INTERNATIONAL STANDARD

ISO 22600-2

First edition
2014-10-01

Health informatics — Privilege management and access control —

Part 2: Formal models

Informatique de santé — Gestion de privilèges et contrôle d'accès —

Partie 2: Modèles formels

Reference number
ISO 22600-2:2014(E)

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

Page

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the WTO principles in the Technical Barriers to Trade (TBT) see the following URL: Foreword - Supplementary information

The committee responsible for this document is ISO/TC 215, *Health informatics*.

This first edition of ISO 22600-2 cancels and replaces ISO/TS 22600-2:2006, which has been technically revised.

ISO 22600 consists of the following parts, under the general title *Health informatics — Privilege management and access control*:

— *Part 1: Overview and policy management*

— *Part 2: Formal models*

— *Part 3: Implementations*

Introduction

The distributed architecture of shared care information systems supporting service-oriented architecture (SOA) is increasingly based on corporate networks and virtual private networks. For meeting the interoperability challenge, the use of standardized user interfaces, tools, and protocols, which ensures platform independence, but also the number of really open information systems, is rapidly growing during the last couple of years.

As a common situation today, hospitals are supported by several vendors providing different applications, which are not able to communicate authentication and authorization since each has its own way of handling these functions. For achieving an integrated scenario, it takes a remarkable amount of money, time, and efforts to get users and changing organizational environments dynamically mapped before starting communication and cooperation. Resources required for the development and maintenance of security functions grow exponentially with the number of applications, with the complexity of organizations towards a regional, national, or even international level, and with the flexibility of users playing multiple roles, sometimes even simultaneously.

The situation becomes even more challenging when inter-organizational communications happens, thereby crossing security policy domain boundaries. Moving from one healthcare centre to another or from country to country, different rules for privileges and their management can apply to similar types of users, both for execution of particular functions and for access to information. The policy differences between these domains have to be bridged automatically or through policy agreements, defining sets of rules followed by the parties involved, for achieving interoperability.

Another challenge to be met is how to improve the quality of care by using IT without infringing the privacy of the patient. To provide physicians with adequate information about the patient, a virtual electronic health care record is required which makes it possible to keep track of all the activities belonging to one patient regardless of where and by whom they have been performed and documented. In such an environment, a generic model or specific agreement between the parties for managing privileges and access control including the patient or its representative is needed.

Besides a diversity of roles and responsibilities, typical for any type of large organization, also ethical and legal aspects in the healthcare scenario due to the sensitivity of person-related health information managed and its personal and social impact have to be considered.

Advanced solutions for privilege management and access control are required today already, but this challenge will even grow over the next couple of years. The reason is the increase of information exchanged between systems in order to fulfil the demands of health service providers at different care levels for having access to more and more patient-related information to ensure the quality and efficiency of patient's diagnosis and treatment, however combined with increased security and privacy risks.

The implementation of this International Standard might be currently too advanced and therefore not feasible in certain organizational and technical settings. For meeting the basic principle of best possible action, it is therefore very important that at least a policy agreement is written between the parties stating to progress towards this International Standard when any update/upgrade of the systems is intended. The level of formalization and granularity of policies and the objects these policies are bound to defines the solution maturity on a pathway towards the presented specification.

The policy agreement also has to contain defined differences in the security systems and agreed solutions on how to overcome the differences. For example, the authentication service and privileges of a requesting party at the responding site have to be managed according to the policy declared in the agreement. For that reason, information and service requester, as well as information and service provider on the one hand, and information and services requested and provided on the other hand, have to be grouped and classified in a limited number of concepts for enabling the specification of a limited number of solution categories. Based on that classification, claimant mechanisms, target sensitivity mechanisms, and policy specification and management mechanisms can be implemented. Once all parties have signed the policy agreement, the communication and information exchange can start with the existing systems if the parties can accept the risks. If there are unacceptable risks which have to be eliminated before the information exchange starts, they shall also be recorded in the policy agreement

together with an action plan stating how these risks shall be removed. The policy agreement also has to contain a time plan for this work and an agreement on how it shall be financed.

The documentation of the negotiation process is very important and provides the platform for the policy agreement.

Privilege management and access control address security and privacy services required for communication and cooperation, i.e. distributed use of health information. It also implies safety aspects, professional standards, and legal and ethical issues. This International Standard introduces principles and specifies services needed for managing privileges and access control. Cryptographic protocols are out of the scope of this International Standard.

This three-part International Standard references existing architectural and security standards as well as specifications in the healthcare area such as ISO, CEN, ASTM, OMG, W3C, etc., and endorses existing appropriate standards or identifies enhancements or modifications or the need for new standards. It comprises of:

— ISO 22600-1: describes the scenarios and the critical parameters in information exchange across policy domains. It also gives examples of necessary documentation methods as the basis for the policy agreement.

— ISO 22600-2: describes and explains, in a more detailed manner, the architectures and underlying models for privilege management and access control which are necessary for secure information sharing including the formal representation of policies.

— ISO 22600-3: describes examples of implementable specifications of application security services and infrastructural services using different specification languages.

It accommodates policy bridging. It is based on a conceptual model where local authorization servers and cross border directory and policy repository services can assist access control in various applications (software components). The policy repository provides information on rules for access to various application functions based on roles and other attributes. The directory service enables identification of the individual user. The granted access will be based on four aspects:

— the authenticated identification of principals (i.e. human users and objects that need to operate under their own rights) involved;

— the rules for access to a specific information object including purpose of use;

— the rules regarding authorization attributes linked to the principal provided by the authorization manager;

— the functions of the specific application

This International Standard supports collaboration between several authorization managers that can operate over organizational and policy borders.

This International Standard is strongly related to other ISO/TC 215 work such as ISO 17090 (all parts), ISO 22857, ISO 21091, and ISO 21298.

This International Standard is meant to be read in conjunction with its complete set of associated standards.

Health informatics — Privilege management and access control —

Part 2: Formal models

1 Scope

This multi-part International Standard defines principles and specifies services needed for managing privileges and access control to data and/or functions.

It focuses on communication and use of health information distributed across policy domain boundaries. This includes healthcare information sharing across unaffiliated providers of healthcare, healthcare organizations, health insurance companies, their patients, staff members, and trading partners by both individuals and application systems ranging from a local situation to a regional or even national situation.

