

First edition
2012-07-01

Health informatics — Clinical document registry federation

*Informatique de santé — Fédération d'enregistrement de documents
cliniques*



Reference number
ISO/TR 13128:2012(E)

© ISO 2012



COPYRIGHT PROTECTED DOCUMENT

© ISO 2012

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

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

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

In exceptional circumstances, when a technical committee has collected data of a different kind from that which is normally published as an International Standard ("state of the art", for example), it may decide by a simple majority vote of its participating members to publish a Technical Report. A Technical Report is entirely informative in nature and does not have to be reviewed until the data it provides are considered to be no longer valid or useful.

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.

ISO/TR 13128 was prepared by Technical Committee ISO/TC 215, *Health informatics*.

Introduction

A federation of Clinical Document Registries (CDRs) provides a view that the entire set of member CDRs can be seen as a single CDR. One can issue a federated query to any of the member CDRs and still receive the consistent answer, which is in effect the same as the union of the individual answers from the CDRs. Therefore, it offers a layer of location transparency.

Within a federation, a pair of CDRs can mutually agree to replicate the entire content of one CDR in the other CDR. The original and replica are systematically synchronized to reflect any changes to the original ever since it was replicated. This is called the “federated replication”. The key benefits of the federated replication are as follows.

- a) It greatly enhances the availability and fault-tolerance of CDRs at the presence of a failure, which may happen due to various kinds of catastrophic events such as fire, earthquake, flood, and even shortage of network connectivity. When a failure of a CDR occurs and is detected, a query to the CDR can be answered alternatively by another CDR if it maintains a federated replica of the failed CDR.
- b) It also improves the performance of processing of federated queries. As previously stated, a federated query can be sent to any of the member CDRs of the intended federation. If the receiving CDR maintains a federated replica of another member CDR, it can answer the query using the replica without relaying it to the original CDR, saving lots of network communication overheads.

The replication is done at the object level only. There is no specification as to how the replica can be used for answering federated queries.

The primary goal of this Technical Report is to define a logical structure for the federation of CDRs (but not repositories) that, therefore, does not require any conformance criteria, and to set out a list of basic functional recommendations and configuration parameters for CDR federations, possibly with replication, as well as to provide several use cases.

Health informatics — Clinical document registry federation

1 Scope

This Technical Report covers aspects of federations of Clinical Document Registries (CDRs) including cases where the whole content of a CDR has been replicated in another CDR of the same federation. More specifically, the following aspects are covered.

- a) Recommendations for the creation and management of federations of CDRs and federation metadata.
- b) Creation and management of federated replicas, including specifications to create, remove and maintain federated replicas as well as metadata for them.
- c) Use cases of federations with query processing examples, including various use cases of federations of CDRs, together with examples of query processing policies to enhance the performance and fault-tolerance.
- d) Processing of queries in the presence of federations and replicas: the presence of federated replicas might affect the semantics of both local and federated queries. Informative examples to define exact behaviours of processing the queries are given.

NOTE It is assumed that the problem of patient identification has been solved in a way that is beyond the scope of this Technical Report, such as IHE XCPD (Cross-Community Patient Discovery) Profile. Some potential issues that will not be addressed in this Technical Report include patient identity management, potential limitations of registries due to jurisdictional policies or requirements, and how replicates are handled.

2 Normative references

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

ISO/TS 27790, *Health informatics — Document registry framework*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/TS 27790 and the following apply.

3.1

registry federation

group of registries that have voluntarily agreed to form a union

NOTE 1 Such a federation may be based on common business interests and specialties that the registries may share.

NOTE 2 Registry federations appear as single logical registry to registry clients.

3.2

object replication

act of duplicating an object within a registry to another registry

3.3

object replica

duplicated copy of an object by object replication

3.4

federated replication

act of duplicating a set of all objects within a registry to another registry belonging to the same federation

NOTE The original and duplicated set of objects should be maintained to be synchronized to each other.

