ARTEMIS: Towards a Secure Interoperability Infrastructure for Healthcare Information Systems

Mike Boniface, Paul Wilken
IT Innovation Centre, University of Southampton
2 Venture Road, Chilworth Science Park, Southampton SO16 7NP, UK
Corresponding author: Mike Boniface, tel: +44 23 8076 0834, fax: 44 23 8076 0833, email: mjb@it-innovation.soton.ac.uk

Abstract. The ARTEMIS project is developing a semantic web service based P2P interoperability infrastructure for healthcare information systems. The strict legislative framework in which these systems are deployed means that the interoperability of security and privacy mechanisms is an important requirement in supporting communication of electronic healthcare records across organisation boundaries. In ARTEMIS, healthcare providers define semantically annotated security and privacy policies for web services based on organisational requirements. The ARTEMIS mediator uses these semantic web service descriptions to broker between organisational policies by reasoning over security and clinical concept ontologies.

Keywords. Healthcare information systems, security, semantic interoperability, web services, P2P

1. Introduction

A typical healthcare provider will use many heterogeneous healthcare information systems to support the delivery of patient care each designed to perform specific function. Typically, these systems are standalone, developed by many different suppliers and are incompatible with one another. The non-interoperability of IT systems represents the biggest single problem in transferring data securely between different parts of a healthcare system [22].

In the ARTEMIS project [1], [9] we are developing a semantic web service based P2P interoperability infrastructure for healthcare information systems that will support new ways of providing health and social care. Healthcare providers join an ARTEMIS

network to access medical web services that enables access to electronic healthcare records maintained by other healthcare organisations. We use ontologies, derived from existing healthcare standards, to describe the semantics of web service operations and data [4]. Using these semantic service descriptions we provide ARTEMIS mediation super peers that enable heterogeneous healthcare information systems to interoperate.

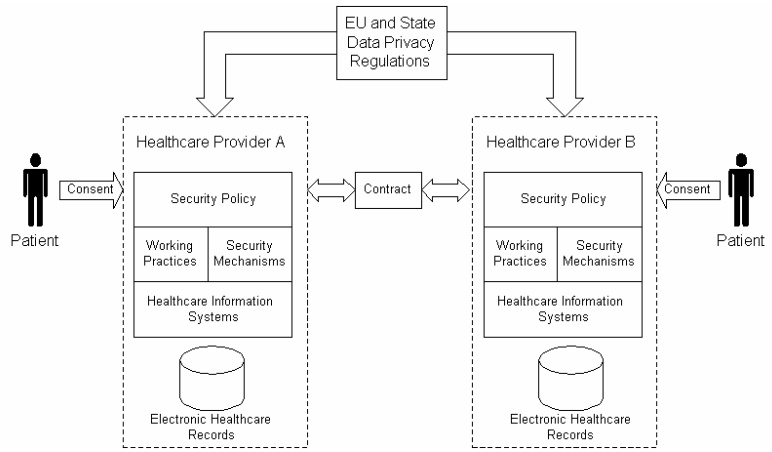

Figure 1: Data privacy regulatory framework

Describing the functional characteristics of services is only part of the story. Resolving non-functional service requirements such as security and privacy are also essential for interoperation. Healthcare information systems operate within a strict regulatory framework that is enforced to ensure the protection of personal data against processing and outlines conditions and rules in which processing is allowed (See Figure 1). There are many such regulations at European level [3],[10] and additional legislation implemented within member states [6]. According to EU Directive 95/46/EC, if a healthcare provider maintains personal data on its patients the healthcare provider is identified as a data controller and is responsible for protecting that data against unauthorised use. Typically, a healthcare provider implements the legislation by authoring a security policy that mandates working practices and security technology requirements (key sizes, algorithms). If a healthcare provider wants to access personal data within another organisation they are identified as a data processor. For the communication to occur between data controller and data processor consent must be obtained from the patient and a contract between the two parties must exist that defines conditions such as the type of data processing and how long the data can be stored by the data processor.

