sphere is also **personal** [6, 7], **personalizable** (the agent chooses what its private sphere contains) [8–10] and **context-dependent** [11, 12].

To represent the possible positions of an agent with respect to the private sphere, we define three roles. The **consumer** role characterizes the agent which asks for sensitive data and uses it. The **provider** role characterizes the agent

which discloses sensitive data[4]. The last role, the **subject**, describes the agent from whom origined sensitive data.

With this definition of the agent's private sphere we install a provider-centred view on the management of sensitive data. This is due to the fact that we mainly have a user-centred view on privacy preserving: user should be confident in the management of the sensitive data they delegate to their personal agent.

2.2 Nine principles for HiMAS

The HiMAS model is inspired by the Hippocratic Databases [14]. In order to preserve privacy, a HiMAS must respect the nine principles described below.

1. **Purpose specification**: The provider must know the objectives of the sensitive data transaction. Therefore it can evaluate the transaction consequences.
2. **Consent**: Each sensitive data transaction requires the provider's consent (and the subject's consent if it is not the same agent).
3. **Limited collection**: The consumer commits to cutting down to a minimum the amount of data for realizing its objectives.
4. **Limited use**: The consumer commits to only use sensitive provider's data to satisfy the objectives that it has specified and nothing more.
5. **Limited disclosure**: The consumer commits to only disclose sensitive data to reach its objectives. Moreover it must disclose it the least number of times possible and to the least number of agents.
6. **Limited retention**: The consumer commits to retain sensitive data only for the minimum amount of time it takes to realize its objectives.
7. **Safety**: The system must guarantee sensitive data safety during storage and transactions.
8. **Openness**: The transmitted sensitive data must remain accessible to the subject and/or the provider during the retention time.
9. **Compliance**: Each agent should be able to check the obedience to the previous principles.

3 Content language for sensitive data transaction

In order to integrate the HiMAS principles in the interaction protocol that we propose, we have chosen to define the semantics of these principles. This study also leads us to determine the different links between these principles. The first step of this work is to group together these principles according to their purpose into the HiMAS agent's reasoning: during the sensitive data transaction; during the other interactions; and in relation to the system implementation. After the study of this semantics, we propose a representation of the required principles in a content language. These two steps are the foundations of the sensitive data transaction protocol that we propose.

[4] We can notice that this vision is the opposite of the centered service vision like for example [13], regarding the consumer and the provider.

3.1 Content language semantics

Let us consider the principles that play a part in a sensitive data transaction. In such a context, the provider defines a **policy** and the consumer a **preference** to define their desires regarding the sensitive data manipulations.

The consumer's policy and the provider's preference are similar to the policy and the preference defined in [2]: these concepts are composed of the transaction objectives[5], the deletion time of collected data, a broadcasting list and the data format (required references).

In order to map a policy to a preference, a sensitive data transaction groups together required sensitive data with the consumer's policy and the provider's consent and preference.

Seven of the nine HiMAS principles play a part in sensitive data transactions:

- **1. Purpose specification**: The consumer asks for provider's sensitive data in order to realize required tasks. Since the consumer must declare his purpose, these tasks should be used to define its objectives. The consumer must send them to the provider.
- **3. Limited collection**: With the definition of its objectives, a consumer can select the sensitive data that is only required for the realization of its objectives.
- **4. Limited use**: The consumer can then determine the possible uses of the collected sensitive data by virtue of its objectives.
- **5. Limited disclosure**: The objectives enable the consumer to determine which agents are allowed to receive the collected sensitive data.
- **6. Limited retention**: The specification of the objectives defines also the sensitive data retention time for the consumer.
- **8. Openness**: The openness implies that the provider and/or the subject are in the broadcasting list.
- **2. Consent**: The mapping between a policy and a preference represents the consent principle that is made after the respect of the principles previously presented.

Principles must be also considered in the different interactions that could take place in the system. We should insure that the consumer respect the **9. Compliance** principle in these interactions.

The last principle, **7. Safety**, has not to be considered in the agents reasoning since it relates to the system design and is therefore not included in the formalization presented in this article.

The semantics of the principles playing a part in the sensitive data transaction. During sensitive data transaction, the central principle for the agent's reasoning is **1. Purpose specification** (Figure 1).

[5] The objectives are close to the concept of goal, like for example in BDI model [15] or [16].