3.5

federated replica

duplicated set of objects by federated replication

3.6

local query

query to a registry that has the limited range of search involving only the objects that have been originally registered to the registry

NOTE It does not involve any replica that is duplicated within the registry or further relayed to other registries.

3.7

local query with replica

local query that has the range of search involving all objects residing in the registry, regardless of whether they are original or duplicated, and that is not further relayed to other registries

3.8

federated query

query to a registry that has the range of search involving all objects within registries of the specified federation

NOTE The queried registry performs a local query (possibly with replica) and broadcasts the same query to all other members of the federation, only as a local query (possibly with replica).

4 Symbols and abbreviated terms

| | |
|-----------|--|
| ANSI | American National Standards Institute |
| CDA | Clinical Document Architecture |
| CDR | Clinical Document Registry |
| ebXML | Electronic Business Extensible Markup Language |
| HL7 | Health Level 7 |
| IHE | Integrating Healthcare Enterprises |
| OASIS | Organization for the Advancement of Structured Information Standards |
| HL7 RIM | HL7 Reference Information Model |
| ebXML RIM | ebXML Registry Information Model |
| RS | Registry Service |
| SOAP | Simple Object Access Protocol |
| XML | Extensible Markup Language |

5 Recommendations for federation

To be seen as a single CDR to a client, member CDRs of a federation should follow the following recommendations.

- a) At least one of the registries forming a federation should be capable of creating one or more federations, and all registries should be capable of joining one or more federations.

- b) A CDR should be capable of maintaining federation-related metadata consisting of metadata for the federation itself and members of the federation. The membership metadata also includes information on federated replicas.
- c) Each CDR that is a member of a federation should maintain metadata about the federation of which it is a member.
- d) A CDR should be capable of creating, maintaining, and removing federated replicas.
- e) Upon receipt of a federated query, a CDR should relay the query to other federation members based on federation metadata and return combined results to the client.
- f) A CDR should be able to take advantage of federated replica that it maintains, to enhance the performance of query processing.

The CDRs that form a federation in principle are assumed to be heterogeneous. Therefore the above recommendations demand additional capabilities and interfaces to existing (local) registries. The metadata associated with federation should not be confused with metadata associate with documents

A CDR federation can be either *centralized* or *decentralized*. In a centralized federation, there is an agent (either one of the member CDRs or a separate entity) that is in charge of maintaining federation-related metadata and routing federated queries. On the other hand, in a decentralized federation, all member CDRs act as agents on their own.

6 Modes of federated replication

Federated replication can be configured several different ways depending upon the following parameters. First, the replication can be either *partial* or *full*. In a partial replication, only part of the content of a CDR is replicated to another CDR, while the whole content is replicated in the full replication mode. Second, the replication can be updated in *real-time* or *scheduled batch* mode. In the former, any changes in the content of a CDR including additions, deletions, or updates of registry objects are immediately reflected to the replica as well. In the latter, the changes are logged and the reflection is delayed until a pre-determined time, after which it is completed. Third, the replication can be done either in a *hierarchical* or *peer-to-peer* (P2P) fashion. In the hierarchical mode, member CDRs are arranged as a tree and the content of a CDR is only replicated to its predecessors in the hierarchy. On the other hand, in the P2P mode, the content of a CDR is replicated to one or more CDRs that have been designated as peers in advance.

7 Use cases for federation

This clause describes some of the use cases where the CDR federation (with or without replication) plays an important role.

7.1 Clinical document sharing via regional federation without replication

Healthcare providers located in a specific region may wish to form a network with a tertiary hospital as the regional centre so that clinical information of their patients can be shared among them, as shown in Figure 1. Assuming each healthcare provider is running its own CDR, organizing the CDRs as a single federation can achieve the goal, without necessitating the consolidation of all clinical data at a single CDR, say the regional centre. When a patient visits any of the member care providers, all clinical documents can be retrieved through a single federated query with the consent of the patient. In this use case, IHE XCA may be used.