It specifies the necessary component-based concepts and is intended to support their technical implementation. It will not specify the use of these concepts in particular clinical process pathways.

This part of ISO 22600 introduces the underlying paradigm of formal high-level models for architectural components. It is based on ISO/IEC 10746 (all parts) and introduces the domain model, the document model, the policy model, the role model, the authorization model, the delegation model, the control model, and the access control model.

The specifications are provided using the meta-languages Unified Modelling Language (UML) and Extensible Markup Language (XML). Additional diagrams are used for explaining the principles. The attributes used have been referenced to the HL7 reference information model (see ISO 21731:2006) and the HL7 data type definitions.

The role model has been roughly introduced referring to ISO 21298.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO 21298:—[1)], *Health informatics — Functional and structural roles*

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

3.1
access control
means of ensuring that the resources of a data processing system can be accessed only by authorized entities in authorized ways

[SOURCE: ISO/IEC 2382-8:1998]

1) To be published.

3.2
accountability
property that ensures that the actions of an entity can be traced uniquely to the entity

[SOURCE: ISO 7498-2:1989]

3.3
attribute authority
AA
authority which assigns privileges by issuing attribute certificates

[SOURCE: ISO/IEC 9594-8:2008]

3.4
attribute certificate
data structure, digitally signed by an attribute authority, that binds some attribute values with identification about its holder

[SOURCE: ISO/IEC 9594-8:2008]

3.5
authentication
provision of assurance of the claimed identity of an entity by securely associating an identifier and its authenticator

Note 1 to entry: See also data origin authentication and peer entity authentication.

[SOURCE: ISO/IEC 15944-5:2008, 3.5]

3.6
authority
entity, which is responsible for the issuance of certificates

Note 1 to entry: Two types are distinguished in this part of ISO 22600: certification authority which issues public key certificates and attribute authority which issues attribute certificates.

3.7
authorization
granting of privileges, which includes the granting of privileges to access data and functions

[SOURCE: ISO 7498-2:1989, modified]

3.8
availability
property of being accessible and useable upon demand by an authorized entity

[SOURCE: ISO 7498-2:1989]

3.9
certificate validation
process of ensuring that a certificate was valid at a given time, including possibly the construction and processing of a certification path, and ensuring that all certificates in that path were valid (i.e. were not expired or revoked) at that given time

3.10
certification authority
CA
certificate issuer; an authority trusted by one or more relying parties to create, assign, and manage certificates

Note 1 to entry: Optionally, the certification authority can create the relying parties' keys. The CA issues certificates by signing certificate data with its private signing key.

Note 2 to entry: Authority in the CA term does not imply any government authorization, only that it is trusted. Certificate issuer can be a better term but CA is used very broadly.

[SOURCE: ISO/IEC 9594-8:2008]

3.11
certification path
ordered sequence of certificates of objects in the DIT which, together with the public key of the initial object in the path, can be processed to obtain that of the final object in the path

3.12
confidentiality
property that information is not made available or disclosed to unauthorized individuals, entities, or processes

[SOURCE: ISO 7498-2:1989]

3.13
credential
prerequisite issued evidence for the entitlement of, or the eligibility for, a role

3.14
delegation
conveyance of privilege from one entity that holds such privilege to another entity

3.15
delegation path
ordered sequence of certificates which, together with authentication of a privilege asserter's identity, can be processed to verify the authenticity of a privilege asserter's privilege

3.16
environmental variables
aspects of policy required for an authorization decision that are not contained within static structures, but are available through some local means to a privilege verifier (e.g. time of day or current account balance)

3.17
identification
performance of tests to enable a data processing system to recognize entities

[SOURCE: ISO/IEC 2382-8:1998]

3.18
identifier
piece of information used to claim an identity, before a potential corroboration by a corresponding authenticator

[SOURCE: ENV 13608-1:2000]

3.19
integrity
property that information is not altered in any way, deliberately or accidentally

3.20
key
sequence of symbols that controls the operations of encipherment and decipherment

[SOURCE: ISO 7498-2:1989]

3.21
non-repudiation
service providing proof of the integrity and origin of data (both in an unforgeable relationship) which can be verified by any party

[SOURCE: ISO 17090-1:2013]

3.22
policy
set of legal, political, organizational, functional, and technical obligations for communication and cooperation

3.23
policy agreement
written agreement where all involved parties commit themselves to a specified set of policies

3.24
principal
human users and objects that need to operate under their own rights

[SOURCE: OMG Security Services Specification: 2001]

3.25
private key
key that is used with an asymmetric cryptographic algorithm and whose possession is restricted (usually to only one entity)

[SOURCE: ISO/IEC 10181-1:1996]

3.26
privilege
capacity assigned to an entity by an authority according to the entity's attribute

3.27
privilege asserter
privilege holder using their attribute certificate or public key certificate to assert privilege

3.28
privilege management infrastructure
PMI
infrastructure able to support the management of privileges in support of a comprehensive authorization service and in relationship with a public key infrastructure

3.29
privilege policy
policy that outlines conditions for privilege verifiers to provide/perform sensitive services to/for qualified privilege asserters

Note 1 to entry: Privilege policy relates attributes associated with the service as well as attributes associated with privilege asserters.

3.30
privilege verifier
entity verifying certificates against a privilege policy

3.31
public key
key that is used with an asymmetric cryptographic algorithm and that can be made publicly available

[SOURCE: ISO/IEC 10181-1:1996]

3.32
public key certificate
PKC
certificate that binds an identity and a public key

[SOURCE: ISO/IEC 9594-8:2008]

Note 1 to entry: The identity can be used to support identity-based access control decisions after the client proves that it has access to the private key that corresponds to the public key contained in the PKC (see RFC 2459).

3.33
role
set of competences and/or performances that are associated with a task

3.34
role assignment certificate
certificate that contains the role attribute, assigning one or more roles to the certificate holder

3.35
role certificate
certificate that assigns privileges to a role rather than directly to individuals

Note 1 to entry: Individuals assigned to that role, through an attribute certificate or public key certificate with a subject directory attributes extension containing that assignment, are indirectly assigned the privileges contained in the role certificate.