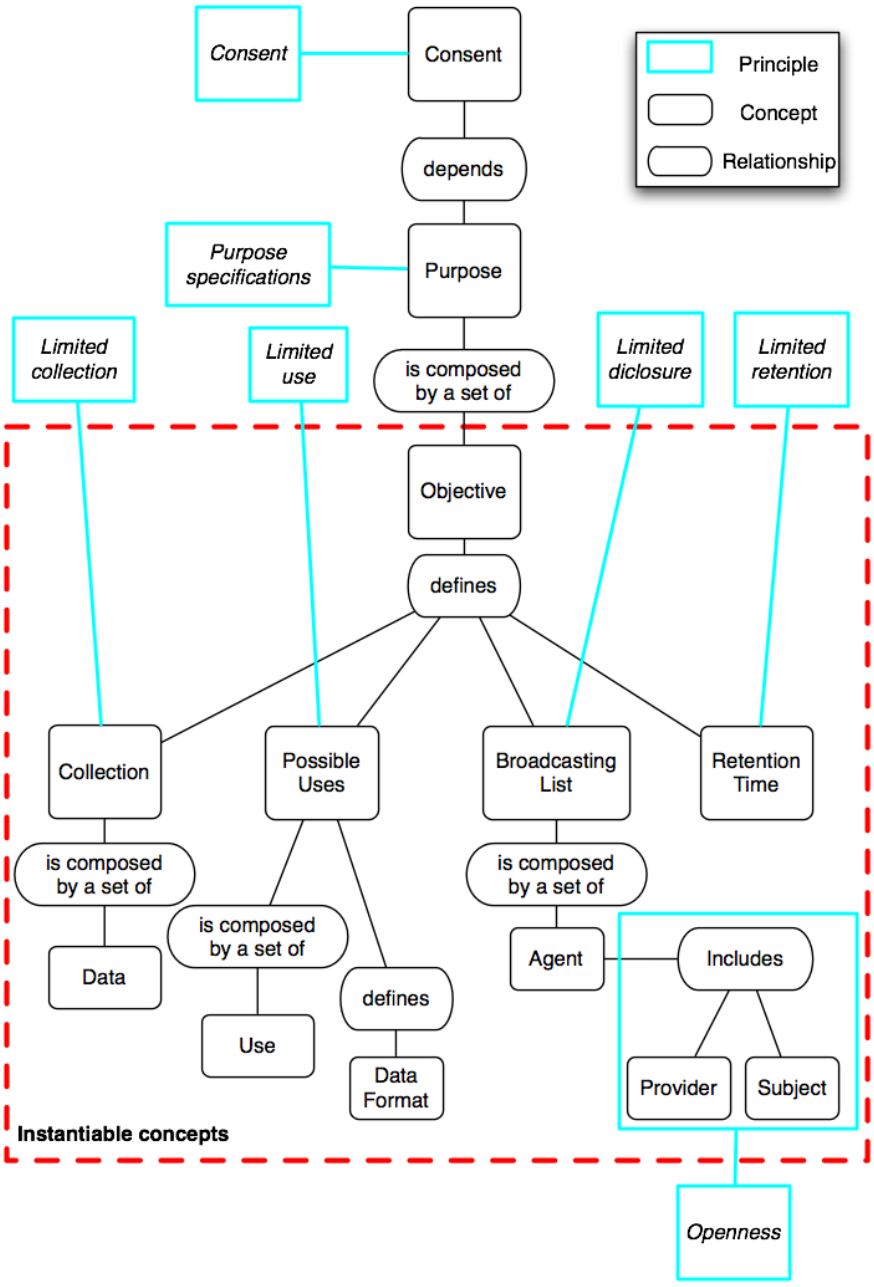

Fig. 1. Conceptual graph representing the semantics of HiMAS principles

Principle	Associated Concept
1. Purpose specification	*Purpose* composed by a set of *Objective*
3. Limited Collection	*Collection* composed by a set of *Data*
4. Limited Use	*PossibleUses* composed by a set of *Use*
5. Limited disclosure	*BroadcastingList* composed by a set of *Agent*
6. Limited retention	*RetentionTime*
7. Openness	*Subject* and *Provider* included in *Agent*
2. Consent	*Consent*

Table 1. Concept representing HiMAS principles.

For each principle (and for the notion of format[6] that is required in our approach) we define an associated concept in a conceptual graph [17] (refer to Table 1 and to Figure 1). Each principle and the notion of format is represented by a concept linked to another according to a semantic relationship. In order to define these, we use an existential positive conjunctive fragment of the first order logic that allows us not to obtain contradictory logical information. We represent each concept by an atomic predicate and each relationship by a binary predicate. The formal description of the conceptual graph presented in Figure 1 is described in Table 2.

