International Journal of Medical Informatics (2004) **73**, 251—257

International Journal of
Medical
Informatics

www.intl.elsevierhealth.com/journals/ijmi

Authorisation and access control for electronic health record systems

Bernd Blobel*

Faculty of Medicine, Institute of Biometrics and Medical Informatics, Otto-von-Guericke University Magdeburg, Leipziger Street 44, D-39120 Magdeburg, Germany

KEYWORDS
Electronic healthcare record;
Security;
Privilege management;
Role management;
Authorisation;
Access control;
Enhanced TTP services

Summary Enabling the shared care paradigm, centralised or even decentralised electronic health record (EHR) systems increasingly become core applications in hospital information systems and health networks. For realising multipurpose use and reuse as well as inter-operability at knowledge level, EHR have to meet special architectural requirements. The component-oriented and model-based architecture should meet international standards. Especially in extended health networks realising inter-organisational communication and co-operation, authorisation cannot be organised at user level anymore. Therefore, models, methods and tools must be established to allow formal and structured policy definition, policy agreements, role definition, authorisation and access control.

Based on the author's international engagement in EHR architecture and security standards referring to the revision of CEN ENV 13606, the GEHR/open EHR approach, HL7 and CORBA, models for health-specific and EHR-related roles, for authorisation management and access control have been developed. The basic concept is the separation of structural roles defining organisational entity-to-entity relationships and enabling specific acts on the one hand, and functional roles bound to specific activities and realising rights and duties on the other hand. Aggregation of organisational, functional, informational and technological components follows specific rules. Using UML and XML, the principles as well as some examples for analysis, design, implementation and maintenance of policy and authorisation management as well as access control have been practically implemented.

1. Introduction

Meeting the challenge of enhancing efficiency and quality of healthcare, measures and outcome of procedures for diagnosis and therapy must be documented, communicated and evaluated carefully. In that context, co-operative and distributed health information and communication systems have to be implemented. Dominant examples for such shared care information systems are electronic healthcare record (EHCR) systems containing all relevant medical information as well as related non-medical information (e.g., materials or billing) derived from the former one. Therefore, healthcare records are the informational basis for any communication and co-operation in, and between, healthcare establishments (HCE). Including non-healthcare processes, the EHCR moves towards an electronic health record (EHR). For providing information and functionality needed, EHR must be structured and operating appropriately. In that context,

*Tel.: +49-391-6713542; fax: +49-391-6713536.
E-mail address: bernd.blobel@mrz.uni-magdeburg.de (B. Blobel).
URL: http://www.med.uni-magdeburg.de/fme/institute/ibmi.

1386-5056/$ — see front matter © 2004 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.ijmedinf.2003.11.018

trustworthiness in doctor—patient as well as in doctor—doctor communication and co-operation based on the consent of the informed patient are basic requirements.

Considering patients' care in job-sharing healthcare establishments such as hospitals, thousands of patients with thousands of items per patient could be shared between some of several thousands of health professionals. This exchange has to be bound on direct or indirect contributions to the patient's care leading to different rights and duties related to that information. These relations are highly dynamic and cannot be managed by an administrator using multidimensional matrixes related to persons, circumstances, information objects, specific items, time restrictions, etc. Therefore, methodologies should be borrowed from system modelling, design, implementation, and maintenance to meet the challenges aforementioned.

2. Future-proof EHR architecture

The framework for a future-proof EHR architecture is based on the generic component model developed in the mid-nineties (e.g., [1—3]). Basis of that architecture are a reference information model (RIM) and agreed vocabularies enabling inter-operability. Referenced to them, domain-specific constraint models will be specified which represent domain-specific knowledge concepts, considering both structural and functional knowledge. The corresponding components have to be established according to all views of the ISO reference model-open distributed processing (RM-ODP) [4], i.e. enterprise, information, computational, engineering, and technology view. A view focuses consideration on one aspect abstracting from all others. The different domain concepts and their view representation is not the task of programmers but of domain experts. For that reason, they will use appropriate expression means such as specific graphical representation (e.g., UML diagrams) or sometimes even verbal templates expressed in XML.

The components can be aggregated to higher level of composition. Contrary to the ISO definition of primitives and composition, in the generic component model at least four level of composition/decomposition have been defined (Fig. 1).