3.36
role specification certificate
certificate that contains the assignment of privileges to a role

3.37
sensitivity
characteristic of a resource that implies its value or importance

3.38
security
combination of availability, confidentiality, integrity, and accountability

[SOURCE: ENV 13608-1:2000]

3.39
security policy
plan or course of action adopted for providing computer security

[SOURCE: ISO/IEC 2382-8:1998]

Note 1 to entry: The set of rules laid down by the security authority governing the use and provision of security services and facilities constitutes its security policy.

3.40
security service
service, provided by a layer of communicating open systems, which ensures adequate security of the systems or of data transfers

[SOURCE: ISO 7498-2:1989]

3.41
source of authority
SOA
attribute authority that a privilege verifier for a particular resource trusts as the ultimate authority to assign a set of privileges

3.42
target
resource being accessed by a claimant

Note 1 to entry: Its sensitivity is modelled in this part of ISO 22600 as a collection of attributes, represented as either ASN.1 attributes or XML elements.

3.43
trust
circumstance existing between two entities whereby one entity makes the assumption that the other entity will behave exactly as the first entity expects

Note 1 to entry: This trust applies only for some specific function. The key role of trust in this framework is to describe the relationship between an authenticating entity and an authority; an entity must be certain that it can trust the authority to create only valid and reliable certificates.

4 Abbreviated terms

AA	Attribute Authority
PKC	Public Key Certificate
UML	Unified Modelling Language
XML	Extensible Markup Language

5 Component paradigm

The framework for a future-proof health information system architecture is based on the generic component model developed in the mid-nineties (e.g. References [1], [2], and [3]). Bases of that architecture are a reference information model (RIM) and agreed vocabularies enabling interoperability. 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/IEC 10746-1 reference model of open distributed processing (RM-ODP), i.e. enterprise view, information view, computational view, engineering view, 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 structured text expressed in XML.

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

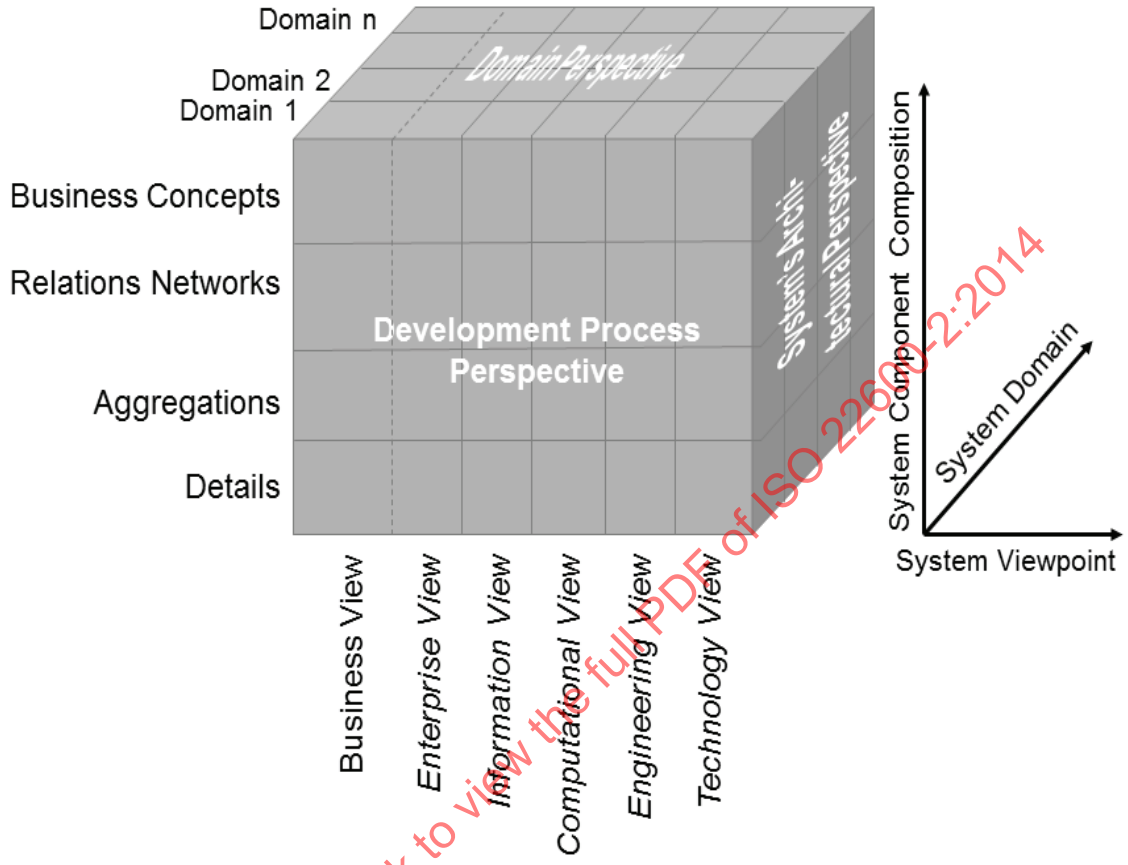

Figure 1 — Generic component model

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 harmonized by meta-languages like XML Metadata Interchange (XMI).[4]

6 Generic models

6.1 Framework

Privilege management and authorization can be based on roles that individual actors or groups of individual actors play. Actors interacting with system components are called principals, which can be a human user, a system, a device, an application, a component, or even an object.

In order to obtain the above-described structure and functionality, there are a number of models, mechanisms, processes, objects, etc. needed, which have to be considered.

Regarding privilege management and access control management, two basic class types shall be dealt with:

- entities:
  - documents;
  - principals;
  - policies;
  - roles;
- acts:
  - policy management;
  - principal management;
  - privilege management;
  - authentication;
  - authorization;
  - access control management;
  - audit.

The following models will be considered in more detail:

- domain model;
- document model;
- policy model;
- role model;
- authorization model;
- control model;
- delegation model;
- access control model.

All specifications in this framework will be kept open, platform-independent, portable, and scalable. Therefore, the models provided are described at meta-model level and at the model level keeping the instance level out of consideration. For expressing systems in such a way, specific languages and meta-languages are used such as UML and XML including means for transfer from one vocabulary to another one.

This specification is defined using UML constructs, UML specifications, UML profiles, and all different diagrams. Regarding XML, several specifications within the XML standard set will be used.

All models being used establish specific kinds of constraints forming constraint models. This concerns all conceivable services or views on systems. A model is a partial representation of reality according to special concepts. The language to be used for graphical models is UML and MOF. The language for verbal (text-based) models is the XML standard set.

