

Integrating the Healthcare Enterprise



5

**IHE IT Infrastructure (ITI)
Technical Framework**

**Volume 1
(ITI TF-1)**

10

Integration Profiles

15

20

**Revision 11.0 – Final Text
September 23, 2014**

25

Please verify you have the most recent version of this document, which is published [here](#).

CONTENTS

1	Introduction	11
	1.1 Overview of the Technical Framework	11
30	1.2 Overview of the IT Infrastructure Volume 1	12
	1.3 Audience	12
	1.4 Relationship to Standards	13
	1.5 Relationship to Real-world Architectures.....	13
	1.6 Conventions	14
35	1.6.1 IHE Actor and Transaction Diagrams and Tables.....	14
	1.6.2 Process Flow Diagrams	14
	1.6.3 Technical Framework Cross-references	14
	1.7 Scope of Changes Introduced in the Current Year	15
	1.8 Security Implications	18
40	1.9 Comments	18
	1.10 Copyright Permission.....	19
	1.11 IHE Technical Framework Development and Maintenance Process	19
2	IT Infrastructure Integration Profiles	20
	2.1 Dependencies among Integration Profiles	20
45	2.2 Integration Profiles Overview.....	23
	2.2.1 This section is reserved.....	23
	2.2.2 Sharing Value Set Integration Profile (SVS).....	23
	2.2.3 Retrieve Information for Display (RID)	24
	2.2.4 Enterprise User Authentication (EUA).....	24
50	2.2.5 Patient Identifier Cross-referencing (PIX).....	24
	2.2.6 Patient Synchronized Applications (PSA).....	24
	2.2.7 Consistent Time (CT)	25
	2.2.8 Patient Demographics Query (PDQ)	25
	2.2.9 Audit Trail and Node Authentication (ATNA)	25
55	2.2.10 Cross-Enterprise Document Sharing (XDS)	26
	2.2.11 Personnel White Pages (PWP)	26
	2.2.12 This section is intentionally left blank.....	26
	2.2.13 Cross Enterprise User Assertion (XUA)	26
	2.2.14 Patient Administration Management (PAM).....	26
60	2.2.15 Cross-Enterprise Document Reliable Interchange (XDR)	27
	2.2.16 Cross-Enterprise Document Media Interchange (XDM)	27
	2.2.17 Retrieve Form for Data Capture (RFD).....	27
	2.2.18 Cross-Community Access (XCA)	27
	2.2.19 Basic Patient Privacy Consents (BPPC).....	27
65	2.2.20 Scanned Documents Integration Profile (XDS-SD).....	27
	2.2.21 This section is reserved.	28
	2.2.22 This section is reserved.	28
	2.2.23 Patient Identifier Cross-referencing HL7 V3 (PIXV3)	28
	2.2.24 Patient Demographics Query HL7 V3 (PDQV3).....	28

70	2.2.25	Multi-Patient Queries (MPQ).....	28
	2.2.26	This section is reserved	28
	2.2.27	XCPD Integration Profile	28
	2.3	Product Implementations	29
3		Retrieve Information for Display (RID)	30
75	3.1	Actors/ Transactions	30
	3.2	Retrieve Information for Display Integration Profile Options.....	32
	3.3	Retrieve Information for Display Process Flow	32
4		Enterprise User Authentication (EUA)	36
	4.1	Actors/ Transactions	36
80	4.2	Enterprise User Authentication Integration Profile Options.....	38
	4.3	Enterprise User Authentication Profile Process Flow	38
	4.3.1	Basic User Authentication Process Flow.....	38
	4.3.2	User Authentication with User Synchronized Applications Process Flow	40
	4.3.3	Fast User Switching with Multiple Applications Process Flow	41
85	5	Patient Identifier Cross-referencing (PIX)	44
	5.1	Actors/ Transactions	46
	5.2	Patient Identifier Cross-referencing Integration Profile Options.....	47
	5.2.1	Pediatric Demographics.....	47
	5.3	Patient Identifier Cross-referencing Profile Process Flows	49
90	5.3.1	Use Case: Multiple Identifier Domains within a Single Facility/ Enterprise	49
	5.3.2	Use Case: Multiple ID Domains Across Cooperating Enterprises	50
	5.3.3	Pediatric Demographic Option Use Cases.....	52
	5.4	Relationship between the PIX Integration Profile and eMPI.....	53
6		Patient Synchronized Applications (PSA)	55
95	6.1	Actors/ Transactions	55
	6.2	Patient Synchronized Applications Integration Profile Options.....	56
	6.3	Patient Synchronized Applications Integration Profile Process Flows.....	56
	6.3.1	Use Case: Simple Patient Switching.....	56
7		Consistent Time (CT).....	59
100	7.1	Actors/ Transactions	59
	7.2	Consistent Time Integration Options	60
	7.3	Consistent Time Process Flow	60
8		Patient Demographics Query (PDQ).....	62
	8.1	Actors/ Transactions	62
105	8.2	Patient Demographics Query Integration Profile Options	63
	8.2.2	Pediatric Demographics.....	63
	8.3	Patient Demographics Query Process Flow.....	64
	8.3.1	Combined Use of PDQ with Other IHE Workflow Profiles	66
	8.3.2	Supplier Data Configuration.....	66
110	9	Audit Trail and Node Authentication (ATNA).....	68
	9.1	Authentication.....	69
	9.1.1	User Authentication	69
	9.1.2	Connection Authentication	69
	9.2	Audit Trails	70

