

Specifying Standard Security Mechanisms in Multi-Agent Systems

Stefan Poslad

Queen Mary, University of London
Mile End Rd.
London E1 4NS
+44 207882 3754

stefan.poslad@elec.qmul.ac.uk

Monique Calisti

Whitestein Technologies AG
Gotthardstrasse 50
CH-8002 Zürich
+41 1 205-5500

mca@whitestein.com

Patricia Charlton

Motorola Labs
Espace technologique St Aubin
91193 Gif-sur-Yvette Cedex - France
+33 1 69 35 48 22

charlton@crm.mot.com

ABSTRACT

Distributed multi-agent systems propose new infrastructure solutions to support the interoperability of electronic services. Security is a central issue for such infrastructures and is compounded by their intrinsic openness, heterogeneity and because of the autonomous and potentially self-interested nature of the agents therein. This article reviews the work that the FIPA agent standards group has undertaken to specify security in multi-agent systems. This enables a discussion about the main issues that developers have to face at different levels (i.e., intra-platform, inter-platform and application level) when developing agent-based security solutions in various domains.

Keywords

FIPA, MAS security, standardization

1. INTRODUCTION

In the same spirit that the Internet developed for access to information, there is a vision to support open service environments with an e-business model that supports dynamic services, automated interaction, rich information exchange and tailored solutions. The research and development of Multi-Agent Systems (MASs) has often targeted providing “open” solutions and can provide some of the necessary infrastructure for supporting e-business solutions within open service environments.

Multi-agent systems represent virtual societies where software entities (agents) acting on behalf of their owners or controllers (people or organizations) can meet and interact for various reasons (e.g., exchanging goods, combining services, etc.) and in various ways (e.g., creating virtual organizations, participating to auctions, etc.). When deployed in an open environment such as the on-line business world, multi-agent systems face particularly challenging

trust and security issues at various levels. However, a major problem is that only very specific areas within the Internet space offer advanced security solutions to protect both service providers and consumers against malicious attacks. Furthermore, these “secure islands” are typically centralized closed systems that heavily rely on human supervision and control. Internet users are becoming increasingly aware of security problems such as experiencing fraudulent transactions even without having used particularly sophisticated on-line financial services. As electronic information and services are handled more automatically on behalf of the user, the user no longer knows how and what data is secure in the electronic exchanges.

On the one hand, as agent technology and the support infrastructure advances, they offer the potential to help support the enhanced security requirements of more open service environments. On the other hand, the problem of security and in particular agent security is a very multi-faceted issue that in the real world involves trade-offs, unseen variables, and imperfect implementations. Any good security design will provide a system architecture supporting the relationships between prevention, detection and reaction [1]. However, highly distributed open service systems, such as MASs currently have no coherent theory, architecture design and implementations to use even classic Internet security in a standard way. If “openness” is to be key, as it brings many advantages to the deployment of services and information, then security that covers the many needs of the environment, services and applications requires some basic security standards.

1.1 Trust, security and privacy

Current research has also demonstrated that we bring our social model to the world when we interact with various inanimate objects from the toaster to the computer within it [2]. For example, our very social and cultural approach to evaluating a first meeting of a service can be strongly influenced by someone’s recommendation if we have attributed a high-level of credibility of knowledge to a person concerning that particular service. Hence, the very success or failure of a service in the physical world could be based on someone’s recommendation. The multifaceted nature of trustworthiness requires support for generic concepts of security, and privacy [3]. These concepts can be defined within a multi-agent systems architecture as follows:

✍ *Trust*: is a social concept for evaluating risk, which is often situated in a cultural environmental and is driven by a

community's need for cooperation through communication and interactions for the perceived survival of that community. The community requires two or more players;

✍ *Security*: is a set of physical realizations that reduce the risk of danger or potential hazards when interacting with the environment. Social trust does not necessarily need to have security; however, security can provide fundamental building blocks for supporting concepts of trust. The mainstream computer network community also uses a concept of trust associated with a network of trusted third parties that are introduced in order to approve unforgeable bindings between names and objects such as public encryption keys, roles and access control lists. It is assumed that this belief in these bindings is complete by all parties. We refer to this specific concept trust as encryption trust.