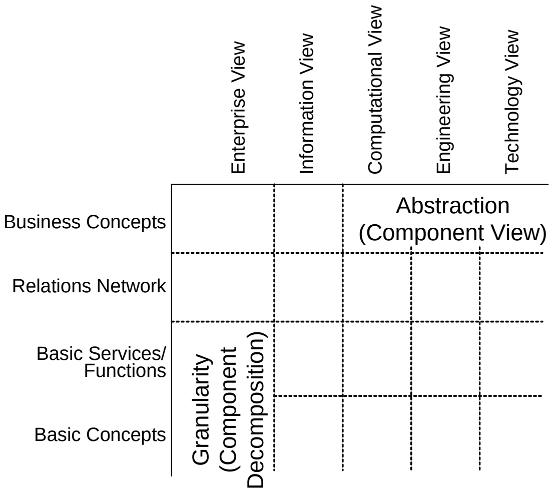

Fig. 1 Generic component model.

The aggregation is performed according to content- or process-related knowledge expressed by logics/algorithms/operations or rules/workflows/procedures/relationships. So, the aggregation of the building blocks ''constraint models'' is controlled by the aforementioned mechanisms or by the communicating or co-operating principal's behaviour. The specification is completely provided at meta-level. Different vocabularies as well as tooling environment and functionality are harmonised by meta-languages like XML metadata interchange (XMI) [5]. Especially the required consistency with the revision of CEN ENV 13606 ''EHR Communication'' predefines this methodology.

3. Modelling authorisation and access control

The model for authorisation and access control in distributed health information systems has to deal with policy description and negotiation including policy agreements, authentication, certification, and directory services but also audit trails, altogether forming the privilege management infrastructure [3].

To make policy specifications standardised and negotiable for achieving policy agreements, policy documents must be formalised. This formalisation concerns grammar and vocabulary used, structure, and content.

3.1. Models used

In our approach, several security-related models have been used: the domain model, the policy model, the role model, the privilege management model, the authorisation model, the access control model as well as the information distance model. Some of the models will be discussed in more detail. For the other models use references such as [3].

3.1.1. Domain model

To keep (complex) information systems that support shared care manageable and operating, principal-related components of the system are grouped by common organisational, logical, and technical properties into domains, according to OMG's definition-forming policy domains, environmental domains, or technology domains. Regarding the generic concept of domains, they can be extended by chaining sub-domains to super-domains forming a common domain of communication and co-operation, which is characterised by establishing an agreed security policy. Such transaction-concrete policy has to be negotiated between the communicating and co-operating principals, which is also called policy bridging [6,7].

3.1.2. Policy model

A security policy is the complex of legal, ethical, social, organisational, psychological, functional, and technical implications for trustworthiness of health information systems. It formulates the concept of requirements and conditions for trustworthy creation, storage, processing and use of sensitive information. A policy can be expressed in three different ways: verbally unstructured, structured using schemata or templates, or formally modelled.

For inter-operability reasons, a policy must be formulated and encoded in a way enabling its correct interpretation and practicing. Therefore, policies have to be constrained regarding syntax, semantics, vocabulary, and operation of policy documents, also called policy statements or policy agreements (agreements between the partners involved). One common way to express constraints is the specification of user-defined schemata such as XML schemata. This schema should be standardised for inter-operability purposes mentioned above. Fig. 2 presents a simple XML instance for a security policy statement. To doubtless refer to a specific policy, the policy instance must be uniquely named and identified via a unique policy ID. The same is true for all the policy components such as domain, targets, operations, and their policies, which have to be named and uniquely identified too [3].

```
<policy>
        <policy_name/>
        <policy_identifier/>
        <policy_authority/>
        <domain_name/>
        <domain_identifier/>
        <target_list>
                <target_name/>
                <target_ID/>
                <target_object>
                        <operations/>
                        <policies/>
                </target_object>
        </target_list>
</policy>
```

Fig. 2 Policy template example.

According to the generic component model approach [2], also the policy concept can be composed/decomposed providing different levels of granularity to fulfil specific needs of definition, expression, and management. Different concept presentations offer different basic components for policies. The PONDER approach [8] defining a declarative, object-oriented language for specifying policies for security and management of distributed systems, introduced basic policy types and composite policies. Basic policy types are:

- authorisation policies that define permitted actions therefore containing subject (except in roles), target, action;
- obligation policies which are event-triggered and define actions to be performed by manager agents therefore containing subject (except in roles), action, event;
- refrain policies that define actions the subjects must refrain from performing therefore containing subject (except in roles), action;
- delegation policies that define what authorisations can be delegated to whom.