115	9.2.1 Audit Messages.....	71
	9.2.2 Backwards Compatibility	72
	9.3 Audit Trail Transport	73
	9.4 Actors/Transactions	73
	9.5 ATNA Integration Profile Options	75
120	9.5.1 ATNA Encryption Option (retired)	76
	9.5.2 Radiology Audit Trail Option.....	76
	9.6 Audit Trail and Node Authentication Process Flow	76
	9.6.1 Normal Node Process Flow	76
	9.6.2 Unauthorized Node Process Flow	79
125	9.6.3 Unauthorized User Process Flow.....	80
	9.7 Relationship between Secure Node, Secure Application, and other Actors	80
	10 Cross-Enterprise Document Sharing (XDS.b)	82
	10.1 Actors/Transactions	84
	10.1.1 Actors	85
130	10.1.2 Transactions.....	86
	10.1.3 XDS Document Contents Support.....	88
	10.2 Integration Profile Options	89
	10.2.1 Document Replacement Option	90
	10.2.2 Document Addendum Option.....	90
135	10.2.3 Document Transformation Option	90
	10.2.4 Folder Management Option.....	90
	10.2.5 Asynchronous Web Services Exchange Option	90
	10.2.6 Reference ID.....	91
	10.2.7 On-Demand Documents Option.....	91
140	10.2.8 Persistence of Retrieved Documents Option.....	92
	10.3 Integration Profile Process Flow.....	92
	10.4 General Principles	98
	10.4.1 EDR-CR Concept	98
	10.4.2 XDS Document Concept	99
145	10.4.3 Submission Request.....	100
	10.4.4 Submission Set Concept.....	100
	10.4.5 Concept of Folder	101
	10.4.6 Example of use of Submission Request, Submission Set and Folder	101
	10.4.7 XDS Registry Data Model and Attributes	103
150	10.4.8 Concept of an XDS Affinity Domain	106
	10.4.9 Patient Identification Management.....	106
	10.4.10 Document Lifecycle	108
	10.4.11 Document Query	110
	10.4.12 Intentionally left blank.	110
155	10.4.13 XDS Document Entry Types	110
	10.5 Implementation Strategies	112
	10.6 Patient Identifier Communication Requirements.....	114
	10.7 Security Considerations	115
	10.7.1 Use of ATNA to address Basic Security	116

