Enabling Technologies for the Interoperable Enterprise

Jean-Henry Morin, Michel Pawlak and Manuel Oriol

Abstract. Today's Enterprise Applications are bound to achieve a high level of interoperability in order to cooperate. Actual tendency is to integrate slowly emerging technologies by relying on old ones. In this article we advocate that new emerging technologies can be used to enhance traditional interoperability techniques. These technologies comprise Digital Right Management/Digital Policy Management systems as well as mobile agents, peer-to-peer and ontologies. To test our approach, we implemented a framework where such technologies have been integrated and report briefly on our experiments

1 Introduction

Today's Enterprise Applications (Business Information Systems, BIS) are characterized by a combination of legacy and best of breed applications. The former category has evolved in the direction of ERP (Enterprise Resource planning) systems while EAI (Enterprise Application Integration) attempts to bridge the gap trying to integrate the whole. The result combined with the advent of Internet open sub-networks and open standards for data exchange has led to growing needs and capabilities towards a massively interconnected eSociety.

Along this path, it appears that the industry is slow at integrating some key emerging technologies while deploying much effort trying to promote Web Services [7] which is the current "buzzword" in the IT industry. It feels like new words and new wrapping of old technology just increasing the complexity of the current IT landscape but falling short in addressing many key and critical issues in the global business information ecosystem. Interoperability in software systems thus appears as a fundamental issue. The problem is that interoperability is not obtained for free. It requires that systems designed for interoperability use it from ground up.

We advocate that the combination and integration of some key emerging technologies can help to address this federating issue of interoperability thus enabling new forms of Business Information Systems (BIS) offering greater flexibility, increased performance and higher security. In particular, we focus on Enterprise Rights and Policy management, secure mobile agent architectures, Peer-to-Peer computing, trust computing and business semantics/ontologies, towards their integration in the emerging interoperable rights and policy enabled Enterprise landscape. In the following sections, we give an overview of these technologies and then explain experiments we have undertaken in the SIMAT [1] project.

2 Overview of the Key enabling technologies

This section is structured around two key interwoven issues. First, Enterprise Digital Rights and Policy Management and second, convergence and integration of key enabling technologies for next generation Business Information Systems.

2.1 Enterprise Digital Rights and Policy Management

Among the most notable changes and evolution the corporate environment is about to witness is probably how to deal with the recurring problem of managing, safeguarding and controlling usage of information assets wherever it resides and especially outside the corporate firewall (e.g., on laptops, memory keys or when transferred outside the administrative corporate domain, etc.) This issue is a strategic corporate wide security issue. The technology addressing this is called Digital Rights Management (DRM) and its strategic managerial dimension is called Digital Policy Management (DPM). Applying DRM to the corporate environment is now referred to and known as Enterprise DRM (eDRM). DRM allows to cryptographically associate rules (i.e., using rights expression languages) that govern the usage of content in a persistent way. As a result, content can have for example an expiry date, require leaving an audit trail upon access or simply require monitoring and metering of its usage. Any attempt to access, render, use or basically do any thing with this content requires interpreting the associated rules prior to granting (or denying) the right to do so.

We strongly believe that systems, organizations, may and will exchange data given a reasonable level of trust among them. DRM/DPM-enabled are likely to let systems trust each other, thus a higher level of interoperability can be achieved by sharing documents controlled by the DRM/DPM infrastructure.

2.1.1 Digital Rights Management: an enabling technical issue

As BIS start breaking the barriers of corporate intranets, and business processes start spanning across multiple corporate structures, the issue of persistent content protection, rule based content access and usage metering are appearing as key requirements. As a result in order to stress the difference and broaden the scope of this technology, the term of Digital Policy Management (DPM) was coined to emphasize the strategic dimension of this field. DRM/DPM and trust computing will be key components of the evolution towards next generation BIS.