Composite policies types are:

- groups that define a scope for related policies to which a set of constraints can apply;
- roles define a group of policies (authorisation, obligation and refrain policies) related to positions within an organisation;
- relationships define a group of policies pertaining to the interactions between a set of roles.

Management structures define a configuration of role instances as well as the relationships between them. Filtered actions extend authorisations allowing the transformation of input and output parameters to be defined. Constraints specify limitations on the applicability of policies while meta-policies define semantic constraints on permitted policies. There are analogies to OMG's object constraint language (OCL) specification [9]. Fig. 3 shows the PONDER base-class diagram.

Another policy decomposition example is given in OMG's CORBA security services specification [5] and in related documents, distinguishing between invocation access policy implementing access control policy for objects, invocation audit policy controlling event type and criteria for audit, and secure invocation policy specifying security policies

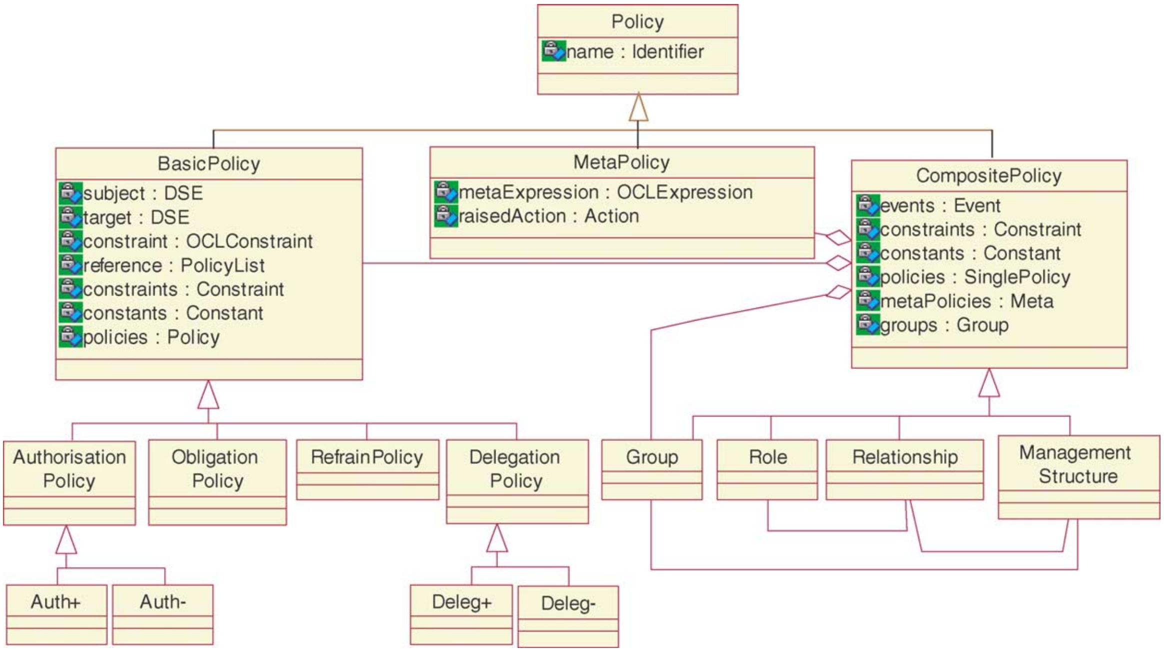

Fig. 3 PONDER base-class diagram.

associated with security associations and message protection. Regarding requirements for different object types, invocation delegation policy, application access policy, application audit policy, and non-repudiation policy have been defined.

3.1.3. Role model

Specification and management of roles are crucial issues in privilege management and access control. Roles provide a means to indirectly assign privileges to individuals. Individuals are issued role assignment certificates that assign one or more roles to them through the role attribute contained in the certificate. Specific privileges are assigned to a role name through role specification certificates, rather than to individual privilege holders through attribute certificates. This level of indirection enables, for example, the privileges assigned to a role to be updated, without impacting the certificates that assign roles to individuals. Role assignment certificates may be attribute certificates or public-key certificates. Role specification certificates may be attribute certificates, but not public-key certificates. If role specification certificates are not used, the assignment of privileges to a role may be done through other means (e.g., may be locally configured at a privilege verifier).