	Principle: 1. Purpose Specification
$\forall p\ Purpose(p)$	$\rightarrow\ \exists x\ composedBy(p,x) \wedge Objective(x)$
	Principle: 3. Limited collection
$\forall x\ Objective(x)$	$\rightarrow\ \exists y\ defines(x,y) \wedge Collection(y)$
$\forall y\ Collection(y)$	$\rightarrow\ \exists z\ composedBy(y,z) \wedge Data(z)$
	Principle: 4. Limited use
$\forall x\ Objective(x)$	$\rightarrow\ \exists y\ composedBy(y,z) \wedge PossibleUses(y)$
$\forall y\ PossibleUses(y)$	$\rightarrow\ \exists z\ composedBy(y,z) \wedge Use(z)$
$\forall y\ PossibleUses(y)$	$\rightarrow\ \exists z\ defines(y,z) \wedge Format(z)$
	Principle: 5. Limited disclosure
$\forall x\ Objective(x)$	$\rightarrow\ \exists y\ defines(x,y) \wedge BroadcastingList(y)$
$\forall y\ BroadcastingList(y)$	$\rightarrow\ \exists z\ composedBy(y,z) \wedge Agent(z)$
	Principle: 8. Openness
$\forall z\ Agent(z)$	$\rightarrow\ \exists w\ includes(z,w) \wedge Subject(w)$
$\forall z\ Agent(z)$	$\rightarrow\ \exists w\ includes(z,w) \wedge Provider(w)$
	Principle: 6. Limited retention
$\forall x\ Objective(x)$	$\rightarrow\ \exists y\ defines(x,y) \wedge RetentionTime(y)$
	Principle: 2. Consent
$\forall c\ Consent(c)$	$\rightarrow\ \exists x\ depends(c,p) \wedge Purpose(p)$

Table 2. Principles formalization.

[6] All the required references.

The implementation of this conceptual graph is made by using an OWL file [18]. Figure 2 presents an example of our implementation. We have chosen to present the instantiation of the relationship *isComposedBy* for the concepts *Collection* and *Data*. This approach uses an extensible knowledge representation language, RDF and RDFS. Each associated concept is represented by a RDFS class and each semantic relationship by an OWL property. RDFS gives a vocabulary to RDF that instantiates RDFS classes and properties. So each instantiation (application context-dependent) of these concepts and these semantic links is in a RDF structure in relation to the vocabulary defined by the RDFS.

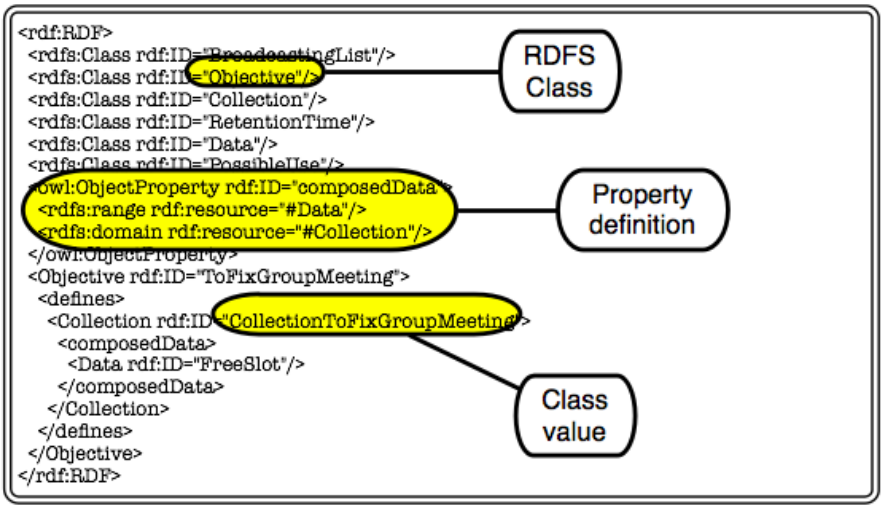

Fig. 2. Example of the conceptual graph implementation

Taking the context of the application into account. HiMAS principles define generic constraints that the agency must satisfy to preserve the private sphere. The previous study semantics that we have just presented, must be linked to the HIMAS application context because of the context-dependent characteristic of the private sphere. An example of the introduction of the context is presented in more details is section 5.

For this integration, we need to instantiate the defined conceptual graph by giving all the possible values for each concept according to the application context and by linking these values (see the dotted block in Figure 1). These values are represented in a RDF structure (see Figure 2).

We have chosen to not instantiate the possible values of two concepts: consent and purpose. Indeed the value of the consent concept can be true or false. Therefore it can be represented by a boolean and we need only to define the se-

mantic links of this concept for the agents' reasoning. We indicate just that the provider must give or not its consent according to the consumer's purpose. This last concept is composed by a set of objectives. Therefore, by defining all the possible values for the objective concept, we define also all the possible values for the purpose.