160	10.8 Intentionally Left Blank	117
11	Personnel White Pages (PWP)	118
	11.1 Actors/ Transactions	118
	11.2 PWP Integration Profile Options	119
	11.3 PWP Integration Profile Process Flow	120
165	12 This is intentionally left blank	121
13	Cross Enterprise User Assertion (XUA) Integration Profile	122
	13.1 Use Cases	123
	13.2 XUA Development	124
	13.3 Intentionally Left Blank	125
170	13.4 Actors/Transaction	125
	13.5 Options	126
	13.5.1 Subject-Role	127
	13.5.2 Authz-Consent	127
	13.5.3 PurposeOfUse	128
175	13.6 Grouping	128
	13.6.1 Audit Trail and Node Authentication (ATNA)	128
	13.6.2 Cross-Enterprise Document Sharing (XDS)	128
	13.6.3 Enterprise User Authentication (EUA)	129
	13.6.4 Any Web-Services Transaction that leverages ITI TF-2x: Appendix V	129
180	13.7 Process Flow	129
	13.8 Security Considerations	130
14	Patient Administration Management (PAM) Integration Profile	131
	14.1 Patient Administration Management Use Cases	131
	14.2 Patient Identity Management Use Case	131
185	14.2.1 Patient Encounter Management Use Case	131
	14.2.2 Actors/ Transactions	132
	14.3 Patient Administration Management Integration Profile Options	133
	14.3.1 Merge Option	134
	14.3.2 Link / Unlink Option	135
190	14.3.3 Inpatient / Outpatient Encounter Management Option	135
	14.3.4 Pending Event Management Option	135
	14.3.5 Advanced Encounter Management Option	135
	14.3.6 Temporary Patient Transfer Tracking Option	135
	14.3.7 Historic Movement Option	135
195	14.3.8 Acknowledgement Support	135
	14.3.9 Maintain Demographics Option	136
	14.3.10 Ambulatory Patient Data Option	136
	14.4 Patient Administration Management Integration Profile Actor Grouping	136
	14.4.1 Actor Grouping of Patient Encounter Supplier	136
200	14.4.2 Actor Grouping with other IHE Actors	137
	14.5 Patient Administration Management Process Flow	137
	14.5.1 Patient Identity Management	137
	14.5.2 Patient Encounter Management	139
	15 XDR Integration Profile	146

205	15.1 Actors/ Transactions.....	146
	15.2 XDR Integration Profile Options	147
	15.2.1 Intentionally Left Blank	147
	15.2.2 Basic Patient Privacy Enforcement Option	147
	15.2.3 Accepts Limited Metadata	147
210	15.3 XDR Process Flow	147
	15.4 Digital communication.....	149
	15.5 Security Considerations	149
	16 Cross-Enterprise Document Media Interchange (XDM) Integration Profile.....	150
	16.1 Actors/ Transactions	150
215	16.2 XDM Integration Profile Options	151
	16.2.1 USB Option	152
	16.2.2 CD-R Option	152
	16.2.3 ZIP over Email	152
	16.2.4 ZIP over Email Response	152
220	16.3 XDM Process Flow.....	152
	16.4 Digital communication.....	153
	16.4.1 Actual Media Type	153
	16.4.2 Virtual Media over a Network.....	154
	16.4.3 Media Content	154
225	16.5 Security considerations	154
	17 Retrieve Form for Data Capture (RFD) Profile	156
	17.1 Use Cases	157
	17.1.1 Investigational New Drug Clinical Trial Use Case	157
	17.1.2 Public Health Reporting Use Cases	159
230	17.1.3 Pharmaco-vigilance Scenario	160
	17.1.4 Cardiology Research Use Cases.....	161
	17.1.5 Radiology Use Case – Clinical Impact Registry	162
	17.1.6 Data Clarification	163
	17.2 Actors/ Transactions.....	164
235	17.2.1 Actors	166
	17.2.2 Transactions	166
	17.3 Retrieve Form for Data Capture Integration Profile Options.....	167
	17.3.1 Archive Form Option	167
	17.3.2 Data Clarifications Option	168
240	17.3.3 XForms Option.....	168
	17.4 Retrieve Forms for Data Capture Process Flow.....	168
	17.5 Security Considerations	173
	17.5.1 RFD Risk analysis Risk Assessment.....	173
	17.5.2 Recommendations	174
245	18 Cross-Community Access (XCA) Integration Profile	176
	18.1 Actors/ Transactions.....	176
	18.2 XCA Integration Profile Options	177
	18.2.1 XDS Affinity Domain Option	177
	18.2.2 Asynchronous Web Services Exchange Option.....	178