✍ *Privacy*: provides both a conceptual and physical space for the social protection of high-valued items such as knowledge, information, objects, services, that a person or community places a high-value on and that these items are respected as such. Often privacy utilizes both concepts of security and trust.

The remainder of this paper is structured as follows: In section 2 the Security requirements are generated from a set of use-cases, section 3 discusses some main issues in standardizing agent security. The security related FIPA specifications are reviewed; the use of the FIPA specifications for secure MAS systems are analyzed in section 4. Finally, a discussion about future directions for standardizing MAS security concludes the paper.

2. REQUIREMENTS AND USE-CASES

Vital problem frameworks for the use of secure MASs are e-business Open Service Spaces (OSSs) that are characterized by:

- ✍ Heterogeneous service components from multiple providers;
- ✍ Dynamic service aggregation;
- ✍ Information that is distributed across insecure environments;
- ✍ Multiple autonomous domains that may become interlinked and lose some of their autonomy.

In order to illustrate some of the pertinent issues and to generate requirements for MAS security within OSSs, some security related scenarios (see Table 1) have been modeled by the FIPA Security workgroup as part of a white paper [4] that reviews the status of MAS security within FIPA.

For example, in the privacy and personalization scenario, several agent roles are identified. A personal agent A manages a person's preferences and characteristics such as tolerance to drugs, gender etc. for a human principal. A doctor service agent B provides a medical help and is able to access these preferences and characteristics in order to slant an instance of a service invocation to that agent, i.e., to treat a patient's medical condition. Other hospital agent services C, D may be used by agent B to carry-out its service and finally other personal agents E and F may also talk with agent A to find out about information about C's service.

Table 1. Some application scenarios and their main security issues.

Scenario	Security issues
Publisher/directory	Authentication, authorization, DoS,
Courier/broker scenario	Message privacy, integrity, authentication, non-repudiation
Task Allocation scenario	Non-repudiation, contract integrity, message privacy
Multi services domains scenario	Propagation of authentication, authority, trust across multiple domains
Personalization and privacy service scenario	Privacy & integrity of user preferences, privacy & integrity of service capabilities, authentication of owner, action, policy integrity & privacy, trust
Mobile agent application scenario	Agent integrity, message integrity

The following security problems can occur in this specific scenario:

- ✍ The service agent B may divulge private information (a user's personal preferences) to other service agents C and D against the wishes of the user agent A;
- ✍ The user agent A may reveal its favorable service offer to other personal agents E and F against the wishes of the service agent B;
- ✍ The identity of A's human-owner or principal may be modified so that A is associated with different characteristics and so receives an ill-matched treatment plan;
- ✍ The personal agent policy for revealing his or her preferences and characteristics to a specific agent such as a doctor agent may become compromised, e.g., the new policy is now that the user agent can reveal information to any other agent;
- ✍ Another agent, who is not qualified to offer a doctor service, may masquerade as an instance of a doctor service type;
- ✍ A may trust a particular doctor B to treat A, but B gets replaced by another instance of the doctor agent.

These simple examples are just a subset of even more complex situations that may occur in a number of various real applications and environments. Basically, the threats in the digital world mirror the threats in the physical world. However, we have systems (not perfect) in the physical world in order to provide the type of protection necessary for the type of service or situation such as trust. The open digital society provides a whole different area that socially, culturally and technically without having such systems in place.

3. Review of FIPA Agent Security

The Foundation for Intelligent Physical Agents or FIPA, a forum of international companies with a strong focus in the

telecommunication industry, was formed in 1996 to promote the uptake of software agents in businesses at large [5].

3.1 The FIPA Specifications

The first FIPA specifications were released in 1997. In 1998, FIPA first became active in specifying agent security [6]. This initial specification has since been made obsolete, but it has provided some useful hooks to model security within a FIPA agent platform.