In general, two types of roles can be distinguished: structural (or organisational) at the one hand, and functional roles at the other hand. Structural roles enable certain services within the generic structure—function relationship. Reflecting human or organisational categories, structural roles describe prerequisites, feasibilities, or competencies for actions. Functional roles are bound to the realisation of actions.

The structural model of a system defines structural roles being rather static. The functional model of a system defines highly dynamic functional roles [3,7]. Examples for structural role are positions in an organisations hierarchy such as director of clinic, head physician, etc., qualifications such as medical doctor, nurse, etc., specialties such as paediatrician, internist, gynaecologist, radiologist, etc. Examples for functional roles are caring doctor, member of a diagnostic team, member of a therapeutic team, attending nurse, prescribing doctor, etc.

In the object-oriented world, semantics for structural and behavioural object models have been specified. Structural models (also known as static models) emphasize the structure of objects in a system, including their classes, interfaces, attributes and relations. Behavioural models (also known as dynamic models) emphasise the behaviour of objects in a system, including their methods, interactions, collaborations, and state histories. Roles might be assigned to any principal. Because principals are actors in use cases, roles have relationship to actors and therefore to actions. Persistence denotes the permanence of the state of the association, marking it as transitory (its state is destroyed when the instance is destroyed) or persistent (its state is not destroyed when the instance is destroyed). In that context, roles which are described as UML association classes, establish

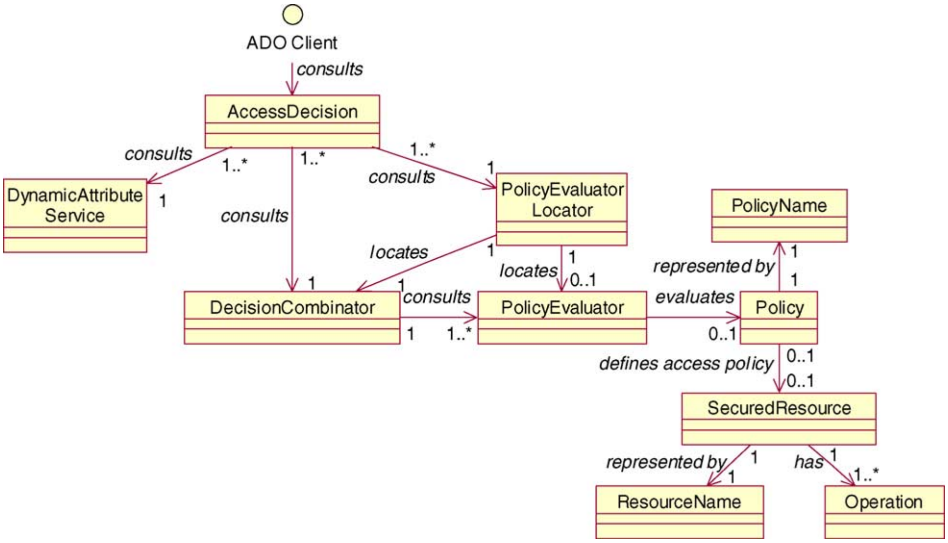

Fig. 4 RAD access decision model [11].

persistent associations in the structural role or certified privilege case, but transient associations in the functional roles environment. Thus, functional role-based actions establish structural roles (e.g., certification act), which themselves may constrain other functional roles. UML does not require the drawing of both association and association class.

3.1.4. Privilege management, authorisation and access control model

The general privilege management model consists of three entities: object, privilege asserter, and privilege verifier. For deciding access to resources, both the assignment of object security attributes and the location and deployment of the appropriate policy must be realised. Fig. 4 presents the resource access decision service (RADS) model used for the CORBA RADS [10]. It should be mentioned that there are three principle decisions made in the privilege management context: request authorised, request denied, and request modified. Fig. 5 demonstrates the CORBA authorisation model. Three different basic access control models have been widely accepted: the owner-based discretionary access control (DAC), the lattice-based mandatory access control (MAC), and the role-based access control model (RBAC), which is especially supported by our approach.