250	18.2.3 Grouping Rules	178
	18.2.4 On-Demand Documents	180
	18.2.5 Persistence of Retrieved Documents	181
	18.3 XCA Process Flow	181
	18.3.1 Use Cases	181
255	18.3.2 homeCommunityId defined.....	182
	18.3.3 Detailed Interactions	182
	18.4 XCA Security Considerations	187
	18.4.1 XCA Risk Assessment	187
	18.4.2 Requirements/Recommendations	188
260	18.4.3 Policy Choices	189
	19 Basic Patient Privacy Consents Integration Profile	190
	19.1 Basic Patient Privacy Consent Use-Cases	191
	19.1.1 Implied Consent vs. Explicit Consent	191
	19.1.2 Wet Signature	192
265	19.1.3 Advanced Patient Privacy Consents	193
	19.2 Creating Patient Privacy Policies.....	194
	19.2.1 Summary of the creation and publication of the policies	195
	19.3 Actors/Transactions	195
	19.3.1 Grouping.....	196
270	19.4 Basic Patient Privacy Consent Profile Options.....	197
	19.4.1 Intentionally Left Blank	197
	19.4.2 Intentionally Left Blank	197
	19.4.3 Basic Patient Privacy Acknowledgement Option.....	197
	19.4.4 Basic Patient Privacy Acknowledgement with Scanned Document Option	197
275	19.4.5 Patient Privacy Acknowledgement View Option.....	198
	19.5 Intentionally Left Blank.....	198
	19.6 BPPC Process Flow in an XDS Affinity Domain.....	198
	19.6.1 Checking for a patient’s acknowledgement of a privacy policy	198
	19.6.2 Recording a patient’s acknowledgement of a privacy policy.....	198
280	19.6.3 Publishing documents against a consent policy	198
	19.6.4 Using published documents.....	199
	19.7 Security Considerations	199
	20 Cross-Enterprise Sharing of Scanned Documents Content Integration Profile	201
	20.1 Use Cases	201
285	20.1.1 Content Use Cases	201
	20.1.2 Content Creator Use Cases	202
	20.1.3 Content Consumer Use Cases	203
	20.2 Actors/ Transactions	203
	20.3 Scanned Documents Content Integration Profile Options	204
290	20.4 Scanned Documents Bindings to XDS, XDR, XDM	204
	20.5 Scanned Documents Content Process Flow.....	205
	21 Sharing Value Sets Integration Profile (SVS).....	206
	21.1 Actors/ Transactions.....	206
	21.1.1 Assumptions and background information	207

295	21.1.2 Value Set Unique ID and Value Set Version	208
	21.1.3 The relationship between ITI SVS and CTS	209
	21.2 SVS Integration Profile Options	215
	21.2.1 Retrieve Multiple Value Sets	216
	21.3 SVS Process Flow	216
300	21.3.1 Overview of the entire process flow	217
	21.3.2 Use Cases	218
	21.4 SVS Security Considerations	226
	22 Document-based Referral Request Integration Profile (DRR)	228
	23 Patient Identifier Cross-referencing HL7 V3 (PIXV3)	229
305	23.1 Actors/Transactions.....	229
	23.2 Patient Identifier Cross-referencing HL7 V3 Integration Profile Options.....	231
	23.2.1 Pediatric Demographics.....	231
	23.3 Patient Identifier Cross-referencing HL7 V3 Integration Profile Process Flows	232
	23.4 Relationship between the PIXV3 Integration Profile and eMPI.....	232
310	23.5 Patient Identifier Communication Requirement	232
	23.6 Security Considerations	233
	24 Patient Demographics Query HL7 V3 (PDQV3).....	234
	24.1 Actors/Transactions.....	234
	24.2 Patient Demographics Query HL7 V3 Integration Profile Options	234
315	24.2.1 Continuation	235
	24.2.2 Pediatric Demographics Option	235
	24.3 Patient Demographics Query HL7 V3 Process Flow.....	236
	24.3.1 Combined Use of PDQV3 with other IHE Workflow Profiles	236
	24.3.2 Supplier Data Configuration	236
320	24.4 Intentionally Left Blank	237
	24.5 Security Considerations	237
	25 Multi-Patient Queries (MPQ).....	238
	25.1 Actors/ Transactions.....	238
	25.2 Multi-Patient Query Integration Profile Options	239
325	25.2.2 Asynchronous Web Services Exchange Option	239
	25.3 MPQ Process Flow.....	240
	25.4 Use Cases	240
	25.4.1 Multi-Patient Query used in Public Health	240
	25.4.2 Technical Use Cases.....	242
330	25.5 Security Considerations	242
	26 Reserved.....	244
	27 XCPD Integration Profile	244
	27.1 Actors/ Transactions.....	244
	27.1.1 Actors	245
335	27.1.2 Transactions	245
	27.2 XCPD Integration Profile Options	246
	27.2.1 Asynchronous Web Services Exchange Option.....	246
	27.2.2 Deferred Response Option	246
	27.3 XCPD Process Flow.....	247