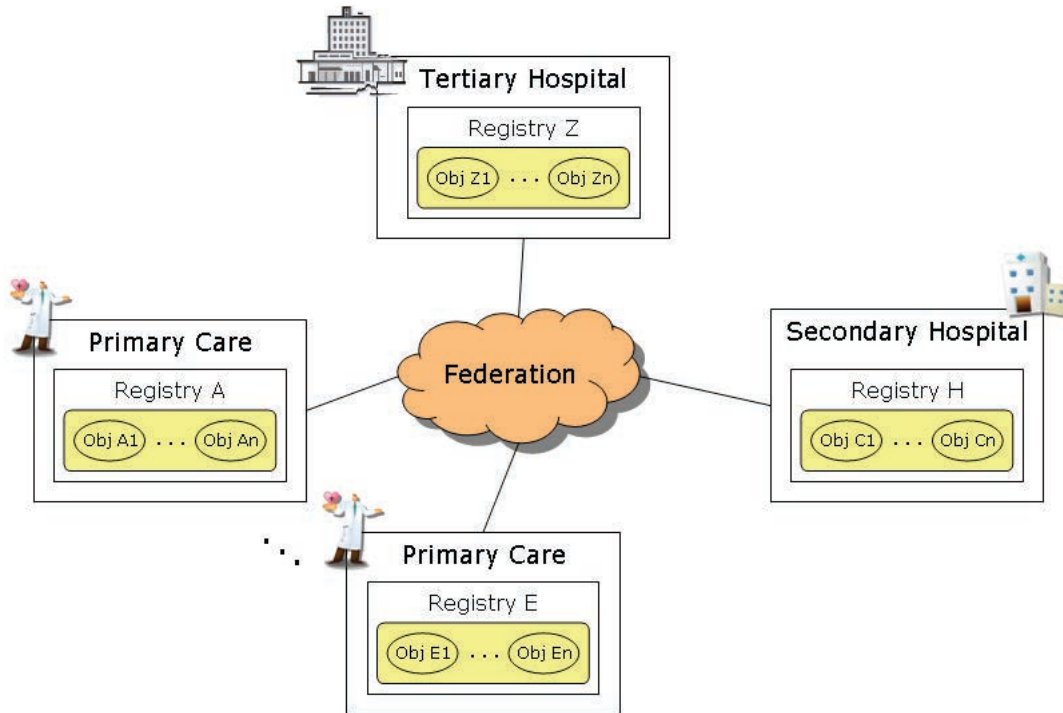


Figure 1 — Use case of clinical document sharing via regional federation without replication

7.2 Clinical document sharing via regional federation with replication

In this use case, the setting is similar to the previous one: healthcare providers located in a specific region form a federation to share the clinical information of their patients. Primary care providers usually cannot afford expensive facilities to protect their patient records as secondary or tertiary hospitals do. If unforeseen disastrous events (such as fire, earthquake, flood, etc.) occur, the whole or part of their data can be lost. Also, the primary care providers very often keep their computers on only during their office hours, beyond which patient records are not available to other care providers, for instance, the emergency room at a tertiary hospital.

If the contents of the CDRs of primary care providers are replicated somewhere else, say a regional centre, as shown in Figure 2, the data can be both protected in the presence of disasters and available even after their office hours, because tertiary or secondary hospitals tend to have better computing facilities that are operational on 7-24 basis.

In this use, IHE XCA can be used for non-federated queries and might be used for federated queries as well (to be further analysed).