Of course, this field deals with security issues. But we must distinguish between two security levels. Namely, the transport level, and the persistent protection and rule or usage based access levels. The first level is now well known and established. It is factored-in in almost all Internet based applications. It essentially provides confidentiality, authentication, integrity and non-repudiation during the transport of data among the communicating parties (i.e., while traveling over open networks such as the Internet) or within corporate firewalls. The second level deals with DRM/DPM. Persistent protection, addresses the issue of securing the content after it has reached its destination. In other words, protecting the content persistently including when

stored on persistent storage. It requires that the application accessing this content be DRM/DPM enabled in order not only to be able to decrypt the content but also to interpret the rules governing its usage. This is where trust computing appears in the picture in order to be able to guarantee a chain of trust from the rendering software down to and including the hardware ensuring that all layers are trustworthy and have not been tampered with. Recent developments in this field include the TCP Alliance led by intel (Trusted Computing Platform Alliance) [10], Microsoft's Palladium [8] operating system security initiative and its more recent flavor: Next Generation Secure Computing Base (NGSCB) [11].

One of the major problems that hampered broader and faster adoption of DRM was the lack of standards and the totally incompatible proprietary solutions that were available (e.g., Microsoft, InterTrust, ContentGuard, IBM, etc.). Recent progress in this field is extremely encouraging in particular with respect to standards. ISO has just ratified MPEG-REL (ISO/IEC 21000-5:2004) Rights Expression Language. It is based on XrML (ContentGuard) and was developed within MPEG-21 [6]. Another encouraging standard in MPEG-21 is about to be ratified by ISO, it addresses the issue of rights interoperability and semantics through RDD (Rights Data Dictionary). Such initiatives are instrumental in this field and represent a prerequisite for broader adoption and interoperability.

2.1.2 Digital Policy Management: a strategic management issue

As we briefly mentioned, policy management is the key strategic issue for the Enterprise. While transport level security is now commonplace both within corporate intranets and over the Internet, nothing is in place to address the issues of persistent protection of content and the management of policies and rights governing content use independently from where it resides. Companies need to be able to define and control who, how, there, when, what and under which condition information can be accessed and used at all time. For example financial statements and reports, design documents, technical specifications, proposals, contracts, legal documents, emails, etc. This also includes the information provided by databases and application servers, dynamically generated and which do not exist statically but are the result of specific queries also bound to usage rights and policies. For example in budget forecasting and simulations we also need to apply usage rules to the resulting reports generated by simulation tools.

DRM has now become mainstream technology addressing these issues shaping the future of corporate content management. Information is a corporate asset and is therefore bound to corporate policies. Today however no technical means are in place to enforce this, thus providing upfront prevention of disclosure or misuse (accidental or malicious) of this content once it has reached a laptop or a removable media such as a CD.

Global corporate information asset management is among the next major challenges facing the enterprise and their Chief Security / Information / Compliance Officers. Their role combined with these technologies in a looming regulatory environment will be instrumental in defining and managing the policies and rules persistently governing the use and access to corporate data and processes. The problem here is that there are basically three levels to be considered: the legal

environment, specific regulatory frameworks often sector bound and internal corporate policies none of which are today instrumented. Their specifications often reside in dusty books and their implementation often left to rule of thumb and experience. Consequently, this field is cruelly lacking models to capture, specify, express, represent, manage these policies prior to any technical DRM project and deployment in a corporate environment. As a result, Digital Policy Management is of strategic nature and must be initiated and driven by corporate managers and not IT / IS people. It is exactly at this point that Enterprise DRM meets Enterprise DPM thus requiring to address the issue in an interdisciplinary space between technology and management science.

2.2 Key enabling technologies, convergence and integration issues

Most of the time, enterprise modeling is centered on the notion of components that constitute the basic building blocks of the targeted application/information system. This implies that an important part is inter-components communication and another part is composition.

In this subsection, we briefly sketch several key enabling technologies: secure mobile agents, peer-to-peer infrastructures and semantic/ontological descriptions. Note that those enabling technologies are respectively located at three different levels of the described systems: component level, inter-component communication and composition. This maps a scale going from small granularity to coarse granularity in the modeling.