3.1.5. Information distance model

Regarding the distance of persons to personal information, three types of person with growing distance to the information can be specified [12]: the originator of information (holder of data), the

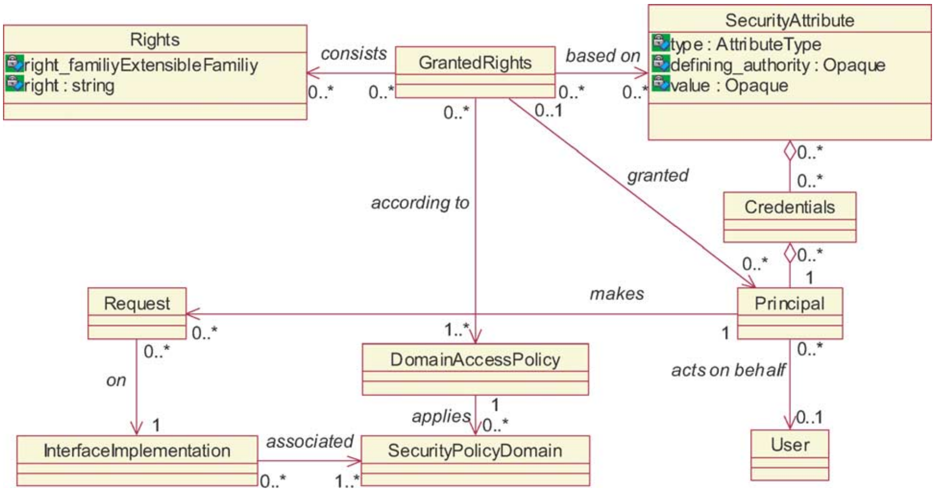

Fig. 5 CORBA authorisation model [11].

producer of information (interpreter of data), and the administrator of information (user of information). In a healthcare environment, the originator of information is normally the patient and the producer of information is the doctor. An example of an information user is a pharmacist. Following the need to know principle, an increasing distance to information causes greater restrictions regarding privileges granted.

3.2. The enterprise security integration framework

The enterprise security integration (ESI) framework provides rapid deployment of secure e-business applications by specifying the interactions amongst security services and application components that use such security services. Using common interfaces, ESI enables the addition of new security technology solutions without changing the existing framework, thus providing openness, flexibility, and scalability. The framework supports security mechanisms to enforce security in security-aware as well as security-unaware applications. ESI core security services concern authentication, authorisation, confidentiality, accountability including audit, and security administration. Additionally, framework security facilities such as profile manager, security association, and security proxy services must be established. For authentication, security tokens such as smart cards also defining basic roles like profession are widely used, therefore also called health professional cards (HPCs). The relationship between principals and token has been certified by ID certificates authorised by certification authorities within a PKI architecture. A trusted third party assigns attribute certificates (key-less certificates) representing special roles of a principal to ID certificates. Fig. 6 presents the role definition according to ESI.

According to the object paradigm, security services as specific methods are provided by implemented components. A component-specific role concept enables access to such a security service component. Both role concepts can be linked by role references [11].

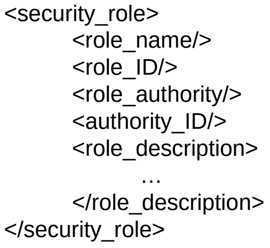

```
<security_role>
        <role_name/>
        <role_ID/>
        <role_authority/>
        <authority_ID/>
        <role_description>
                ...
        </role_description>
</security_role>
```

Fig. 6 Role definition.

4. Authentication mechanisms and attribute certificates

The authentication framework has been specified, e.g., in ISO 9798 and ISO 10181. The authentication procedure is based on a PKI. Verification of the principal's authentication provided by a random number encoded with the principal's private key is performed using that principal's public key which is bound to the principal by a certificate signed by the PKI CA. The authentication certificate follows the X.509V3 specification.

There are several ways for binding key-related ID certificates to key-less attribute certificates: the monolithic approach, the autonomic approach, and the approach of chained signatures. In the monolithic approach, the attribute certificate is part of the ID certificate. In the autonomic approach, some relevant information in the ID certificate is referred to bind with the attribute certificate. In the binding approach using chained signatures, the ID certification authority's signature is referred to bind with the attribute certificate. The ISO TDS 17090 fixed the first approach [3]. Claimant privileges are conveyed as attributes, in either a public key certificate (in the subject directory attributes extension), or (more frequently) in an attribute certificate. The syntax of an attribute certificate is specified in X.509. The usage of attribute certificate components is described in the corresponding standards [13].