It is expected that many documents will be expressed using XML. The structure for such a document is defined in a document type definition (DTD) or an XML schema instance. A privilege policy can act

directly on the XML elements (e.g. by comparing attributes in an authorization certificate to elements in the document).

6.2 Domain model

To keep (complex) information systems that support shared care manageable and operating, principal-related components of the system are grouped by common organizational, logical, and technical properties into domains. Following OMG's (Object Management Group) definition, this could be done for common policies (policy domains), for common environments (environment domains), or common technology (technology domains). Any kind of interoperability internally to a domain is called an intra-domain communication and co-operation, whereas interoperability between domains is called an inter-domain communication and co-operation. For example, communication could be realized between departments of a hospital internally to the domain hospital (intra-domain communication), but externally to the domain of a special department (inter-domain communication). Regarding security requirements, security policy domains are of special interest.

A domain is characterized by a domain identifier, a domain name, a domain authority, and a domain qualifier. The provided data type definition resembles the HL7 version 3 data type definition.[5]

Table 1 — Security policy domain attributes

Attribute	Type	Remarks
domain_identifier	SET < OID >	Set of ISO ObjectIdentifier
domain_name	BAG < EN >	Bag of EntityName
domain_authority_ID	OID	ISO ObjectIdentifier
domain_authority_name	ST	String
domain_qualifier	CS	CodedSimpleValue

Security policy domain class inherits attributes from domain class, plus the attributes policy identifier and policy name.

A policy describes the legal framework including rules and regulations, the organizational and administrative framework, functionalities, claims and objectives, the principals involved, agreements, rights, duties, and penalties defined as well as the technological solution implemented for collecting, recording, processing, and communicating data in information systems. For describing policies, methods such as policy templates or formal policy modelling might be deployed.

Domains are specified generically in this part of ISO 22600, and their definition in practice can be flexible. A domain might consist of sub-domains (which will inherit and might specialize policies from the parent domain). The smallest-scale domain might be an individual workplace or a specific component within an information system. Domains can be extended into super-domains, by chaining a set of distinct domains and forming a common larger-scale domain for communication and co-operation.

This co-operation between domains requires the definition of a common set of policies that applies to all of the collaborating domains. It shall be derived from all of the relevant domain-specific policies across all of those domains. These common policies are derived (negotiated) through a process known as policy bridging (see Figure 2). The eventual agreed policies need to be documented and signed by all of the domain authorities (see ISO 22600-1, Annex A). Ideally, this whole process will be capable of electronic representation and negotiation, to permit real-time electronic collaboration to take place within a (pre-agreed) permitted and regulated framework. The policy negotiation or verification would then take place at every service interaction.

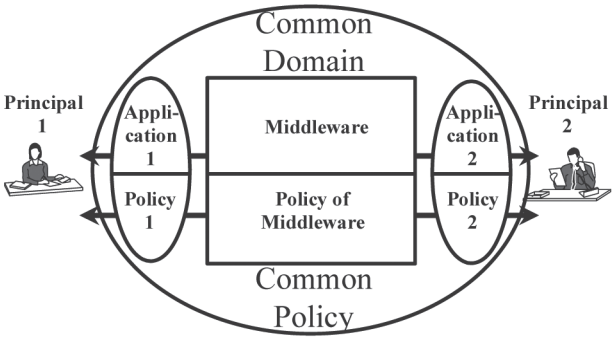

Figure 2 — Policy bridging

This collaboration will introduce the need for components between the principals. Middleware concepts are being introduced increasingly into the new(er) healthcare information systems. Middleware components can enable interoperability through direct invocation (middleware communication services) or chained invocation (including middleware application services). The latter is characterized by different models of delegation (see 6.8).

Such an architecture can be represented by chains of different domains as shown in Figure 3.

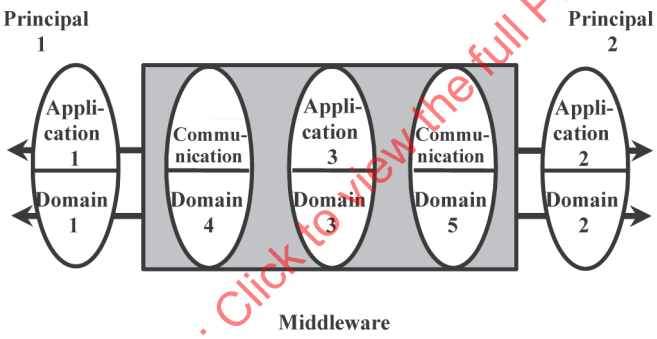

Figure 3 — Domain concept with middleware services

From the security point of view, a domain ensuring intra-domain communication according to its own policy is commonly considered with need of protection only at its boundary to external domains with their specific policies (or even the policy-free domain of the Internet). This is done by e.g. firewalls, proxy servers, etc. Regarding the external environment, a domain is therefore often considered closed. The internal domain is mistakenly assumed to be secure, often neglecting internal threats and attacks which are among the majority of all security attacks.

Regarding the specific requirements and conditions of healthcare, the underlying security model shall consider the whole spectre of security services and mechanisms, which can be accomplished by secure micro domains.

6.3 Document model

Processes, entities, roles, etc. shall be documented, resulting in a special information object, which has to be signed for provably expressing the particular relations between entities and processes. The combination of processes and relations leads to multiple signatures (e.g. in the case of delegation).

This part of ISO 22600 uses the cryptographic message syntax to support multiple signatures on a document. Each signature is computed over the document content and optionally a set of signed attributes specific to the particular signature. These attributes can include timestamps, signature purpose, and other information.

6.4 Policy model

A security policy is the complex of legal, ethical, social, organizational, psychological, functional, and technical rules for ensuring trustworthiness of health information systems. A policy is the formulation of the concept of requirements and conditions for trustworthy creation, collection, storage, processing, disclosure, retention, transmission, and use of sensitive information. A policy can be expressed

— verbally unstructured,

— structured using schemata or templates, or

— formally modelled.

For interoperability reasons, a policy shall be formulated and encoded in a way that enables its correct interpretation and practice. 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).