3.2 Content language syntax

We sum up first all the requirements for sensitive data transaction in a HiMAS represented in Figure 3. Then we present the syntax of such a transaction.

In [1], we have shown that HiMAS agents have to determine risk-taking for a sensitive data transaction. During sensitive data transaction, the consumer (resp. provider) builds its policy (resp. preference) according to its intention. Before building such a transaction, the HiMAS agents pass a judgement on the other HiMAS agents regarding their reliability. For example, this function can be implemented by a processus of trust management like in [19]. If the consumer and the provider are reliable, then the transaction can begin.

We begin the description of the content language elements according to the chronological order of a sensitive data transaction: the design of the policy, the sensitive data transaction and the design of the preference.

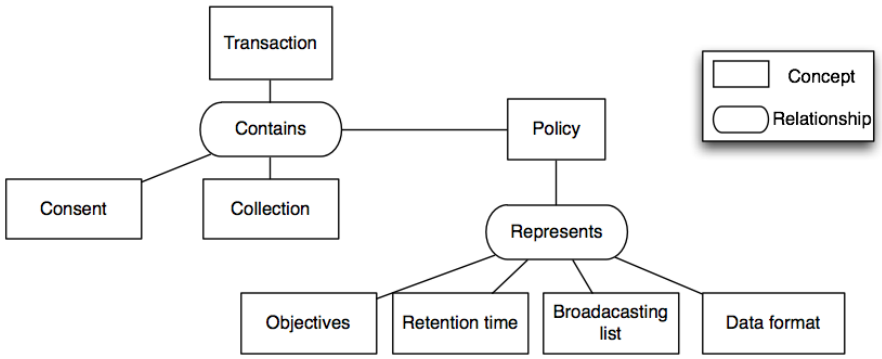

Fig. 3. Sensitive data transaction representation

Policy A policy must contain the objectives, the retention date, the broadcasting list and the data format for each asked data (Figure 3 and Table 3).

Once the consumer has determined its objectives and the concepts representing them, it builds a policy syntactically (using an XSD schema) and semantically valid (using an OWL file).

$\forall y\ policy(y)$	$\rightarrow$	$\exists z\ represents(y,z) \wedge objective(z)$
$\forall y\ policy(y)$	$\rightarrow$	$\exists z\ represents(y,z)\ \wedge format(z)$
$\forall y\ policy(y)$	$\rightarrow$	$\exists z\ represents(y,z)\ \wedge broadcastingList(z)$
$\forall y\ policy(y)$	$\rightarrow$	$\exists z\ represents(y,z)\ \wedge retentionTime(z)$

Table 3. Policy formalization.

Sensitive data transaction We have defined in [1] such a transaction set up a policy, a preference, the provider's consent and the sensitive data requested by the consumer. Notice that the formalization presented in Figure 3 does not refer to the provider's preference. Indeed a preference and a policy are based on the same concepts and we represent the provider's preference by the modifications that the provider induces from the consumer's policy if there is no agreement on the constraints defined in the policy.

All values for all elements of the transaction are defined in the content language that allows the consumer to build a valid transaction with regard to the privacy preservation.

$\forall y\ transaction(y)$	$\rightarrow$	$\exists z\ contains(y,z) \wedge consent(z)$
$\forall y\ transaction(y)$	$\rightarrow$	$\exists z\ contains(y,z) \wedge collection(z)$
$\forall y\ transaction(y)$	$\rightarrow$	$\exists z\ contains(y,z) \wedge policy(z)$

Table 4. Sensitive data transaction formalization.

In order to build a sensitive data transaction that is syntactically valid, we use the same approach as for the policy. We formally define such a transaction in Table 4 and in Figure 3.

4 Sensitive data transaction protocol

In this section, we propose to formalize a sensitive data transaction protocol based on the content language previously defined. This approach also allows us to provide a guideline about the design of the policy and preference for the HiMAS agents.

In our content language, the consumer's objectives are semantically linked to the principles playing a part in a sensitive data transaction. This content language includes all the possible values for each class representing one HiMAS principle. The consumer can therefore know if it violates the private sphere or not by verifying that the elements contained in its policy are included in the content language and by verifying that it respects the semantic links between these elements.