3.2 Architectural Security Elements

The abstract architecture specification [7] covers some of the general properties for security, but it stopped short of proposing one or more (abstract) functional architectural elements for security such as secure channels or authentication services. The security concepts in the abstract architecture are summarized here: The central requirements for security that must permeate the FIPA architecture are:

- ✍ *Identity*: the ability to determine the identity of the various entities in the system.
- ✍ *Access Permissions*: based on the identity of an entity, determine what access policies apply to the entity.
- ✍ *Content Integrity Validity*: the ability to determine whether a piece of software, a message, or other data has been modified since being dispatched by its originating source.
- ✍ *Content Privacy*. The ability to ensure that only designated identities can examine software, a message or other data. To all others the information is obscured.

Security related architectural elements in other specifications are concerned more services such as:

- ✍ Message transport service: [8];
- ✍ Agent management service: [9];
- ✍ Security support services: [6].

Each of these will be discussed in turn in the following sections.

3.3 Message Transport Service

The MTS transport specification [8] specifies an optional tag called “encrypted” for defining how an ACL message can be encrypted for exchange between two agents. The use of this encryption follows the IETF RFC822 model [10]. The value of the envelope-encrypted field is optional. If it is set, it indicates that the enclosed ACL message payload is encrypted as defined in RFC822. The MTSs implemented in current FIPA platforms based on such as JADE, FIPA-OS etc. do not support RFC822.




Figure 1. Methods of Communication between Agents on different Agent Platforms as defined in the FIPA Message Transport Specification. The numbers are explained in the main text. The ACC represents the Agent Communication Channel.

Under the RCF822 scheme, messages must contain header fields as well as the contents. Consequently, it is necessary that they remain unencrypted, so that mail transport services may access them. Since names, addresses, and “Subject” field contents may contain sensitive information, this requirement limits total message privacy.

The message envelope encryption model sets the encrypted field on a per message basis. There is no higher-level abstraction to specify message security for a group of messages such as on a per session or on a per interaction sequence or with respect to a policy.

Next we consider how an agent sends a message to another agent resident on a remote AP, there are three options: (see numbered arrows in Figure 1):

1. Agent A sends the message to its local ACC using a proprietary or standard interface. The ACC then takes care of sending the message to the correct remote ACC using a suitable MTP.
2. Agent A sends the message directly to the ACC on the remote AP on which Agent B resides. This remote ACC then delivers the message to B.
3. Agent A sends the message directly to Agent B, by using a direct communication mechanism. This communication mode is not covered by FIPA.

Security for the communication is not end-to-end in the sense of being application to application. Messages are encrypted in the message transport service in the Agent Communication Channel (ACC): the transfer of the messages to the transport layer service may be unencrypted.

It is easy to eavesdrop on messages during their transfer from the agent to the ACC if they are unencrypted particularly if the message is transferred unencrypted to a remote ACC via interaction pattern 2 (Figure 1). Hence, interaction pattern 2 would not be secure.

A final issue is that the RFC822 model does not define any levels of granularity for the encryption, for example, digitally signed messages without encryption versus encryption without signing.

3.4 Agent Management Service

The current FIPA agent management specification [9] defines the concept of an agent platform, a physical infrastructure in which agents are deployed. It also specifies the use of AMS (name service and agent life-cycle management service) and DF (directory service) agents.

The main security issues are that: the AMS and DF and other agents have no credentials available for use to verify an agent's identity; the AMS directory has no access control; the DF directory has no access control; AMS registration specifies an ownership (a principal responsible for the agent) field in the service description frame of the agent management ontology – this ownership field has no integrity check and can be easily forged.

The main recommendations to improve the security of agent management security are the specification of (1) credentials for authenticating agents; and (2) access control schemes for agents such as the AMS and DF and integrity checks for the ownership field.

3.5 Agent Security Support Service

The agent management security specification was proposed [6] as a secure extension to the agent management specification (Figure 2). It defined a FIPA Agent Security Manager through which all communication passes; it enhanced the DF and AMS agent services and proposed fields in the transport envelope to set separate levels for confidentiality and integrity.