After the out-of-bound legislative conditions for data processing have been agreed there are still technical challenges in terms of security and privacy mechanisms that need to be resolved before electronic healthcare records can be automatically shared between healthcare information systems. In most cases healthcare providers have different security policies that state a diverse set of security requirements and capabilities. Authentication and authorisation mechanisms for healthcare professions

may also be different. In this paper, we describe the ARTEMIS architecture and an approach for mediating between security and privacy policies using a combination of industry supported web service standards and reasoning over semantics web service descriptions.

2. Web service standards and interoperability

Web services have promised to provide a solution to complex interoperability problems through the use of open standards developed by organisations such as the W3C [25] and OASIS [20]. However, integrating heterogeneous systems based on standards does not equate directly to interoperability. The first difficulty is that the standards themselves can be complex, interpreted in different ways and implementations can provide different levels of compliance. In addition, the recent proliferation of sometimes competing web service standards for security and privacy such as WS-Security, WS-SecurityPolicy, WS-Authorisation, WS-Privacy, WS-Trust and WS-SecureConversation only increases the possibility of incompatible systems.

In the healthcare sector, where numerous standards already exist, this problem is well known and initiatives such as IHE [12] dictate how complex standards such as HL7 [11] and DICOM [8] should be utilized when implementing hospital workflows. In the web service community, to ensure some level of interoperability, leading vendors formed a group called WS-Interoperability (WS-I) [27]. WS-I defines so-called "profiles" – constrained ways to use Web Service standards. For example, WS-I Basic Profile 1.0 [28] specifies that only certain transport protocols should be used (even though WSDL [26] can accommodate others), so that vendors don't have to implement all possible protocols in their frameworks. Even though many web service standards exist, only WS-I Basic Profile 1.0 (soon to be joined by Basic Security Profile 1.0 [29]), are widely supported by vendor's toolkits [13], [14].

The second barrier to interoperability is that current standards have focused on syntactic issues, which in most cases still require human readable specifications for service integration. To improve interoperability semantics are needed to allow software to understand the meaning of data and a service's function allowing improved service discovery, automated orchestration and mediation. The importance of semantics for interoperability is well documented [5], however, semantic web technologies are not currently mainstream with little adoption by leading vendors.

The level of support for web service standards and semantic web technologies in existing industrial toolkits is a key constraint for the integration of healthcare information systems (HIS) with the ARTEMIS infrastructure. If the toolkits do not support the standards and technologies the cost of joining could be prohibitive for HIS vendors. In ARTEMIS, we take a pragmatic approach by requiring only web services that conform to industrial implemented standards and manage semantics within the middleware.

3. Web Service Descriptions

ARTEMIS middleware provides tools for authoring web services to provide access to existing healthcare information system services and electronic healthcare records. The discovery and advertisement of services by healthcare organisations are managed through a Peer-to-Peer network structure: each health information system is represented by an Artemis Peer node that communicates directly with an 'ARTEMIS Mediator' or 'Superpeer.' The web service descriptions are held at the mediator using standard web service repositories, however, the service descriptions are annotated with semantics so that the service operation, meaning of data and non-functional requirements can be understood by the infrastructure.