The sensitive data transaction protocol that we propose is presented in Figure 4. The content language implementation must be common to all the HiMAS

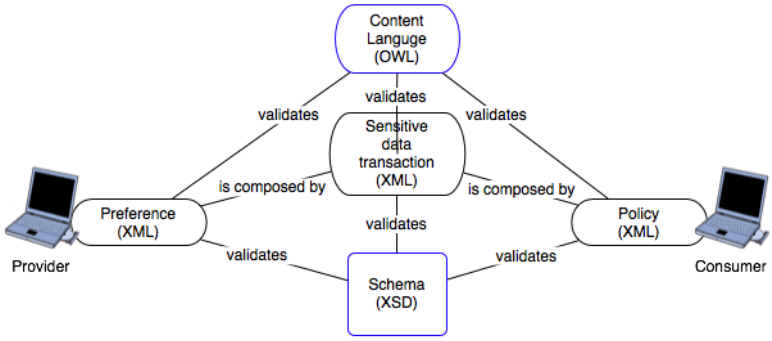

Fig. 4. Sensitive data transaction protocol

agents so that each agent can base its reasoning on the same vocabulary and the same semantics. We have chosen to represent this as external to the agents and available for the consultation by the agency. With this approach, many HiMAS can refer to the same language if their application context is the same. Moreover, this technique allows us to consider the openness between many HiMAS having the same context. At a design level, the possible modifications for this language require only one control entity and there are no propagation problems.

Each consumer and each provider validate their policy and their preference using the content language previously presented in order to build and to execute a sensitive data transaction.

4.1 Steps of the interaction protocol

We present now the three steps of the interaction protocol that we propose in a chronological order: the design of the policy, the sensitive data transaction and the design of the preference. These steps are represented in Figure 5.

Design of the policy A consumer builds its policy according to its objectives by using the content language. In this way, it can be understood by the other agents. Moreover the consumer's behavior respects the private sphere if its policy validates the content language.

A first constraint of our protocol imposes that the XSD file validates the XML file to ensure the syntax of such a transaction.

A second constraint of our protocol imposes that the values of the XML file must be included in the conceptual graph previously defined (see Figure 1) to ensure the semantics.

Sensitive data transaction Once the consumer has defined and validated its policy, the sensitive data transaction can begin.

To inform the provider about its request, the consumer must build a sensitive data transaction. This transaction contains its policy and must be validated by the content language.

Once the sensitive data transaction file built and validated, the consumer can send it to the provider in order that the provider could know its request. This step is represented by the first interaction of Figure 5.

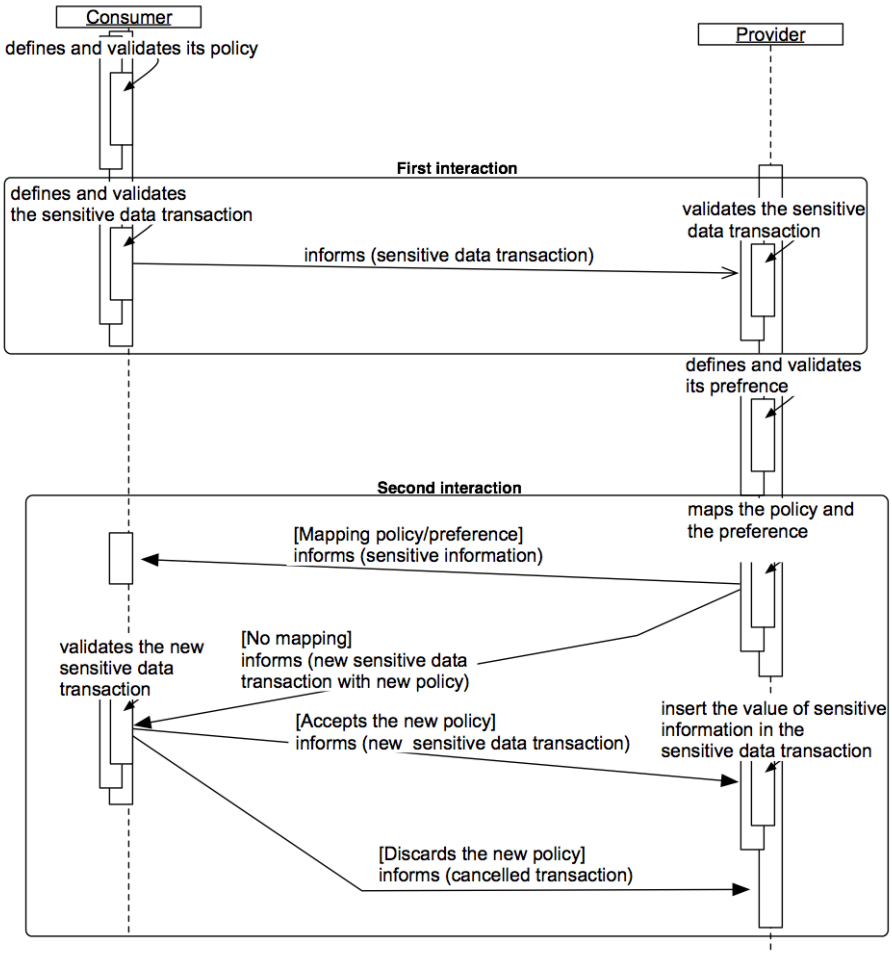

Fig. 5. Sensitive data transaction protocol

Design of the preference From the management rules of its private sphere, a provider establishes the conditions of the use, the disclosure, the retention of its

sensitive data. Once it received a sensitive data transaction, these rules allow it to accept or not the consumer's policy.