340	27.3.1 Use Cases	247
	27.3.2 Detailed Interactions (Informative)	248
	27.4 XCPD Security Considerations.....	252
	27.4.1 Requirements/Recommendations.....	252
	27.4.2 Policy Choices.....	253
345	Appendix A: Actor Descriptions	254
	Appendix B: Transaction Descriptions.....	257
	Appendix C: IHE Integration Statements	262
	C.1 Structure and Content of an IHE Integration Statement.....	262
	C.2 Format of an IHE Integration Statement	263
350	Appendix D: User Authentication Techniques - Passwords, Biometrics, and Tokens.....	264
	Appendix E: Cross Profile Considerations	265
	E.1 Combined use of RID, EUA and PIX Integration Profiles.....	265
	E.2 XDS Integration with RID.....	265
	E.3 XDS Integration with PIX	266
355	E.4 XDS Integration with PWP	267
	E.5 XDS Integration with PDQ.....	268
	E.6 XDM Integration with XDS, Content Integration Profiles, PIX, and DSG	268
	E.7 XDM/XDR Distinction.....	269
	E.8 XDR Integration with XDS, Content Integration Profiles, PIX, and DSG	269
360	E.8.1 XDR Integration with XDS	269
	E.8.2 XDR Integration with XDS Content Integration Profiles.....	270
	E.8.3 XDR Integration with PIX.....	270
	E.9 XCA Integration with XDS and non-XDS communities	271
	E.9.1 An XDS Affinity Domain.....	271
365	E.9.2 A Non-XDS Affinity Domain.....	271
	E.9.3 A Collection of XDS Affinity Domains (informative).....	272
	E.9.4 A Collection of Non-XDS Affinity Domains (informative).....	273
	E.9.5 An XDS Affinity Domain with a “Transparent” XCA Gateway (informative)	274
	E.10 XCA and Patient Identification Management.....	274
370	E.10.1 Patient Identification using PIX.....	275
	E.10.2 Patient Identification using PDQ	275
	Appendix F: Request to Standards Development Organizations.....	277
	Appendix G: Security Considerations	278
	G.1 Cross Profile Considerations	278
375	Appendix H: Intentionally Left Blank	280
	Appendix I: Intentionally Left Blank.....	281
	Appendix J: Content and Format of XDS Documents.....	282
	Appendix K: XDS Concept Details	284
	K.1 XDS Document Concept.....	284
380	K.2 Concept of an XDS Affinity Domain.....	285
	K.3 Other Principles of XDS	286
	K.4 Document Identification.....	287
	K.5 Example of Document Relationship	287
	Appendix L: XDS Affinity Domain Definition Checklist.....	289

385	Appendix M: Cross-Enterprise Document Sharing and IHE Roadmap	290
	M.1 Document Content Integration Profiles for XDS	290
	M.2 Cross-Enterprise Dynamic Information Sharing	290
	M.3 Collaborative Workflow Process Management	290
	M.4 Security and Privacy Management	290
390	M.5 Federation of XDS Affinity Domains	291
	Appendix N: Intentionally Left Blank	292
	Appendix O: Intentionally Left Blank	293
	Appendix P: Privacy Access Policies (Informative).....	294
	P.1 Consents in a sensitivity labeled and role based access control environment	295
395	P.3 Possible checklist for implementations.....	296
	P.4 Potential obligations	297
	P.5 Dynamic Use Models.....	298
	GLOSSARY	299

400 **1 Introduction**

Integrating the Healthcare Enterprise (IHE) is an initiative designed to stimulate the integration of the information systems that support modern healthcare institutions. Its fundamental objective is to ensure that in the care of patients all required information for medical decisions is both correct and available to healthcare professionals. The IHE initiative is both a process and a forum for encouraging integration efforts. It defines a technical framework for the implementation of established messaging standards to achieve specific clinical goals. It includes a rigorous testing process for the implementation of this framework. And it organizes educational sessions and exhibits at major meetings of medical professionals to demonstrate the benefits of this framework and encourage its adoption by industry and users.