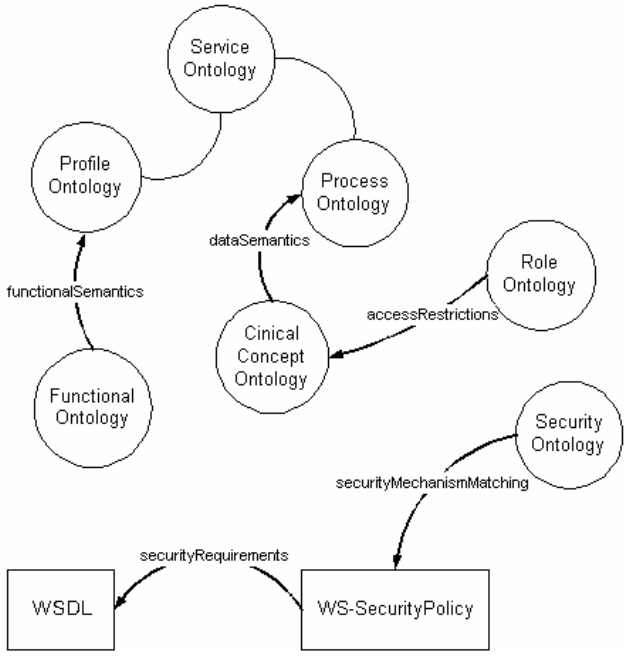

Figure 2: ARTEMIS web service descriptions

Web services advertised on an ARTEMIS network are described using standard WSDL and WS-SecurityPolicy [30] annotated with semantics as shown in Figure 2. The semantic descriptions are stored within the ARTEMIS middleware and are structured using the OWL-S [21] ontology augmented with medical data and security/privacy ontologies. Services are functionally classified using the OWL-S service profile that is extended with functional concepts derived from HL7 trigger events. The data semantics are represented using the OWL-S process model. The input and output process model parameters are associated with concepts from a clinical concept ontology (CCO). The CCO has been derived from the UMLS semantic network [18] and metathesaurus [19] to provide a rich set of terminology for describing medical data semantics. The use of HL7 and UMLS is not mandatory, as

the ARTEMIS infrastructure can support arbitrary functional and clinical concept ontologies required by healthcare providers.

4. Security and Privacy Policy Mediation

A core requirement in ARTEMIS is for very robust, but highly flexible approach to security and privacy. The approach supported by ARTEMIS is to allow healthcare providers to codify their particular preferences and requirements for data security (confidentiality, integrity) and privacy (authorisation and anonymisation) in accordance with overarching organisational security policies. Healthcare providers exposing web services require requesters to conform to certain security requirements. For example, a security policy may state that the requester is authenticated using X509 certificates or SAML assertions and that data integrity is verified using SHA1 digital signature algorithm. However, in practice both requester and provider may have different security requirements and capabilities and to achieve interoperability brokering between security policies may be required.

Figure 3 shows how the ARTEMIS infrastructure supports mediation between security policies. Healthcare providers create standard WS-SecurityPolicy's that define the security requirements and capabilities. The security policies contain references to standard identifiers for algorithms such as *http://www.w3.org/2000/09/xmldsig#rsa-sha1* for RSA-SHA1 digital signature algorithm. These identifiers are mapped to concepts within security ontologies that enable reasoning between various credential and security mechanisms at a semantic level. Ontologies already exist for this purpose [7].

When a web service is invoked, the ARTEMIS mediators act as brokers between requester and provider. For example, in Figure 3, healthcare organisation B may advertise a web service requiring Triple-DES encryption with chain blocking cipher (CBC). These security requirements are specified using WS-SecurityPolicy assertions as shown below:

```
<wsp:Policy xmlns:wsp="..." xmlns:wsse="...">
<wsse:SecurityToken wsp:Usage="wsp:Required">
  <wsse:TokenType>wsse:x509v3</wsse:TokenType>
</wsse:SecurityToken>
<wsse:Confidentiality wsp:Usage=”wsp:Required”>
       <wsse:Algorithm Type-“wsse:AlgEncryption”
              URI=http://www.w3.org/2001/04/xmlenc#3des-cbc
       ..
</wsse:Confidentiality>
..
</wsp:Policy>
```

If Healthcare organisation A wants to invoke the service using existing web service standards, Triple-DES would have to be supported. However, ARTEMIS allows request and provider to support different encryption mechanisms by allowing

compatibility between algorithms to be expressed using ontology. In this case, healthcare organisation A only supports AES encryption. The service can still be invoked as an ARTEMIS mediator can infer that AES encryption offers a superior level of protection when compared to Triple-DES. If we look at the message sequence, (A) encrypts the request message using AES and sends it to its mediator super peer. The mediator forwards the message to the mediator at (B). Mediator (B) decrypts the request message and re-encrypts using Triple-DES before invoking the service at (B). This process is then reversed when the response message is returned from (B) to (A).