5. Classification of security objects

For managing privileges and access control as well as for deciding corresponding rights, the accessing principals and the requested and used information must be grouped. Otherwise, the highly dynamic matrix to be managed would go beyond any limits. For classifying personal medical data in a shared care environment, the military classification scheme is insufficient. Therefore, the standards mentioned define an extended and refined schema considering medical workflow as well as the legal, ethical and social implications of personal medical information.

6. Component-based secure EHR over the internet

The European HARP project has specified and implemented open portable EHR systems enriched with enhanced TTP services and comprehensive development strategies for establishing fine grained application security services. Constraints specified can be bound to components at runtime, enabling different views or supporting specific domain knowledge concepts. By binding attribute certificates to components, appropriate policies can be enforced. These constraints such as, e.g., certificates are interpreted at both server and client side using authorization services. the harp cross-security platform is solely based on standards including the XML standard set for the establishment of EHR clients and servers as well as their communication [14].

7. Conclusions

The establishment of shared care must be supported by distributed, inter-operable information systems. Multiple uses of personal medical data by shared care partners require a trustworthy PKI-based communication and co-operation platform as well as a privilege and access control infrastructure. The components of such an infrastructure have to be specified regarding syntax, semantics, protocols and scenarios. Currently, these specifications are provided as new standards of ISO TC 215 as well as CEN TC 251, both dedicated to health informatics.

Following the stream defined in the Generic Component Model paradigm, newer publications (e.g., [15]) confirm correctness and feasibility of the approach proposed. The specifications mentioned have been deployed already in the practical environment of the first German health network ''ONCONET Saxony-Anhalt''.

For enabling international communication of personal health information, security and safety-related basic concepts, aggregations, relations network, and business concepts must be specified and modelled by international domain experts such as members of IMIA Working Group 4 ''Security'' to be included into component-based solutions according to our widely accepted and standardised approach.

Acknowledgements

The author is in debt to many members of the HARP Consortium, CORBA, ISO TC 215 WG 4, CEN TC 251 WG 3 as well as ASTM Committee E 31 for kind co-operation.

References

[1] B. Blobel, M. Holena, Comparison, evaluation, and possible harmonisation of the HL7, DHE, and CORBA middleware, in: J. Dudeck, B. Blobel, W. Lordieck, T. Bürkle (Eds.), New technologies in hospital information systems, Series Studies in Health Technology and Informatics, vol. 89, IOS Press, Amsterdam, 1997, pp. 40—47.
[2] ISO/IEC 10746-2, Information technology: open distributed processing—reference model. Part 2. Foundations.
[3] W3C, XML Metadata Interchange: http://www.w3c.org.
[4] B. Blobel, Analysis, design and implementation of secure and interoperable distributed health information systems, Series Studies in Health Technology and Informatics, vol. 89, IOS Press, Amsterdam, 2002.
[5] OMG Inc, The CORBA Security Specification, Object Management Group Framingham Inc., 1997.
[6] B. Blobel, F. Roger-France, A systematic approach for analysis and design of secure health information systems, Int. J. Med. Inform. 62 (3) (2001) 51—78.
[7] B. Blobel, Application of the component paradigm for analysis and design of advanced health system architectures, Int. J. Med. Inform. 60 (3) (2000) 281—301.
[8] N. Damianou, N. Dulay, E. Lupu, M. Sloman, Ponder: a language for specifying security and management policies for distributed systems: the language specification, version 2.3, Imperial College Research Report DoC 2000—2001, 20 October 2000.
[9] OMG Inc. Object Constraint Language Specification: http://www.omg.org.
[10] OMG Inc. Resource access decision facility specification, version 1.0, April 2001: http://www.omg.org.
[11] B. Hartmann, D.J. Flinn, K. Beznosov, Enterprise Security with EJB and CORBA, Wiley, New York, 2001.
[12] K. Yamamoto, et al., The awareness of security issues among hospitals in Japan, in: Proceedings of the IMIA Workshop Caring for Health Information Safety, Security and Secrecy, Hemskerk, The Netherlands, 13—16 November 1993.
[13] ISO DTS 17090 Health Informatics—Public Key Infrastructure, 2001.
[14] HARP Consortium: http://www.ist-harp.org.
[15] G. Neumann, M. Strembeck, A scenario-driven role engineering process for functional RBAC roles, in: Proceeding of the ACM Workshop on Role-Based Access Control, 2002.