To reliably refer to a specific policy, the policy instance shall be uniquely named and identified via a unique policy ID. The same is true for all the policy elements such as domain, targets, operations, and their policies, which have to be named and uniquely identified too. In summary, a policy is characterized by a policy identifier, a policy name, a policy authority, a domain identifier, a domain name, a target list, target identifier, target name, target object, operations allowed, and related policies.

For readability reasons, domain-related attributes have been included in Table 2 even if policy inherits from domain. The provided data type definition resembles the HL7 version 3 data type definition.[5]

Table 2 — Basic security policy attributes

Attribute	Type	Remarks
policy_identifier	SET <II>	Set of InstanceIdentifier
policy_name	CS	CodedSimpleValue
policy_authority_ID	OID	ISO ObjectIdentifier
policy_authority-name	ST	String
domain_identifier	SET <OID>	Set of ISO ObjectIdentifier
domain_name	BAG <EN>	Bag of EntityName
target_list	LIST <INT>	List of INT
target_ID	SET <II>	Set of InstanceIdentifier
target_name	EN	EntityName
target_object	II	InstanceIdentifier
operation_code	CE	CodedWithEqivalents
policies	CD	ConceptDescription

Where allowed by the laws of the jurisdiction, health information systems such as the EHR should at minimum have a policy for patients to control access to their health information, a policy with common access rules by the organization, policies defined by laws and regulations, and one policy per structural role as well as one policy per functional role.

Every creation, collection, access or modification, use, disclosure, retention, and destruction of an EHR component shall be covered by one or more policies. The reference model of the EHR Extract includes a policy ID attribute within the record component class to permit references to such policies to be made at any level of granularity within the EHR hierarchy. The policies that apply specifically to an EHR can be included within the EHR Extract, eventually including any bridged policies.

As any other component, also policy components can be composed or decomposed according to the generic component model. Using HL7 version 3 data type definitions, the policy class can be specialized

into basic policy, meta policy, and composite policy (see Figure 4), which have been verbally explained in detail in Table 3 and Table 4, respectively.[6]

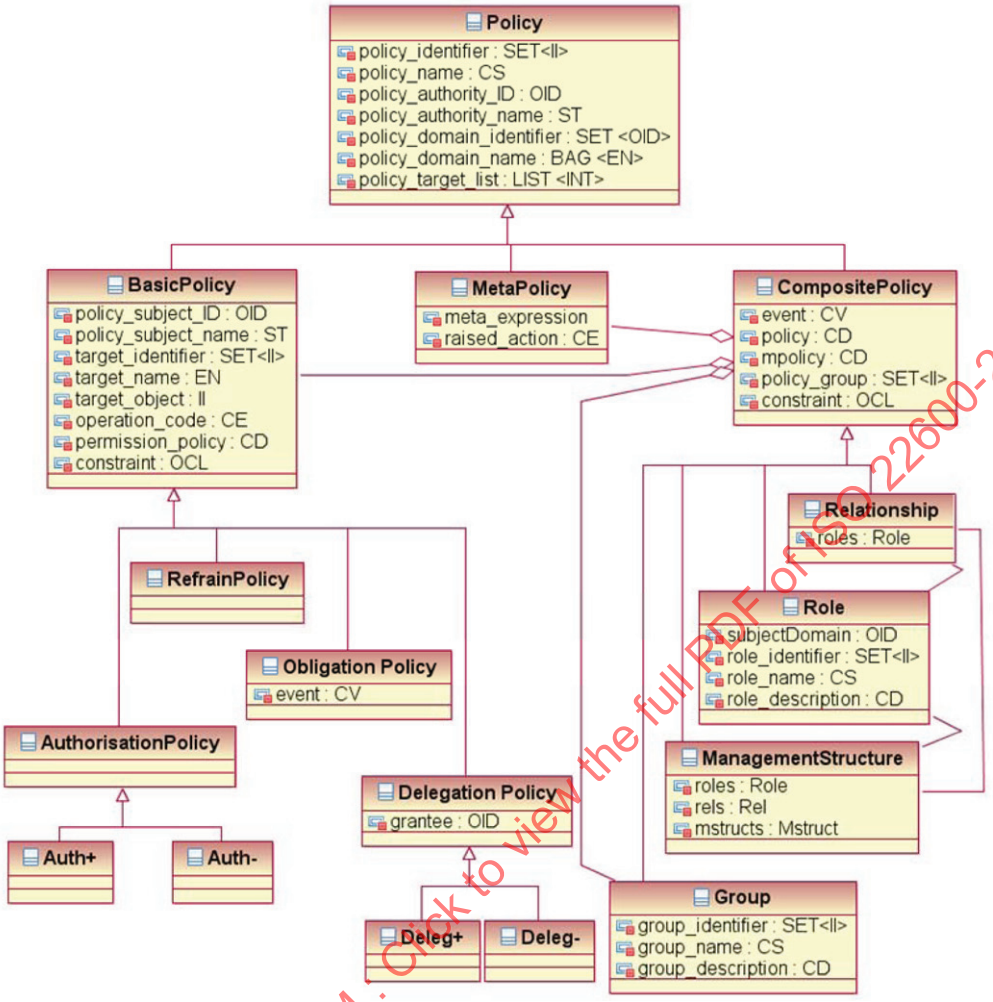

Figure 4 — Policy base-class diagram

The specializations of the composite policy abstract class are interrelated in a complex way, which has been indicated in outlines as simple association.

Table 3 — Basic policy types[6]

Basic policy type	Purpose	Content
Authorization policies	Define permitted actions	Subject (except in roles), target, action
Obligation policies	Are event triggered and define actions to be performed by manager agents	Subject (except in roles), action, event
Refrain policies	Define actions the subjects shall refrain from performing	Subject (except in roles), action
Delegation policies	Define what authorizations can be delegated to whom	

Table 4 — Composite policy types[6]

Composite policy type	Purpose
Groups	Define a scope for related policies to which a set of constraints can apply
Roles	Define a group of policies (authorization, obligation, and refrain policies) (For details on roles, see 6.5 and Annex A.)
Relationships	Define a group of policies pertaining to the interactions between a set of roles

Another way for policy decomposition has been provided by the OMG's security services specification distinguishing between the following policies:

— invocation access policy implementing access control policy for objects;

— invocation audit policy controlling event type and criteria for audit;

— secure invocation policy specifying security policies 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.