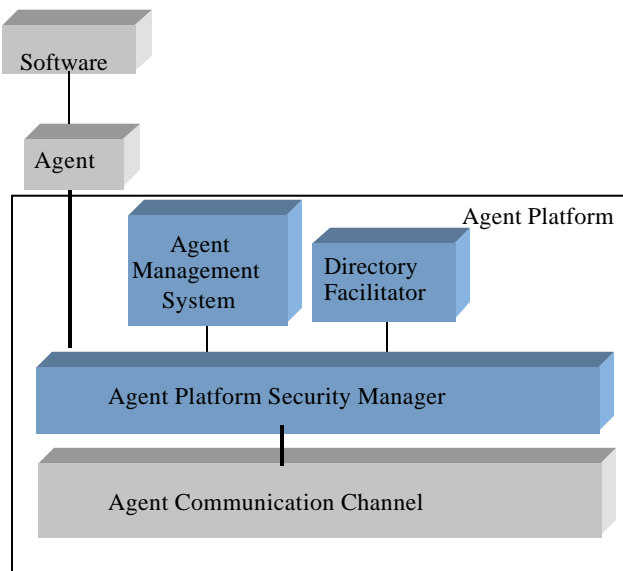


Figure 2. Proposed security extension to the FIPA Agent Management Specification (obsolete).

The strengths of the agent security specification model are:

- ✍ The specification depicts abstractions for levels of privacy and integrity that are technology independent, i.e., they are specified as high, medium or low;
- ✍ Message privacy is specified independently of message integrity;

✍ Multi-level model of confidentiality and privacy can underpin adaptive models of security, i.e., the agent can configure or reconfigure privacy and integrity according to application requirements or management policy.

This specification did not address the main recommendations to improve the security requirements of the AMS and DF mentioned earlier, mainly because FIPA at that time considered these to be an issue for the supporting infrastructure and not for the agents themselves.

3.6 Review of FIPA Security Systems

Security in agent mobility has been well researched, although no single or de facto standard has been developed. It is believed that mobile agents offer a greater opportunity for misuse and abuse [11]. This has led to the hypothesis that if we can solve the problems of mobile agent security then these solutions can be confidently applied to solve the security problems of other types (static) of agent system [12]. The main issues of security for mobile agents are that mobile agents must be protected from attacks by remote platforms and that remote platforms must be protected from attacks by mobile agents.

As MASs of communicative agents reach out more into the untrusted heterogeneous environment of other MASs, communicative agents will likely face similar threats to those threats in mobile agent systems. There are however, important differences between MASs of communicative agents and mobile agents: the protection of the agent code against code modification whilst being an obvious concern in mobile agent systems is not a major threat in MASs of communicative agents. Communicative agents are also more prone to communication threats than mobile agents. Multi-agent systems of communicative agents offer a comparable challenge to mobile agent systems, but to an extent, a different opportunity for misuse and abuse.

Whilst the current FIPA specifications contain minimal support for agent security; several researchers have reported adding security to FIPA based MASs. MASs most often reported use encryption-based mechanisms to protect systems. Two key architectural elements are added: a secure channel to provide message privacy and a certification authority (CA) to provide authentication [13] [14] and [15].

Zhang et al. [13] have added security to the FIPA-OS MAS for mobile agents and communicative agents. The security service is modeled by two agents: a Secure Agent Communication Channel (SACC) agent to perform mutual authentication, and a Negotiator Agent to negotiate about the level of encryption to be used and to exchange symmetric keys for bulk encryption. Poggi et al. [14] report a security model for the JADE (Java Agent Development) FIPA MAS. Their approach uses a Certification authority, a distributed authorization model, security models and a secure channel based on SSL. Hu [15] has used the FIPA ACL combined with PKI for authentication and uses the SPKI (Simple PKI) model for authority delegation.

4. Some thoughts on future directions for FIPA MAS Security

The following are suggested as future research areas for FIPA:

- ✍ Specifying multiple levels of security and the use of adaptable security;
- ✍ Security, trust and Privacy Policies; Modeling security at the ACL level;
- ✍ Architectural Abstractions, services and design issues for MAS security.