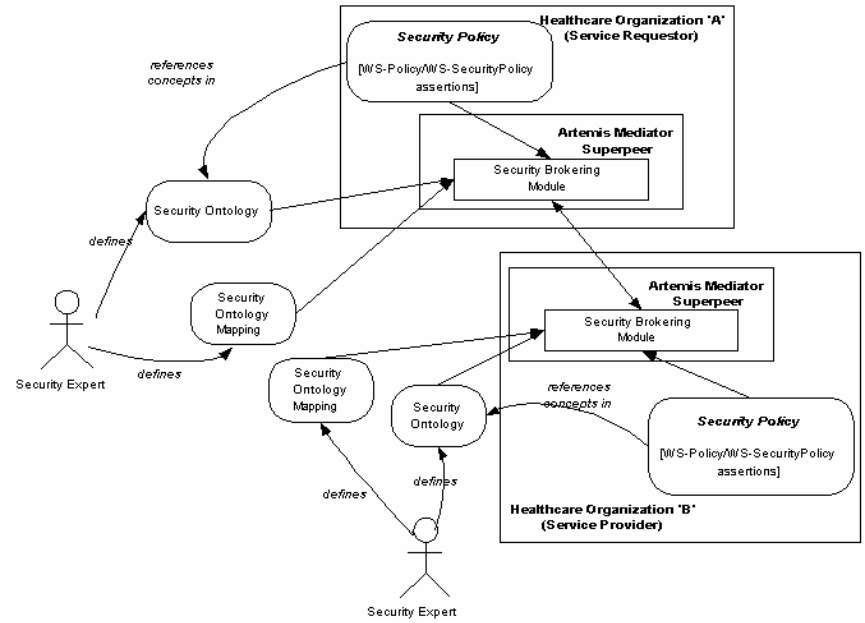

Figure 3: Security policy brokering

In addition to specifying security policies, healthcare providers need to restrict access to medical data to only authorised healthcare professionals. In ARTEMIS, healthcare providers define privacy policies that state which healthcare professionals are able to access specific medical data. Our initial approach is to allow a healthcare provider to develop role ontology that defines the clinical occupations for healthcare professionals within their organisation. These roles are then attached to concepts in the clinical concept ontology. As medical data is described using the clinical concepts, authorisation is enforced based on the role of the healthcare professional and the clinical concept being accessed. The ARTEMIS mediator can broker between privacy policies using ontology mappings linking organisational role ontologies with clinical concept ontologies. When a web service is invoked the mediator translates the role of the healthcare professional in the requesting organisation to the equivalent role in the providing organisation. Authorisation decisions are presented to the web service using SAML assertions that are signed by a trusted mediator.

Roles based authorisation is insufficient to model access restrictions in all but simple healthcare scenarios. Therefore, future work will investigate workflow context access control, where authorisation decisions are based on the healthcare professional and the context in which they are accessing data.

5. Pilot Application

An ARTEMIS pilot application is being clinically deployed by healthcare providers located in two European countries to demonstrate the interoperability of healthcare information systems across organizational and country boundaries. The pilot application includes healthcare providers South East Belfast Healthcare Trust (SEBT) in Belfast, Northern Ireland and Hacettepe Hospital in Ankara, Turkey. Each healthcare provider operates within distinct legislative domains and has different healthcare information systems to support patient care.