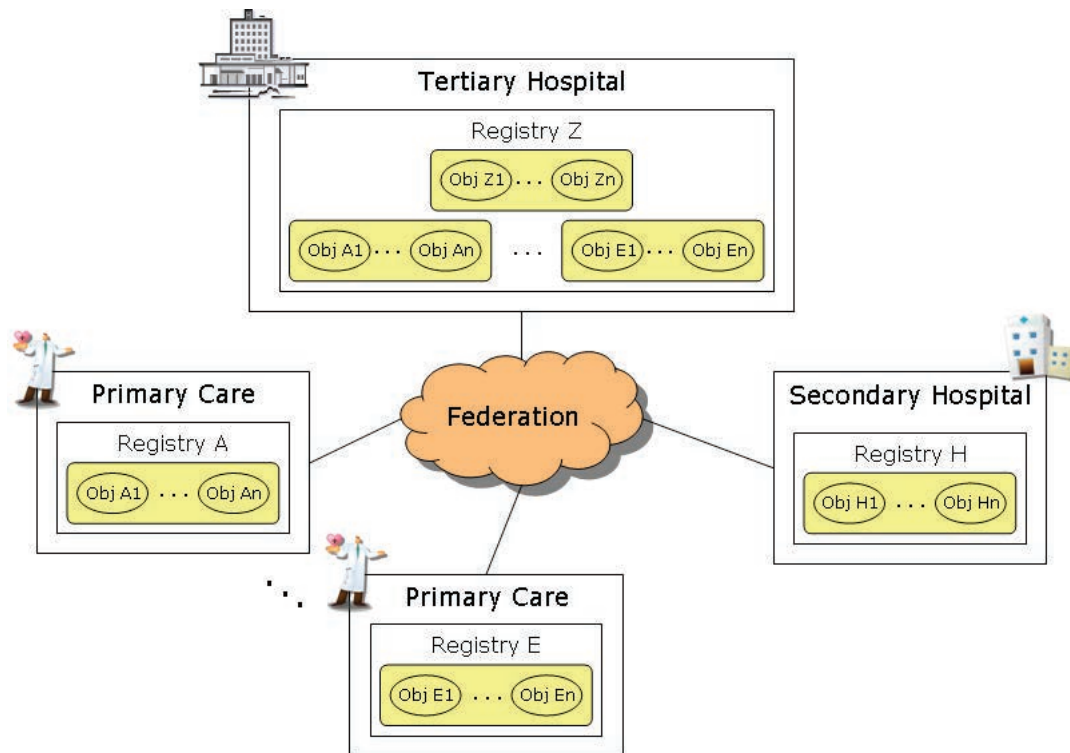


Figure 2 — Use case of clinical document sharing via regional federation with replication

7.3 National federation of hospitals for research purposes with replication

Suppose a government is planning to set up a network of hospitals with repositories of human-originated specimens such as tissue or blood samples. Each hospital acts as a regional centre and maintains a CDR for CDA documents containing information about the specimen they store. The government wishes to maintain copies of the CDRs of the hospitals, as well as its own CDR for specimens sampled at government-sponsored institutions, as shown in Figure 3. The primary purpose of the network is to provide nation-wide specimen information to researchers so that they can search and ask for allocations of the specimens of interests for their experiments. A key recommendation is that the whole network of the CDRs should be seen as a single unified CDR so that anyone can cast a query to any of the hospitals and the government and still get the same result. Hence a federation of CDRs is recommended. This is an exemplary case of hierarchical replication.