One common way to express constraints is the specification of user-defined schemata such as XML schemata. This schema should be standardized for interoperability purposes mentioned above.

Figure 5 presents a simple XML instance for a security policy statement.

```
<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>
```

Figure 5 — Policy template example

Policies shall be managed and stored in standardized trustworthy policy repositories.

6.5 Role model

For managing relationships between the entities mediated by an activity, two different roles have to be defined: organizational or structural roles at the entity's side and functional roles at the act's side (see Figure 6).

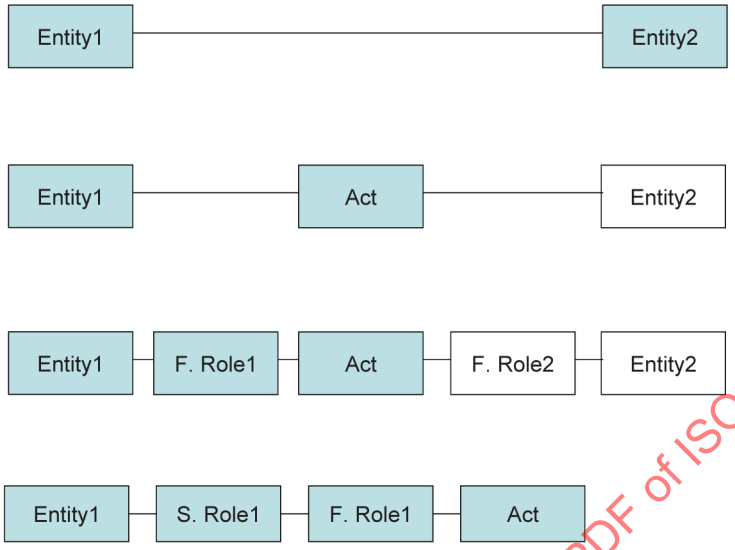

Figure 6 — The generic role concept

Because the role concept is deployed in many different contextual relationships such as the professionals' administration, certification procedures, roster management, etc., the role model and its deployment has been separately defined in ISO 21298.

For enabling the deployment of this International Standard, the currently available main aspects of ISO 21298 are presented in Annex A. Being under development, ISO 21298 is a living document. Therefore, the reader is encouraged to inform himself about the matured role aspects by considering the version of ISO 21298 available at the respective time.

6.6 Authorization model — Role and privilege assignment

Credentialing, privileging, and authorization are performed by connecting roles to policies.

Roles provide a means to indirectly assign privileges to individuals. Individuals are issued role assignment certificates assigning one or more roles to them by role attributes. Privileges are assigned to a role by role specification certificates, rather than to individuals. The indirect assignment enables the privileges assigned to a role to be updated without impacting the certificates that assign roles to individuals. Role assignment certificates can be attribute certificates or public key certificates. Role specification certificates cannot be public key certificates, but shall use other technologies such as attribute certificates or SAML technology. If role specification certificates are not used, the assignment of privileges to a role can be done through other means (e.g. can be locally configured at a privilege verifier).

The following scenarios are all possible.

— Any number of roles can be defined by any attribute authority (AA).

— The role itself and the members of a role can be defined and administered separately, by different AAs. This implies that roles can be local within a domain, e.g. on an organizational, regional, or national level.

— Role membership, just as any other privilege, can be delegated.

— Roles and membership can be assigned any suitable lifetime.

If, e.g. the role assignment certificate is an attribute certificate, the role attribute is contained in the attributes component of the attribute certificate. If the role assignment certificate is a public key certificate, the role attribute is contained in the subjectDirectoryAttributes extension. In the latter case, any additional privileges contained in the public key certificate are privileges that are directly assigned to the certificate subject. Thus, a privilege asserter can present a role assignment certificate to the privilege verifier demonstrating only that the privilege asserter has a particular role (e.g. "manager" or "purchaser"). The privilege verifier can know a priori, or can have to discover by some other means, the privileges associated with the asserted role in order to accept/reject/modify a request. The role specification certificate can be used for this purpose.

A privilege verifier shall have an understanding of the privileges specified for the role. The assignment of those privileges to the role can be provided within the privilege management infrastructure (PMI), e.g. by a role specification certificate, or outside the PMI (e.g. locally configured). For role privileges asserted in a role specification certificate, mechanisms for linking that certificate with the relevant role assignment certificate for the privilege asserter are provided in this part of ISO 22600. The issuer of the role assignment certificate can be different from the issuer of the role specification certificate and these certificates are administered (e.g. creation, expiration, revocation) entirely separately. The same certificate (attribute certificate or public key certificate) can contain a role assignment certificate as well as contain assignment of other privileges directly to the same individual. However, a role specification certificate shall be a separate certificate.

NOTE The use of roles within an authorization framework can increase the complexity of path processing, because such functionality essentially defines another delegation path which has to be followed. The delegation path for the role assignment certificate can involve different AAs and can be independent of the AA that issued the role specification certificate.

The general privilege management model consists of three entities: the object, the privilege asserter, and the privilege verifier. Request might be authorized, denied, or modified.

6.7 Control model

Access control is the process which determines if a claimant's privileges permit him/her/it to access a service provided by a target component. In this context, access is broader than acquiring some data. It might refer to any service offered by a target component (e.g. deletion, computation, transfer).

The control model illustrates how control is exerted over access to a sensitive object operation. There are four components in the model: the claimant, the verifier, the target, and the control policy (see Figure 7). The claimant has privilege attributes, contained in an attribute certificate. The target has constraining attributes, which might be contained in a security label, attribute certificate, or as record or list in a local database. The techniques described here enable the verifier, who might be the accountable party of the target or an independent authority, to control access to the target by the claimant, in accordance with the control policy and optionally taking other environmental variables or components into account (e.g. local time).