Before analyzing the consumer's policy, the provider must first verify the transaction validity at a syntactic and semantic level, using the content language. These two validations allow to determine if a consumer has a malicious behavior on the limitations imposed by HiMAS principles and on the sensitive data transaction protocol.

If the sensitive data transaction is validated, then the provider can make a mapping between its preference and the consumer's policy. If no mapping is found, the provider can propose to the consumer some adaptations of its policy.

Once the consumer and the provider have agreed on the policy, the provider completes the transaction with the values of requested sensitive data. If no agreement is found, the transaction is canceled and the provider can not answer to the consumer's request. The second interaction of Figure 5 represents these steps.

4.2 Synthesis

One of the first advantages of this approach is the possibility to verify the constraints defined by the principles of the HiMAS thanks to the content language. The consumer (resp. provider) can design its policy (resp. preference) with respect to the constraints defined by HiMAS principles. This obedience is made by the semantic links between the concepts representing the HiMAS principles.

Each transaction between the consumer and the provider can be represented by the "inform" communicative act of FIPA [20]. Indeed, these two agents exchange only one specific data: a sensitive data transaction that will be completed during such a transaction.

This protocol is provider-centred and is opposite to all the most of transaction protocols that are in general service-centred. It defines the same principles as the P3P [2] and sensitive data transaction as an interaction in ISLANDER [21]. In order to preserve completely the private sphere, this protocol must be integrated in a secure communication medium (principle **7. Safety**) which is not purpose in this paper.

5 Application

In order to illustrate the HiMAS model and the sensitive data transaction, we consider a decentralized calendar management application [4]. In this context, each user is represented by an agent in charge of the scheduling of events, either tasks or meetings. Timetables can be shared with other agents. When agents do not share their timetables, a negotiation system is necessary to fix the meetings.

We have chosen a simple example for the illustration of the sensitive data transaction protocol: a consumer wants to fix a group meeting with a provider and other agents (group G) in a given period of time (interval between two slots of time). We consider as sensitive data the free and occupied slots of time in users' calendar. Figure 6 represents this example.

In order to fix such a meeting, we define the following constraints:

– The sensitive data that the consumer can collect is the free slots of time for a given period.
– The consumer can disclose this sensitive data to the group G and it must guarantee that the provider is able to access to this data.
– If the sensitive data was disclosed, all the possible references can be disclosed.
– The consumer can not retain collected data after a given time.
– The possible uses of the collected sensitive data are storage, negotiation and sharing.

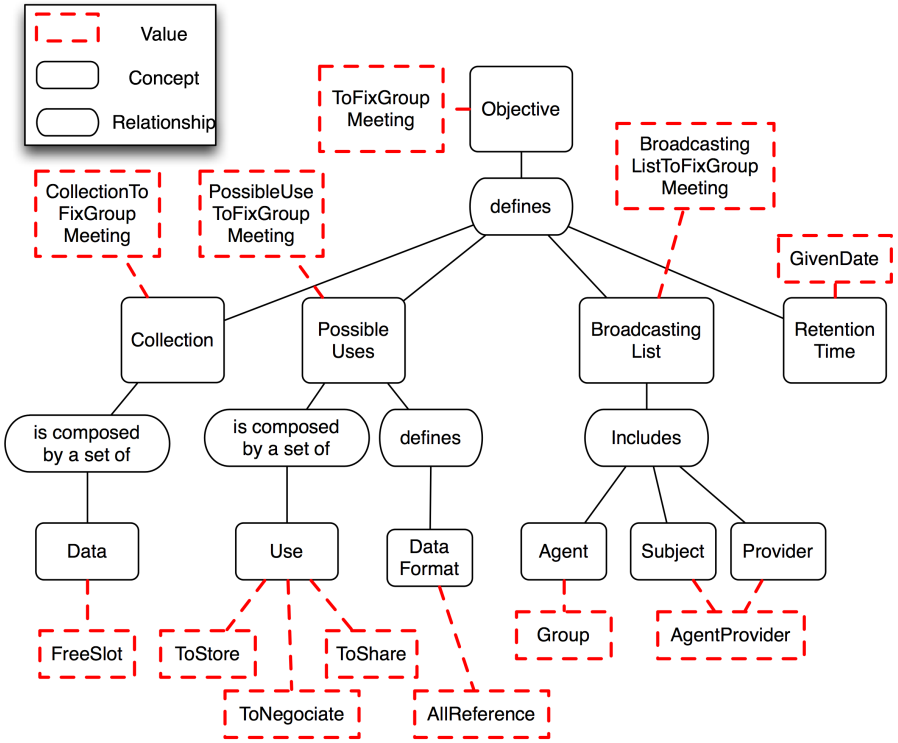

Fig. 6. Content language in context of calendar management and objective "to fix a group meeting"