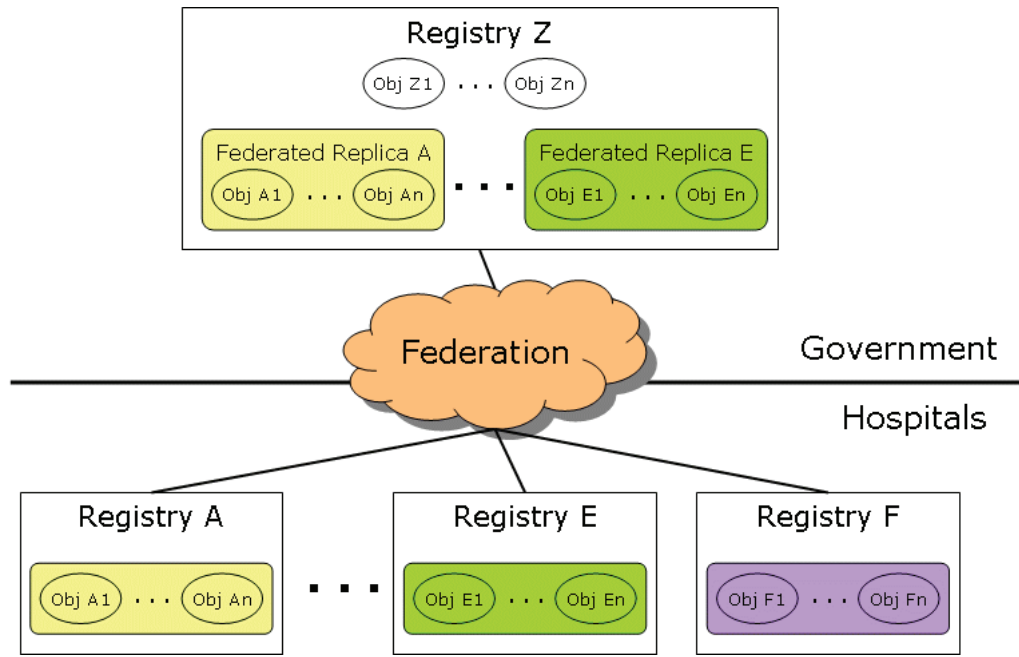


Figure 3 — Use case of national federation of hospitals for research purposes

7.4 National federation of regional registries for clinical purposes with replication

Clinical documents are shared by hospitals and clinics within a region via a centralized regional CDR. The regional CDRs across a nation are then federated. The network infrastructure may be poor as it is often the case in developing countries, and therefore federated queries are either very slow or cannot even be processed due to frequent network disconnections. To mitigate the problem, it is recommended that registry objects be replicated as illustrated in Figure 4. Synchronization of the original and replica is performed in scheduled batch mode to reduce network traffic during business hours. Regions can adopt different federation policies (centralized or decentralized) depending upon specific regional requirements.

Note that the fault-tolerance achieved by replicating registry objects applies only to registry objects (metadata about documents), but not to documents. Achieving fault-tolerance on the documents stored in repositories is beyond the scope of this technical report.