Multi-level security and adaptive security

It is anticipated that graded and adaptive levels of security will be needed, based on the different services or application domains and their requirements. Therefore one would have to define different groups of mechanisms that would be used in given situations. Some examples of different grades of security requirements could include:

?The choice between public but integrity verifiable messages (i.e. readable by all but with certainty that they have not been tampered with), versus encrypted as well as integrity verifiable messages (i.e. readable only by the intended recipient in addition to the certainty that they have not been modified).

?The choice between public lookups of directory information (i.e. services and registered agents available for all to see), versus authenticated lookups (i.e. lookups restricted to some privileged agents).

The requirements for a minimal level of security includes authentication; message privacy; detectable unauthorized message integrity violations and access control to key agent services.

4.1 Policies

Policies explicitly define the type of conditions a particular set of computational services will adhere to when operating in a particular context. This approach provides more openness to the service architecture as the computational services explicitly declare their intention to join a particular policy rather than this being implicitly defined within the communicative acts and protocols. Policies can be defined as a set of ontologies where the matching of policies can be done through a set of meta-constraint satisfaction rules. Examples of policies include policies for: new-user registration, error handling, information sharing, delegation policy and control. The notion of policies can be applied to various concepts within an agent architecture, such as dynamic participation in service teams [16]. More substantial work has been done in defining policies of trust [17].

4.2 ACL security

When considering the impact of security on agent communication paradigms within an agent system, we need to consider at what layers of the communication infrastructure, security should be accessible. If we consider the ACL as a set of four layers: transport level, speech-act or communicative level, ontology level and interaction protocol level, we examine the issues that should

be considered with respect to providing security at each of these levels.

For the purposes of this discussion, a conversation is the set of related communicative acts (akin to a session) that comprise an interaction between two agents, and follows a given interaction protocol. A message contains a speech act and is associated with a single utterance within an interaction, and transport is the means by which a message gets from the sender to the receiver.

4.2.1 Transport Level issues

There is already much existing work in the area of message transport between processes, especially in the context of client-server models. Our security solution should take advantage of these as much as possible. For instance, it may be possible to fold transport-level security services under the umbrella of the transport service in the abstract architecture.

With that caveat, we also mention that sending messages between agents is not necessarily relegated entirely to some existing transport, so existing transport-level security may not necessarily cover agent message-passing. For instance, agents may use email or forward messages through gateway or proxy agents. Therefore, it is not clear that relying entirely on existing transport-level security is desirable.

Finally, the lower down the network protocol stack, encryption occurs e.g., the IP layer, the less transparent it may appear to the agent. In addition, very low-level network layer encryption is not likely to be end-to-end.

4.2.2 Communicative Act issues

The addition of new communicative acts to access the security service has the advantage of simplicity. It has been proposed in several research papers, for example, [17] have proposed adding new speech acts to KQML for apply-certificate, issue-certificate, renew-certificate, update-certificate and revoke-certificate. This approach could have been adopted for agent management in the agent management specification. The disadvantage is that FIPA has resisted adding service or application specific speech acts, for example for security, in order to keep the core set of speech acts generic and to a minimum. Rather than new speech acts, an ontological approach is introduced as a powerful alternative approach.

Foner [18] was one of the first agent researchers to discuss the problem that many semantic models proposed for agent communication, require one agent to leak or reveal information about its internal state to another agent. For example, when one FIPA agent informs another agent that it is raining then the semantics of the inform communicative act require that the sender agent believe it is raining, and believe that the receiving agent does not yet believe it's raining and that after sending the message the receiving agent will come to believe it is raining. There is a trade-off in maintaining privacy versus using agent communication protocols that support rich knowledge exchange involving intentions, goals and plans. However, it is also possible to define some semantics for communication that does not depend on the sender and receiver sharing the same internal state.

4.2.3 Ontology Level

Making use of the existing FIPA speech acts and interaction protocols but referencing one or more security ontologies would minimize the changes to the existing ACL specifications to support security. It may be beneficial if FIPA seeks to reuse existing security schema from the mainstream computer network community.

4.2.4 Interaction Protocol Level

One key argument for providing security at the level of the interaction protocol is that conversations naturally provide a scope for session keys. To wit, one natural paradigm is that an agent, wishing to interact with another agent in the context of some task, can authenticate itself to that agent; the agents can then share public keys that are valid for the duration of the interaction. This may be accompanied by the negotiation of policies at the interaction level – “This interaction takes place under the umbrella of this security policy ... encryption method is ...”.