2.2.1Secure Mobile Agents

A mobile agent is a program that can have its site of execution changed during the course of its lifetime. An agent encapsulates data and behavior and can be sent to any network site to fulfill a task specified by its owner. This is a different approach to that taken by standard middleware and Web Services architectures, which use a client-server approach where the client sends a request to a server and awaits a response. This exchange can develop into a dialogue between the client and server over the network. In the agent approach, the agent is programmed to encapsulate client logic. The agent is then sent to a remote node - where the agent executes as a normal program - and the “client-server” dialogue takes place locally at the remote site between the “server” and the agent.

The mobile agent paradigm has several advantages for Internet applications. First, since an agent encapsulates behavior, new application functionality can be dynamically distributed and linked to clients and servers via agents. Second, migrating client programs towards servers, or server programs towards clients can reduce the amount of network traffic generated by the application. In some cases the client may even disconnect himself from the network while his agent operates on his behalf at a server. This property is particularly useful in today's operating environments where disconnection may be a feature while processes continue to be active and running as background tasks. The ultimate goal of agent technology is that an agent becomes capable of acting autonomously on a user's behalf, in the same way

that a human agent (like a travel or sales agent) acts autonomously on behalf of others.

Up until now, a critical lack of agent security has hampered the adoption of agent technology in industrial applications. Allowing the possibility to users of executing foreign code on their sites opens these sites to the risk of viruses and other unwanted behaviors. Clearly, software platforms still have difficulty in controlling the data manipulations of programs that they run. It is for this reason that a mobile agent platform must ensure that mobile agents execute in secure isolation from the rest of the host platform.

Agents also represent a true opportunity as a programming paradigm shift. Namely, the frontier between the corporate intranet and the world at large is fading towards the virtual enterprise where business partners and processes are reconfigured on an ongoing basis. It is thus, by using mobile agents, to have scenarios where interoperability is not achieved through message passing but rather through computation units exchanges. This is completely different from using messages to let interoperability happen as usually occurs. This lead opens a new possibility to interoperability platforms designers.

2.2.2 Peer-to-Peer Computing, Protocols and Architectures

Peer-to-Peer is more an architectural paradigm than a technology per se. Like client-server architectures when first introduced they don't refer to specific technologies but rather to models of interoperation for network oriented applications. Peer-to-Peer appears to be the natural evolution of client-server fuelled by the need to depart from a world where actors are categorized once and for all as being either "clients" or "servers". As a matter of fact, given today's IT environment this distinction has blurred giving place to a networked ecosystem (society) where the actors are nodes which are "equal" and can eventually be replaced. Moreover, the need for a common network agnostic abstraction (i.e., independent of the network transport layer and topology) has emerged. Connectivity is a resource and it is assumed nowadays. When developing productivity tools and applications for BIS, no one really cares to know how to workout the underlying network transport details such as sockets, NATs and other network engineering details.

Peer-to-Peer Computing departs from these technical details providing a common abstraction (overlay) allowing peers to interoperate naturally by providing services for lookup, discovery, organizing in groups, credential management, group and individual services, advertising, monitoring, etc. Such architectures rely heavily on the use of asynchronous messaging and open communication protocols thus enabling disconnected and autonomous operation among the peers. Project JXTA [2] is a noteworthy example of such systems that has the advantage of being a protocol-based approach. Finally, peer-to-peer computing appears to have brought to the application layer the simple constructs that served to design the Internet protocols at the time. Available communication primitives are open and allow the different systems to interoperate in a simple way.

2.2.3 Semantics and Ontologies

Ontologies have recently received significant attention particularly in the BIS community because they hold a promise towards the ability of processing the semantics rather than the "plumbing". The W3C is working in these domains with the initiative on the Semantic Web [4] especially with RDF (Resource Description Framework) [5] and OWL (Web Ontology Language) [9]. MPEG-21 has finalized a standard: Rights Data Dictionary (RDD) [6] currently under review for standardization at the ISO.

In the scope of BIS and DRM, being able to define, share, map and use common and agreed upon semantics will be critical in the near future due to a necessary high degree of interoperation. Such metadata will not only require domain specific ontologies but also standards for their processing. To this extent, RDF and OWL provide promising frameworks based on XML. Ultimately, the goal is to ease automated information exchange among actors at the semantic level and thus enable dynamic application interoperability at the data level. This is a key requirement for next generation BIS coupled with DRM that need to capture the nature and dynamics of trading partners.