The implementation of this HiMAS is made by instantiating the classes of Figure 2 with the values of Figure 6. For example, the class *Objective* is instantiated by the value "ToFixMeetingGroup" and this value defines the value "PossibleUseToFixGroupMeeting" (composed by the values "ToStore, ToNegociate, ToShare" linked to the class *Use*) for the class *PossibleUses*.

Once the content language is defined, the consumer and the provider can build, according to their intentions, a sensitive data transaction, regarding the privacy preservation.

The consumer builds its policy by parsing the content language. It first finds the objective corresponding to the goal "to fix a group meeting". After it chooses the values of its policy among the values proposed in the content language for its objective and sends a sensitive data transaction to the provider. The provider check the policy thanks to the content language in order to verify the consumer intentions. If it agrees with this policy, it informs the consumer of the required sensitive data. Else it can modify the consumer's policy by other values of the content language, according to its preference, and it informs the consumer of its modification. In this case, the consumer accepts or not this new policy.

6 Related work

The principles playing a part in the sensitive data transaction allow HiMAS agents to define their policy and their preference. This vision can be associated with the policy about policy that are the metapolicies. We propose in this section a global vision of this notion in order to present its main aspects.

Metapolicies are a notion introduced by Hosmer in [22, 23] that describe this like a set of policies about policies. These metapolicies are used in order to define a set of rules and assumptions about the policies of security in a given system for the policies interaction coordination.

Some other works use this notion like Kühnhauser [24] that uses metapolicies for the interfacing and the cooperation of complex policies, and for conflict resolution between the security policy. An other kind of work is the PONDER system [25, 26], where metapolicies are used in order to describe the security policies and to resolve the conflicts.

Generally the main objective of metapolicies is to define and to manage a set of policies of security for a given system regarding to the resolution of conflicts.

HiMAS principles define guidelines for the agents' reasoning about their policy and preference. These principles represent metapolicies for the agents behavior in relation to the communication and the manipulation of sensitive data. However the policy in our study case is not the same as in the work about security. HiMAS principles allow the agents to reason about a set of behavior constraints and do not allow to manage the set of agents' policies. We may link these principles to the notion of metaknowledge introduced by Pitrat [27].

7 Conclusion and perspectives

Our sensitive data transaction protocol allows us to apply seven HiMAS principles: **1. Purpose specifications**, **2. Consent**, **3. Limited collection**, **4. Limited use**, **5. Limited disclosure**, **6. Limited retention** and **8. Openness**. This protocol is generic and can be personalizable according to the kind of sensitive information that is exchanged.

The obedience to these principles consists in the consideration of our protocol at two levels. The first one is the definition of the content language. These principles are semantically and syntactically defined in a content language. The second one represents the use of the content language by the agents to build a sensitive data transaction.

The semantic links between the HiMAS principles allow us to determine in a content language, the maximal set of the sensitive data processing that a consumer can do on the collected data. A provider can also verify if a consumer respects the principles that limit the collection, the use, the disclosure and the retention, by referring to the content language. To ensure that all the principles are taken into account, we also formalize the sensitive data transaction that contributes to the malicious agent detection (agents that do not adhere to this formalization).

The content language of our protocol solves the main problem of the P3P [28]. Indeed, the mapping between a policy and a preference based on the same content language, a provider is able to understand the consumer's intention contrary to the P3P where this mapping is not guaranteed. Another advantage is the possibility to define the limitations imposed by HiMAS principles.

As a perspective, we want to focus on the principle of **9. Compliance** which is related to the problem of the interaction between agents. A first hint would be to implement a social order [29] in relation to the judgment function of HiMAS agents. We plan to implement this function using some trust management technics.

Acknowledgments : This work is supported by Web Intelligence project, financed by the ISLE cluster of Rhône-Alpes region. We thank France Telecom R&D for supporting the research related to trust mentioned in this paper.