We note also that a given security implementation may have the potential to influence the interaction protocols themselves. For instance, if authentication becomes a part of every interaction among FIPA agents, this could either become some sort of a policy or could be embedded in the interaction protocols themselves. Also, the interaction with a security service may not naturally follow a pre-existing interaction protocol; therefore new interaction protocols may need to be defined for such interactions (this may be true for services in general).

5. Conclusions

If services such as automated negotiation, personalized access and local context awareness are to be supported by agent technology then security becomes necessary. It is needed to support the legal concerns for data protection, the use of personal preferences, social and moral concerns, and the general security requirements for e-business. FIPA's Security WG is currently active in these areas.

6. ACKNOWLEDGMENTS

We thank all our colleagues within FIPA who have contributed to or reviewed the security white paper. We also thank people for their response to the FIPA Security WG Request for Information, issued to the agent and security community. In addition, the FIPA Security WG wishes to thank the membership for its input during the FIPA meetings and to others for contributions to the email list. The views expressed in this article are those of the authors.

7. REFERENCES

- [1] Schneier B. *Secrets and Lies: Digital Security in a Networked World*, Wiley, (2000).
- [2] Nass C and Reeves B. *The Media Equation: How People Treat Computers, Televisions, and New Media as Real People and Places*. Cambridge University Press, (1996).
- [3] Falcone R., Singh M., and Tan Y. (Eds.) *Trust in Cyber-societies: Integrating the Human and Artificial. Perspectives*, LNAI 2246 Springer, (2001).
- [4] FIPA MAS Security White paper, reference f-out-00113, <http://www.fipa.org/repository>.
- [5] FIPA, The Foundation for Intelligent Physical Agents, Home Web-page. <http://www.fipa.org>.
- [6] FIPA 98 Part 10 Version 1.0: Agent Security Management Specification (obsolete). <http://www.fipa.org/repository/obsoletespecs.html>
- [7] FIPA Abstract Architecture Specification., Version J, <http://www.fipa.org/repository>.
- [8] FIPA Agent Message Transport Service Specification. [http://www.fipa.org/repository.FIPA Agent Management Specification](http://www.fipa.org/repository.FIPA%20Agent%20Management%20Specification). <http://www.fipa.org/repository>.
- [9] Crocker D.H. Standard for the format of ARPA Internet Text Messages. IETF Request for Comments 822.
- [10] Jansen W and Karygiannis T. *Mobile Agent Security*, National Institute of Standards and Technology Special Publication 800-19 (August 1999)
- [11] Ghanea-Hancock R, Gifford I. Top secret multi-agent systems. 1st Int. Workshop on security of mobile multi-agent systems (SEMAS-2001), 5th Int. Conf. Autonomous Agents, Montreal, Canada (2001).
- [12] Ghanea-Hancock R, Gifford I. Top secret multi-agent systems. 1st Int. Workshop on security of mobile multi-agent systems (SEMAS-2001), 5th Int. Conf. Autonomous Agents, Montreal, Canada (2001).
- [13] Zhang M, Karmouch A and Impey R. Towards a Secure Agent Platform based on FIPA. Proc. MATA 2001. Springer-Verlag. LNCS, (2001), Vol. 2164, 277-289.
- [14] Poggi A, Rimassa G and Tomaiuolo M. Multi-User and Security Support for Multi-Agent Systems. Proc. of WOA 2001 Workshop, Modena, (Sep 2001).
- [15] Hu Y-J. Some thoughts on Agent Trust And delegation. Proc. 5th Int. Conf. on Autonomous Agents, AA2000, Montreal, (2000) 489-496.
- [16] P. and Cattoni R. Evaluating the Deployment of FIPA Standards when Developing Application Services”, *International Journal of Pattern Recognition and Artificial Intelligence*, Vol. 15, No. 3, (2001) 551-577.
- [17] Foner LN. A security architecture for multi-agent match-making. Proc. ICMAS (1996).