3 Experiment report: The SIMAT project

The objective of the SIMAT project was to evaluate the use and acceptance of secure mobile agent technology in the context of middleware services for both Web and virtual enterprise applications. More specifically, one part of the project aimed at defining a set of business scenarios and to explore the integration of key emerging technologies. By encapsulating business objects and their corresponding behavior into mobile agents, the resulting Active Business Objects become part of processes that go through organizations, traveling across the various activities required to successfully fulfill their tasks. Furthermore, they become information providers able to release relevant information about their own usage, execution and state to the information. The next subsection describes main aspects of the SIMAT architecture (Figure 1).

3.1 SIMAT Architecture

Starting from bottom up, as hardware architectures and operating environments are likely to remain heterogeneous, platform independence is required through a common abstraction. To this extent, the Java Virtual Machine has proven to be a way of dealing rather successfully with this issue.

Moving up the architecture, we designed a secure mobile agent platform. The architecture is using the Java based implementation of Lana secure mobile agents platform [3]. It offers several key security properties, the most important being that a Lana agent executing on a host cannot gain access to resources on that machine nor to other agents' data, unless explicitly authorized by the security policy of that host or agent.

The DRM/DPM component is placed above the secure mobile agent framework. However, considering that this issue also relates to trust computing, it is important to note that it should probably span the whole architecture with hooks at all levels down to and including the hardware.

On the communication side, the SIMAT project uses the Java reference implementation of JXTA peer-to-peer protocols for the dynamic lookup and discovery of Lana secure agent platforms.

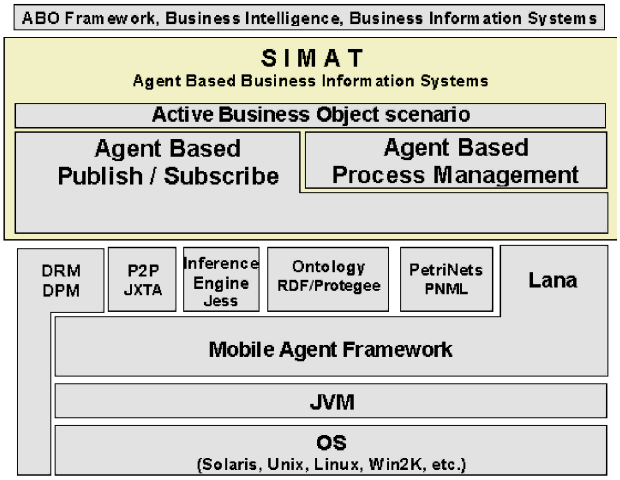

Figure 1. Overall SIMAT Architecture

As we wanted to provide an infrastructure that could allow easy adaptation to external changes, we integrated the Jess (Java Expert System Shell) inference engine, RDF as a lightweight ontology system and PNML as the interchange format for Petri nets to structure the processes. This unique combination of technologies served as the basis for the development of our framework.

3.2 SIMAT Scenarios

Three incremental Business Information System scenarios were designed and implemented to evaluate and assess the Lana secure agent platform as well as its related agent paradigm and technologies. In this context, Mobile Agents were used to encapsulate business objects together with their specific behaviors and security requirements. We have defined and implemented these scenarios as incremental stages, each of which supports the next.

The first scenario was a distributed agent-based publish-subscribe service. Among the major shortcomings that we wanted to address using mobile agents, peer-to-peer and inference engines are several issues such as scalability, security, availability, personalization, dynamic behavior, etc. Such issues required careful attention in the scope of novel architectures for business information systems.

The second scenario enhanced the first one to include distributed agent-based process management. The result showed an agent based distributed process automation system where processes were represented using colored Petri nets.