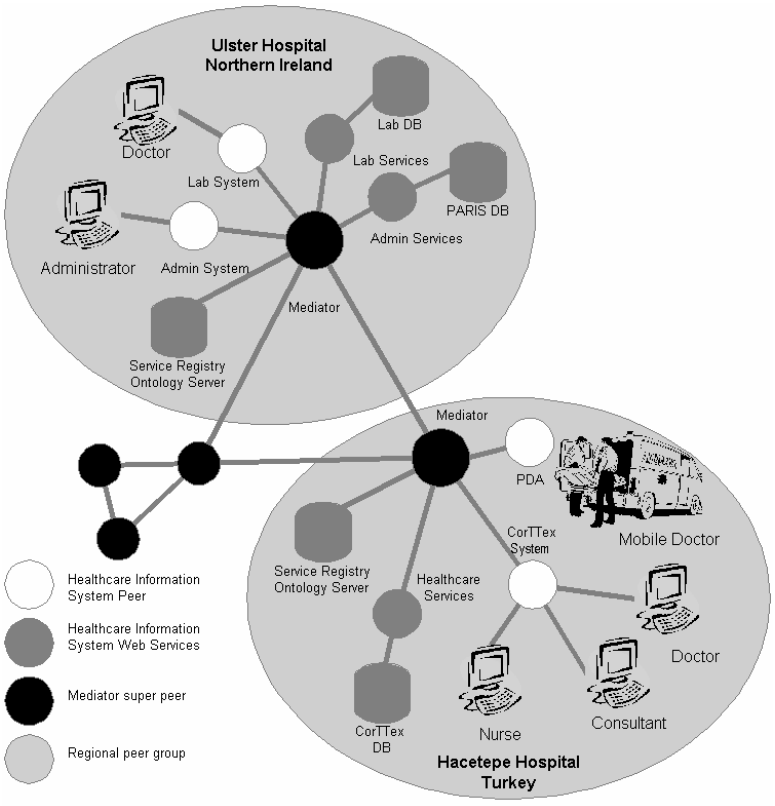

Figure 4: ARTEMIS pilot application deployment

The clinical scenario focuses a Belfast businessman on a 5-day business trip to Ankara in Turkey. The businessman has an ongoing heart condition and is receiving treatment in Belfast from the SEBT District Nursing team following an episode in a Belfast Acute hospital. During the meeting, the man suddenly becomes ill and is admitted to Hacettepe Hospital in Ankara. The doctor knows the patient is from Northern Ireland and searches the ARTEMIS network for healthcare providers within Northern Ireland that hold records for the patient [2]. The doctor receives a response from SEBT stating that records for the patient are held. Once patient records have been located, the doctors at Hacettepe hospital negotiate, out-of-band, access to the patient data with the SEBT data guardian (a person responsible for authorizing access to patient records). The data guardian then authorizes the doctor from Hacettepe hospital to access the businessman's healthcare records by updating the ARTEMIS privacy policy. The doctor then directly requests specific documents such as the assessment, case notes and medication record from SEBT using the Retrieve Information for Display service provided by the ARTEMIS Network [14].

Figure 4 shows the deployment of the core ARTEMIS components to support the client scenario. Each healthcare provider operates autonomously within a distinct region represented by an ARTEMIS peer group. The autonomy inherent in P2P networks complements the need for healthcare providers to maintain regional policies for the privacy of medical data. Each peer group contains a mediator peer that provides semantic interoperability and routing, a semantic service registry that stores web service advertisements, client peers that represent healthcare information system clients and semantically annotated web services that provide access to electronic healthcare records. Web services have been developed for the CorTTex [4] information system deployed at Hacettepe hospital and the PARIS [23] information system deployed within SEBT. The web services for both systems were semantically annotated using the ARTEMIS semantic annotation tools and advertised in a semantic web service registry.

The ARTEMIS middleware is based on open source technology. Each peer is deployed within a standard Tomcat/Axis web service container that we assume is hosted behind a firewall with restricted access on only well-known ports http and https. Connectivity with the P2P network is achieved using JXTA [15], an open source technology supported and managed by Sun Microsystems that defines a set of XML protocols that implement typical P2P functionalities. We have extended JXTA to support policy-driven secure P2P messaging using XML Encryption and XML Signature. We have also developed Axis web service handlers that implement WS-Security driven from security requirements defined in WS-SecurityPolicy documents. The handlers provide interoperability with .NET [17]or Java web services based on the WS-I Security Profile. The semantic web service descriptions are stored at the mediator within the Sesame ontology server [24]. Security and privacy policy mediation has been implemented using a semantic processor based on the MAFRA toolkit [16].

6. Conclusions

In this paper, we presented a secure interoperability infrastructure for healthcare information systems being developed in the ARTEMIS project. The architecture enables the communication of medical data across healthcare provider boundaries through mediation between semantic security and privacy policies on the condition that out-of-band contracts and patient consent has been agreed. The use of semantic web service descriptions enables differences in security requirements and capabilities to be resolved. The initial authorisation infrastructure allows access control based on mediation between clinical roles defined by different organisations. Role based authorisation is insufficient for most clinical scenarios and future work will look to incorporate workflow context into authorisation decisions.