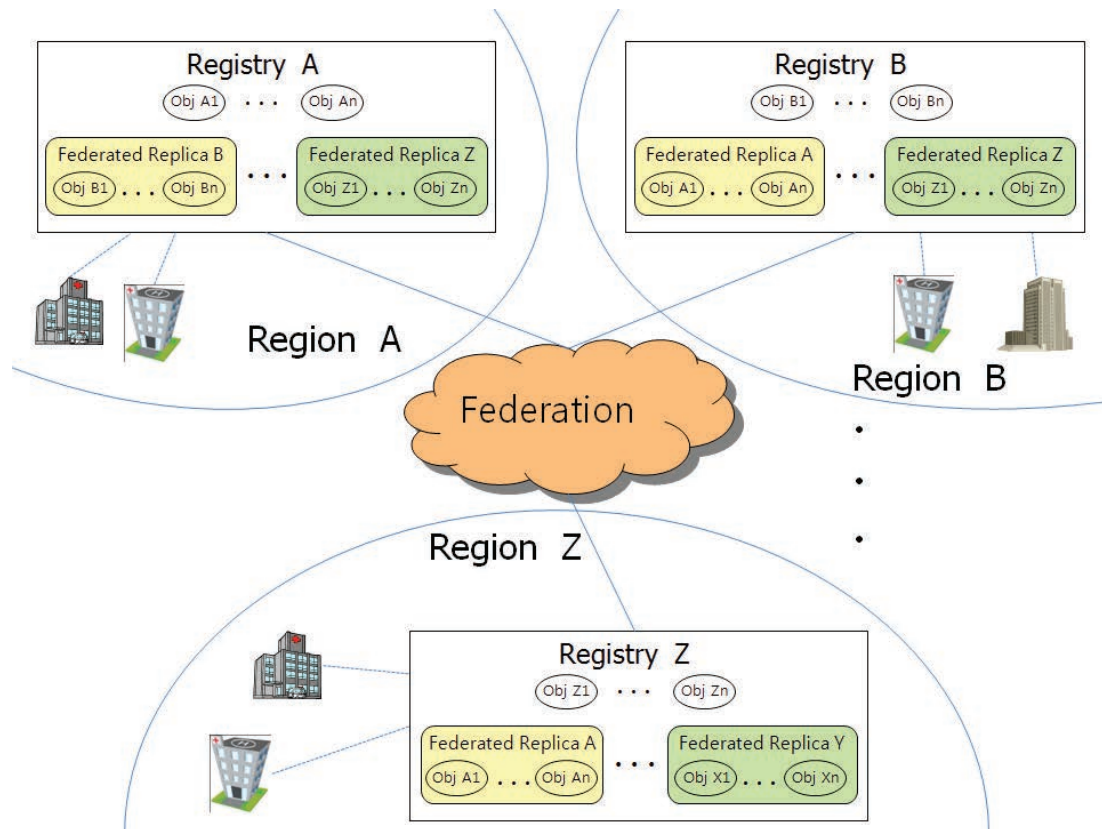


Figure 4 — Use case of national federation of regional registries for clinical purposes

8 Query processing

Queries received by a CDR are classified into two categories: local versus federated queries. The exact semantics of both types of queries depends on specific needs, yet some policies can be exemplified as follows.

a) Local query

- **LQ()**, **LQ(replicaSearch = false)**: Search only original local objects.
- **LQ(replicaSearch = true)**: Search both original local objects and federated replicas, if any.
- **LQ(replicaSearch = true, replica set)**: Search both original local objects and the specified set of federated replicas only.
 - **replica set**: set of federated replica IDs.
 - E.g., **LQ(replicaSearch = true, replica A | replica B)**
- **LQ(replicaSearch = true, replica set, originalSearch = false)**: Search the specified set of federated replicas, but not the original local objects.

b) Federated query

- FQ(specific-policy)
- Policy 1: Perform a local query on the local registry and route the query to all members of the federation as a local query. If a registry does not respond, it is skipped.
 - local registry: **LQ()**

- federation member registries: **LQ()**
- Policy 2: Perform a local query on the local registry and route the query to all members of the federation as a local query. If a registry does not respond and its replica exists in another registry, route the query to the registry that maintains the replica (it could be the local registry). This policy enhances fault-tolerance by answering using replicas on behalf of the un-responding registries.
 - local registry: **LQ()**
 - federation member registries: **LQ()**
 - replica holder: **LQ(replicaSearch = true, replica N, originalSearch = false)**
- Policy 3: Perform a local query with replica on the local registry and route the query to only those members of the federation whose replicas do not exist in the local registry. This policy enhances performance by avoiding unnecessary query routings. If a registry does not respond, it is skipped.
 - local registry: **LQ(replicaSearch = true)**
 - federation member registries that have not been replicated: **LQ()**
- Policy 4: Perform a local query with replica on the local registry and route the query to only those members of the federation whose replicas do not exist in the local registry. If a registry does not respond and its replica exists in another registry, route the query to the replica holder. This policy is a mix of Policy 2 and 3, and enhances both fault-tolerance and performance.
 - local registry: **LQ()**
 - federation member registries that have not been replicated: **LQ()**
 - replica holder: **LQ(replicaSearch = true, replica N, originalSearch = false)**

Bibliography

- [1] OAISIS standard, ebXML Registry Services and Protocols Version 3.0
- [2] IHE IT Infrastructure Technical Framework, Cross-Enterprise Document Sharing (XDS.b) Profile
- [3] IHE IT Infrastructure Technical Framework, Cross Community Access (XCA) Profile
- [4] IHE IT Infrastructure Technical Framework, Cross Community Patient Discovery (XCPD) Profile
- [5] IHE White Paper, Cross Community Information Exchange including Federation of XDS Affinity Domains

ICS 35.240.80

Price based on 9 pages