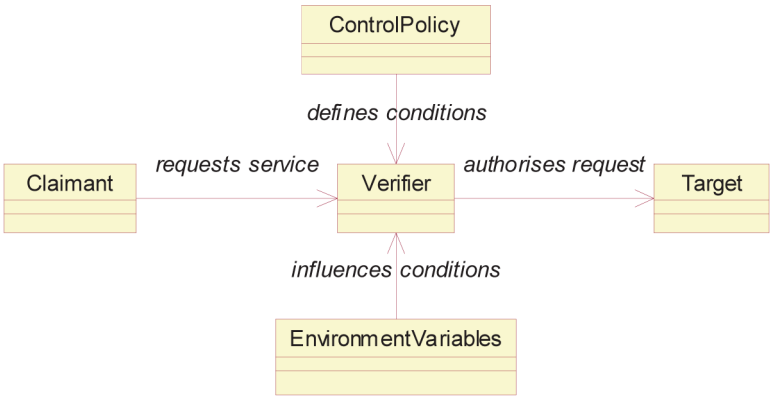

Figure 7 — Control model

The claimant's privileges are typically encapsulated in its attribute certificate. This can be presented to the verifier in the service request (push strategy), or it can be distributed by some other means, such as via a directory (pull strategy). The control policy shall be protected for integrity and authenticity and, for this purpose, it might sometimes be combined with the claimant's privilege in an attribute certificate. Normally, however, it will be declared separately.

The claimant can be an entity identified by a public key certificate, or an executable object identified by the digest.

6.8 Delegation model

In addition to the control model, there is a need for a delegation model. There are three components of the delegation model: the verifier, the source of authority, and the claimant (see Figure 8).

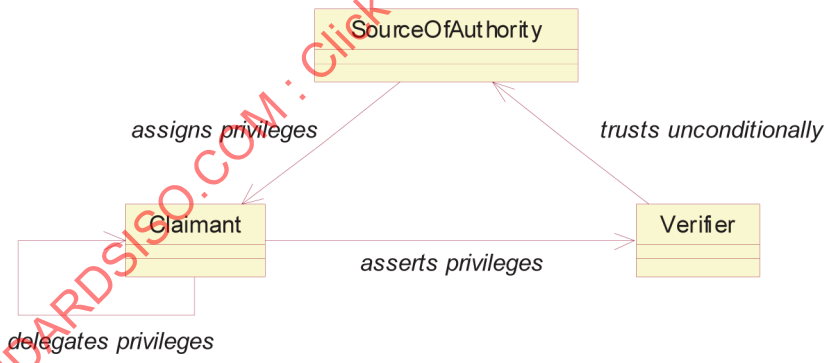

Figure 8 — Delegation model

The verifier endows an entity known as the source of authority with global privilege within the context a delegation occurs. The source of authority is known as attribute authority. It delegates privilege to claimants by issuing attribute certificates. The claimant asserts its delegated privileges by providing its credentials. This can be done by proving its knowledge of a private key whose public counterpart is contained in a public key certificate referenced by an attribute certificate which includes the claimed privilege.

Optionally, the claimant can delegate its privilege to another claimant. The verifier shall confirm that all entities in the delegation path possess sufficient privilege to access the target requested by the direct claimant.

The source of authority can also process a request from an entity to delegate its privilege by issuing an attribute certificate to another entity. However, this process is outside the scope of this part of ISO 22600.

The claimant and the verifier can be entities in different security domains. In such cases, the source of authority might be located in the verifier's domain, and a continuous section of the delegation path, which includes the direct claimant, shall be in the other security domain.

The delegation path is distinct from the certificate validation path used to validate the public key certificates of the entities involved in the delegation process. However, the quality of authenticity offered by the public key certificate validation process shall be commensurate with the sensitivity of the target being protected.

Specifying interoperability between distributed objects or components respectively, the Object Management Group has defined an alternative delegation model within its CORBA security services specification. In an object system, a client calls on an object to perform an operation, but this object will often not complete the operation itself, so it will call on other objects to do so. This will usually result in a chain of calls on other objects. (For further details, see www.omg.org)

In privilege delegation, the initiating principal's access control information (i.e. its security attributes) can be delegated to further objects in the chain to give the recipient the privileges to act on its behalf under specified circumstances (see the OMD delegation schemes given in Table 5).

Another authorization scheme is reference restriction where the permission to use an object under specified circumstances is passed as part of the object reference to the recipient. Reference restriction is not included in this part of ISO 22600.

The following terms are used in describing OMG's delegation options:

— initiator: the first client in a call chain;

— final target: the final recipient in a call chain;

— intermediate: an object in a call chain that is neither the initiator nor the final target;

— immediate invoker: an object or client from which an object receives a call.

Table 5 — Delegation schemes (OMG)

Intermediate performs	Target		Constraints
a) one method on one object			
b) several methods on one object			
c) any method on:	1) one object		none
	2) some object(s)		target restrictions
	3) any object		no target restrictions
	using	no privileges	simple delegation
		a subset of the initiator's privileges	
		both the initiator's and its own privileges	composite delegation
		received privileges and its own privileges	combined or traced delegation, depending on whether privileges are combined or concatenated
	during some validity period		part of time constraints
	for a specified number of invocations		part of time constraints

Communication of health information is frequently connected with a supplier chain performing this activity (e.g. involvement of secretaries, clerks, service departments, but also any other principals). This delegation model shall be used for any such chaining of services.

6.9 Access control model

The use of roles can greatly simplify user account administration and access control. Additionally, administration constraints might need to be enforced. For example, the separation of duties might be introduced as authorization constraint widely used.

Basic elements for access control management are principals, roles, permissions, operations, and objects. The access control management is characterized by the following components:

— definition of roles and role constraints;

— user-role assignment;

— role-permission assignment;

— assignment of constraints for activation of user assigned roles.

Harmonizing the role models specified in section 6.5 and Annex A and advanced access control models such as the NIST standard role-based access control, Figure 9 has been developed presenting an adapted role-based access control schema.