Acknowledgements

The ARTEMIS project has received research funding from the EC's Sixth Framework Programme (project IST-2103 STP under the eHealth Action Line of the Information Society Technologies Programme).

References

[1] Aden T, Eichelberg M, Thoben W, "A fault-tolerant cryptographic protocol for patient record requests", Proceedings of EuroPACS-MIR 2004

[2] Artemis Project, http://www.srdc.metu.edu.tr/webpage/projects/

[3] Council Of Europe – Committee of Ministers, Recommendation No. R(97)5 of The Committee Of Ministers to Member States on the Protection Of Medical Data, Council of Europe Publishing, Strasbourg, 12 February 1997

[4] CorTTex, http://www.corttex.nl/

[5] Paul Cowles, "Web Services and the Semantic Web: Making information accessible and usable", Web Services Journal, Volume 02 Issue 12, http://www.sys-con.com/webservices/article.cfm?id=419

[6] Data Protection Act 1998, UK parliament, http://www.hmso.gov.uk/acts/acts1998/19980029.htm

[7] Denker, G. Kangal, L. "Security Annotation for DAML web services", In Proc. 2nd International Semantic Web Conference (ISWC2003), Sanibel Island, Florida, USA, October 20-23, 2003.

[8] Digital Imaging and Communications in Medicine (DICOM), http://medical.nema.org/

[9] Dogac, A., Laleci, G., Kirbas, S., Kabak, Y., Sinir, S., Yildiz, A., "Artemis: Deploying Semantically Enriched Web Services in the Healthcare Domain", Information Systems Journal (Elsevier) special issue on Semantic Web and Web Services, accepted for publication

[10] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, OJ L, 23 Nov. 1995, http://europa.eu.int/com-m/internal_market/privacy/law_en.htm
[11] Health Level 7 (HL7), http://www.hl7.org
[12] HIMSS and RSNA, Integrating the Healthcare Enterprise (IHE) – IT Infrastructure Technical Framework - Volume 1, ITI TF-1 Integration Profiles, Rev. 1.0, http://www.rsna.org/IHE/tf/ihe_tf_index.shtml
[13] IBM WebSphere Application Developer (WSAD) http://www-306.ibm.com/software/awdtools/studioappdev/,
[14] Integrating the Healthcare Enterprise, Integration Profiles, Volume 1 http://www.rsna.org/IHE/tf/ihe_iti_tf_1.1_vol1_FT.pdf
[15] JXTA, http://www.jxta.org/
[16] MAFRA Toolkit, http://sourceforge.net/projects/mafra-toolkit/
[17] Microsoft .NET Framework 1.1, http://msdn.microsoft.com/netframework/
[18] National Library of Medicine, Unified Medical Language System, Semantic Network, http://www.nlm.nih.gov/research/umls/meta3.html
[19] National Library of Medicine, Unified Medical Language System, Metathesaurus, http://www.nlm.nih.gov/research/umls/meta2.html
[20] OASIS, http://www.oasis-open.org
[21] OWL-S 1.0, http://www.daml.org/services/owl-s/1.0/
[22] N. Outram et al, Workshop on Biomedical Informatics and collaboration with the Research Infrastructure projects Report, 18-19 March, Brussels
[23] PARIS, http://www.in4tek.com/
[24] Sesame, http://www.openrdf.org/
[25] World Wide Web Consortium (W3C), http://www.w3.org/
[26] Web Service Description Language (WSDL), http://www.w3.org/TR/wsdl
[27] WS-Interoperability (WS-I), http://ws-i.org/
[28] WS-I Basic Profile 1.0, http://www.ws-i.org/Profiles/BasicProfile-1.0-2004-04-16.html
[29] WS-I Basic Security Profile 1.0, http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0-2004-05-12.html
[30] WS-SecurityPolicy 1.0, http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnglobspec/html/ws-securitypolicy.asp