Finally, the third scenario put the pieces together and aimed at showing a preliminary scenario setting the groundwork for the development of an Active Business Object framework. In this scenario, we built upon the two previous scenarios using the agent based distributed process and the agent based publish-subscribe services. To illustrate this scenario, we have considered a simple ABO based scenario involving several actors in a distributed negotiation-decision process. This scenario also involved circulating content as support information in the process

3.3 SIMAT Insights

Three secure agent based Business Information System scenarios were designed and implemented. Thus providing valuable insights and results with respect to both the feasibility of the approach and the soundness of using secure mobile agents as a new programming paradigm as well as a technical solution in this field.

While there remains uncertainty with respect to the timing for broad recognition and adoption of secure mobile agents, the project has clearly revealed that it represents a major paradigm shift in software engineering and a truly innovative approach to next generation distributed Business Information Systems and electronic services.

Project results have been presented and discussed in both academic and industrial settings such as the 5th International Conference on Enterprise Information Systems and among industrial contacts and partners.

However, the experience acquired during SIMAT project also showed us that the integration of inference engines and Petri nets to the architecture was not as useful at it seemed to be at first. The significant advantages, in the scope of BIS, that these technologies were bringing were annihilated by the effort needed to switch to radically different conceptual design. Actually, one of the goals of the SIMAT architecture was to integrate in a unique overview principal business preoccupation. If DRM/DPM, P2P networks, ontologies and the need for security and mobility represent a high level technical abstraction easily reflecting existing issues, both Petri nets and inference engines represent formal solutions leading to inconvenient frequent paradigm switching

4 Conclusions and Future Directions

The field of Business Information System and Business Intelligence applications builds on several innovations and advances drawn from technology and management. In particular, the emergence of the Internet as a medium for corporate interoperation and data exchange is a fact that changed the way Enterprise systems are now built. In particular, the need for interoperability has grown over the time. At this need has leveraged the need for integration, mobility, ubiquity, security, trust, etc.

These issues represent the core of the background leading tomorrow's architectures in this field. In this paper we advocate that all those points are useful and may be used in the context of enterprise systems and interoperability. One of the first steps toward such approaches was the to build the SIMAT infrastructure around those key technologies.

Future step is to define an abstraction able to give an overview of enterprise needs and aiming at helping modeling interoperable enterprise systems. The resulting model should represent a global view putting together existing standards and protocols in order to be as much as possible independent from any platform, language or proprietary technology. We believe that such abstraction model has to be built around the key enabling technologies described in this paper while being flexible enough to allow future integration of new key concepts.

References

[1] J.-H. Morin and J. Sievering, “Towards Agent Based Business Information Systems and Process Management”, in Proceedings of the 5th International Conference on Enterprise Information Systems, ICEIS 2003, Angers, France, April 22-26, 2003, pp. 288-295.

[2] L. Gong, Project JXTA: A Technical Overview, Technical Report, Sun Microsystems, April 2001.

[3] Bryce, C. Razafimahefa and M. Pawlak, Lana: An Approach to Programming Autonomous Systems, in ECOOP 2002 -- Object-Oriented Programming, Vol. 3-540-43759-2, pp. 281-308, 16th European Conference on Object-Oriented Programming, Malaga, Spain, June 10-14 LNCS 2374, 2002.

[4] T. Berners-Lee and E. Miller, The Semantic Web lifts off, in ERCIM News No. 51, 2002.

[5] W3C, Resource Description Framework, (Miller, E.). Retrieved May 2004, from http://www.w3.org/RDF/

[6] MPEG-21, Retrieved May 2004, from http://www.itscj.ipsj.or.jp/sc29/29w42911.htm#MPEG-21

[7] W3C, Web Services Activity, (Haas, H.). Retrieved May 2004, from http://www.w3.org/2002/ws/

[8] A. Carroll, M. Juarez, J. Polk and T. Leininger, "Palladium": A Business Overview. Retrieved Oct 2002, from http://www.microsoft.com /PressPass/features/2002/jul02/0724palladiumwp.asp

[9] W3C, Web Ontology Language, Retrieved May 2004, from http://www.w3.org/2001/sw/WebOnt/

[10] Trusted Computing Platform Alliance, Retrieved May 2004, from http://www.trustedcomputing.org/

[11] Microsoft Corporation, Next Generation Secure Computing Base, http://www.microsoft.com/ngscb/