References

1. Crépin, L., Vercouter, L., cois Jaquenet, F., Demazeau, Y., Boissier, O.: Hippocratic multi-agent systems. In: Proceedings of the 10th International Conference of Entreprise Information Systems. (2008) 301–308
2. W3C: Plateform for privacy preferences, http://www.w3.org/p3p/. (2002)
3. Cranor, L.F.: Web Privacy with P3P. O'Reilly (2002)
4. Demazeau, Y., Melaye, D., Verrons, M.H.: A decentralized calendar system featuring sharing, trusting and negotiating. In: Proceedings of the International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems. Volume 4031 of Lecture Notes in Computer Science., Springer (2006) 731–740
5. Thomson, J.J.: The right of privacy (1975) Philosophy and Public Affairs 4: 295-314.
6. Demeulenaere, P.: Difficulties of private life characterization from a sociologic point of view. In: Privacy in Information Society. Volume 11. (2002)
7. Baase, S.: A Gift of Fire: Social, Legal, and Ethical Issues in Computing. Prentice-Hall (2003)
8. Westin, A.F.: Special report: legal safeguards to insure privacy in a computer society. Commun. ACM **10**(9) (1967) 533–537

9. Warren, S.D., Brandeis, L.D.: The right to privacy. Wadsworth Publ. Co., Belmont, CA, USA (1985)
10. Lessig, L.: Code and Other Laws of Cyberspace. Basic Books, New York (2000)
11. Bellotti, V., Sellen, A.: Design for privacy in ubiquitous computing environments. In: Proceedings of the European Conference on Computer Supported Cooperative Work (ECSCW), Kluwer Academic Publishers (1993) 77–92
12. Palen, L., Dourish, P.: Unpacking "privacy" for a networked world. In: Proceedings of the 2003 Conference on Human Factors in Computing Systems, ACM (2003) 129–136
13. Rezgui, A., Ouzzani, M., Bouguettaya, A., Medjahed, B.: Preserving privacy in web services. In Chiang, R.H.L., Lim, E.P., eds.: In Proceedings of the Workshop on Web Information and Data Management, ACM (2002) 56–62
14. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic databases. In: Proceedings of the International Conference Very Large Data Bases, Morgan Kaufmann (2002) 143–154
15. Bratman, M.E.: Intention, plans, and practical reason. O'Reilly, Harvard University Press: Cambridge,MA (1987)
16. Sichman, J.S., Demazeau, Y.: Exploiting social reasoning to deal with agency level inconsistency. In: Proceedings of the First International Conference on Multiagent Systems, The MIT Press (1995) 352–359
17. Sowa, J.F.: Conceptual Structures: Information Processing in Mind and Machine. Addison-Wesley (1984)
18. W3C: Owl web ontology language, http://www.w3.org/tr/owl-features/. (2004)
19. Damiani, E., di Vimercati, S.D.C., Paraboschi, S., Samarati, P.: P2P-based collaborative spam detection and filtering. In: In Proceedings of 4th International Conference on Peer-to-Peer Computing, IEEE Computer Society (2004) 176–183
20. FIPA: Fipa communicative act library specification, http://www.fipa.org/specs/fipa00037/index.html. (2002)
21. Esteva, M., de la Cruz, D., Sierra, C.: Islander: an electronic institutions editor. In: Proceedings of the First International Joint Conference on Autonomous Agents & Multiagent Systems, ACM (2002) 1045–1052
22. Hosmer, H.H.: Metapolicies I. ACM SIGSAC Data Management Workshop **10**(2-3) (1991) 18–43
23. Hosmer, H.H.: Metapolicies II. In: Proceeding of the 15th National Computer Security Conference, Elsevier Advanced Technology Publications (1992) 369–378
24. Kühnhauser, W.E.: A paradigm for user-defined security policies. In: Symposium on Reliable Distributed Systems. (1995) 135–144
25. Lupu, E., Sloman, M., Dulay, N., Damianou, N.: Ponder: Realising enterprise viewpoint concepts. In: Proceeding of the 4th International Enterprise Distributed Object Computing Conference, IEEE Computer Society (2000) 66–75
26. Twidle, K.P., Lupu, E.: Ponder2 - policy-based self managed cells. In: Proceeding of the First International Conference on Autonomous Infrastructure, Management and Security (AIMS). Volume 4543 of Lecture Notes in Computer Science., Springer (2007) 230
27. Pitrat, J.: Métaconnaissance, Futur de l'Intelligence Artificielle. Hermés (1990)
28. Thibadeau, R.: A critique of P3P: Privacy on web, dollar.ecom.cmu.edu/p3pcritique/. (2000)
29. Castelfranchi, C.: Engineering social order. In: Proceeding of the First International Workshop Engineering Societies in the Agent World. Volume 1972 of Lecture Notes in Computer Science., Springer (2000) 1–18