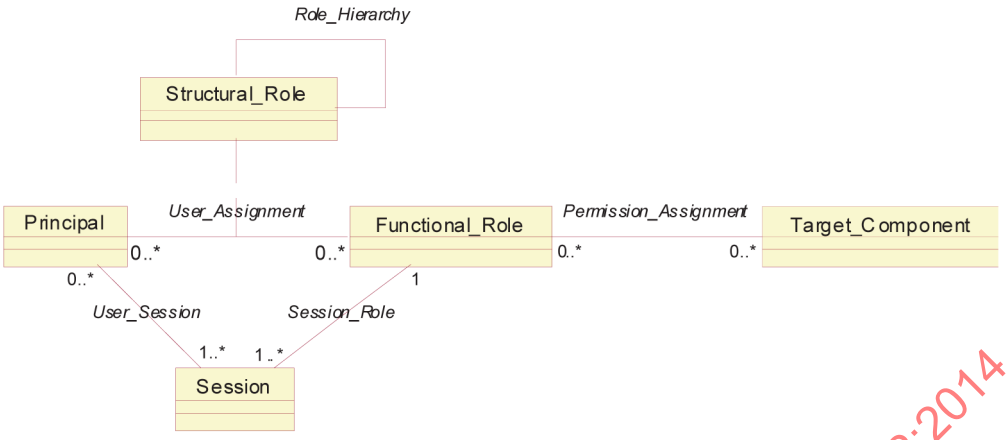

Figure 9 — Role-based access control schema

The RBAC schema defines the assignment of permissions dedicated to a functional role which has been assigned to a principal within a certain session. The functional role might be qualified by a set of structural roles assigned to the same principal. Simplified RBAC models just define the three components "user" (specialized principal), "role", and "permission" associated by the UserAssignment and the PermissionAssignment, respectively. The permission addresses the performance of certain operations on certain target objects. For mandatory access control (MAC), a top-level authority defines a static permission framework — the mandates, discretionary access control enables the accountable party of a target object to assign privileges to a certain user. Here, the privilege to assign privileges is the highest permission level, which should be carefully deployed.

Each model component is defined by the following subcomponents:

— a set of basic element sets;

— a set of RBAC relations involving those element sets (containing subsets of Cartesian products denoting valid assignments);

— a set of mapping functions that yield instances of members from one element set for a given instance from another element set.

Defining constraints on roles, processes, target objects, and related privileges by policies, Figure 9 turns into Figure 3 of ISO 22600-3, here numbered Figure 10.

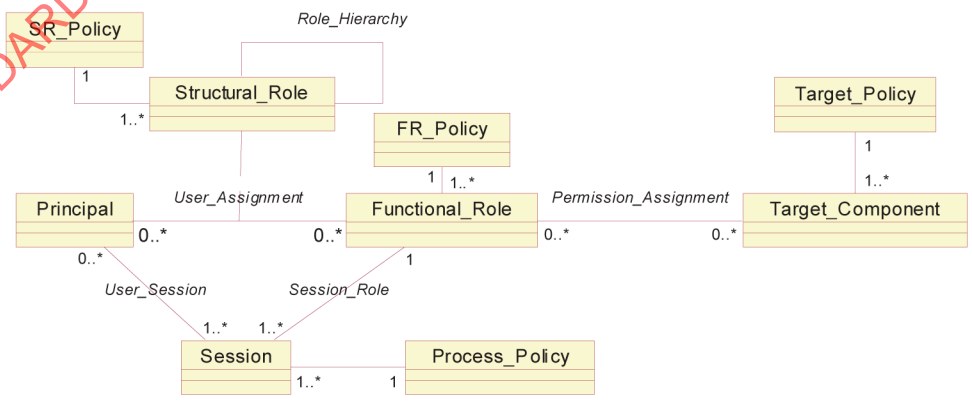

Figure 10 — Policy-driven RBAC schema

Annex A
(informative)

Functional and structural roles

A.1 Healthcare-related roles

For managing relationships between the entities, roles can be assigned to any principal. Principals are the actors in healthcare. Therefore, roles are associated to actors and to acts.

In general, two types of roles can be distinguished: structural roles and functional roles. Structural roles reflect the structural aspects of relationships between entities. Structural roles describe prerequisites, feasibilities, or competences for acts. Functional roles reflect functional aspects of relationships between entities. Functional roles are bound to the realization/performance of acts, where actions might be concatenated to an activity or even to a process.

Considering both structural roles and functional roles in the same context, structural roles provide the prerequisites/competences for entities to perform interactions (an act) within their specific functional roles. Qualifications, skills, etc. influence both the assignment of the structural roles and the performance of activities according to their functional roles (see Figure A.1).

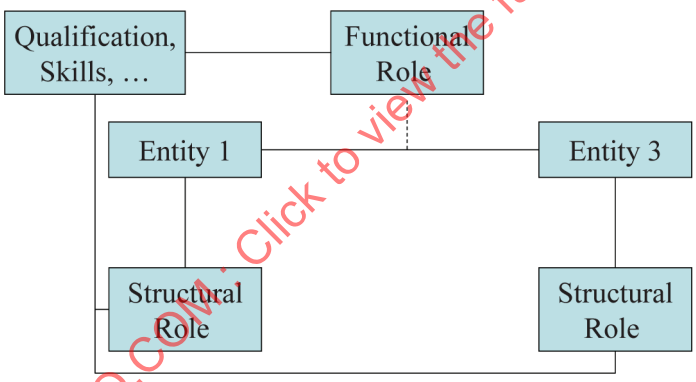

Figure A.1 — Generic role concept

Possible examples for structural roles of healthcare professionals are

— medical director,

— director of clinic,

— head of the department,

— senior physician,

— resident physician,

— physician,

— medical assistant,

— trainee,

— head nurse,

— nurse, and

— medical student.

Examples for “atomic” functional roles are

— prescriber,

— signer, and

— investigator (test performer).

They can be aggregated in relation to a process (admission, care of a certain subject, etc.). Possible examples for functional roles of healthcare professionals are

— caring doctor (responsible doctor),

— member of diagnostic team,

— member of therapeutic team,

— consulting doctor,

— admitting doctor,

— family doctor, and

— function-specific nurse.

A.2 Functional role model

Regarding the healthcare business process, functional roles can be defined in levels of authorizations and permissions in the following generic way reusing slightly changed definitions established in the Australian HealthNet Project, cross-referenced against other works:

— subject of care (normally the patient);

— subject of care agent (parent, guardian, carer, or other legal representative);

— responsible (personal) healthcare professional (the healthcare professional with the closest relationship to the patient, often his GP);

— privileged healthcare professional

  — nominated by the subject of care or

  — nominated by the healthcare facility of care (there is a nomination by regulation, practice, etc.);

— healthcare professional (involved in providing direct care to the patient);

— health-related professional indirectly involved in patient care, teaching, research, etc.;

— administrator (and any other parties supporting service provision to the patient).

This list fixes the set functional roles applied to manage the creation, access, processing, and communication of health information.

Additionally, functional roles can be grouped according to the relation to the information created, collected, stored, used, disclosed, retained, and destroyed:

— composer;

— committer;