410 The approach employed in the IHE initiative is to support the use of existing standards, e.g., HL7, ASTM, DICOM, ISO, IETF, OASIS and others as appropriate, rather than to define new standards. IHE profiles further constrain configuration choices where necessary in these standards to ensure that they can be used in their respective domains in an integrated manner between different actors. When clarifications or extensions to existing standards are necessary, 415 IHE refers recommendations to the relevant standards bodies.

This initiative has numerous sponsors and supporting organizations in different medical specialty domains and geographical regions. In North America the primary sponsors are the Healthcare Information and Management Systems Society (HIMSS) and the Radiological Society of North America (RSNA). IHE Canada has also been formed. IHE Europe (IHE-EUR) is supported by a large coalition of organizations including the European Association of Radiology (EAR) and European Congress of Radiologists (ECR), the Coordination Committee of the Radiological and Electromedical Industries (COCIR), Deutsche Röntgengesellschaft (DRG), the EuroPACS Association, Groupement pour la Modernisation du Système d'Information Hospitalier (GMSIH), Société Française de Radiologie (SFR), Società Italiana di Radiologia Medica (SIRM), and the European Institute for health Records (EuroRec). In Japan IHE-J is sponsored by the Ministry of Economy, Trade, and Industry (METI); the Ministry of Health, Labor, and Welfare; and MEDIS-DC; cooperating organizations include the Japan Industries Association of Radiological Systems (JIRA), the Japan Association of Healthcare Information Systems Industry (JAHIS), Japan Radiological Society (JRS), Japan Society of Radiological Technology (JSRT), 430 and the Japan Association of Medical Informatics (JAMI). Other organizations representing healthcare professionals are invited to join in the expansion of the IHE process across disciplinary and geographic boundaries.

1.1 Overview of the Technical Framework

This document, the IHE IT Infrastructure Technical Framework (ITI TF), defines specific implementations of established standards to achieve integration goals that promote appropriate sharing of medical information to support optimal patient care. It is expanded annually, after a period of public review, and maintained regularly through the identification and correction of errata. The current version, Revision 11.0 for Final Text, specifies the IHE transactions defined and implemented as of September 2014. The latest version of the document is always available 440 via the Internet at http://www.ihe.net/Resources/Technical_Frameworks.

445 The IHE IT Infrastructure Technical Framework identifies a subset of the functional components of the healthcare enterprise, called IHE actors, and specifies their interactions in terms of a set of coordinated, standards-based transactions. It describes this body of transactions in progressively greater depth. The present volume (ITI TF-1) provides a high-level view of IHE functionality, showing the transactions organized into functional units called integration profiles that highlight their capacity to address specific IT Infrastructure requirements.

450 Volumes 2a, 2b, and 2x of the IT Infrastructure Technical Framework provide detailed technical descriptions of each IHE transaction used in the IT Infrastructure Integration Profiles. Volume 3 contains content specification and specifications used by multiple transactions. These volumes are consistent and can be used in conjunction with the Integration Profiles of other IHE domains.

The other domains within the IHE initiative also produce Technical Frameworks within their respective areas that together form the IHE Technical Framework. For example, the following IHE Technical Framework(s) are some of those which are available:

- IHE IT Infrastructure Technical Framework
- 455 • IHE Cardiology Technical Framework
- IHE Laboratory Technical Framework
- IHE Patient Care Coordination Technical Framework
- IHE Radiology Technical Framework

460 Where applicable, references are made to other technical frameworks. For the conventions on referencing other frameworks, see ITI TF-1: 1.6.3.

1.2 Overview of the IT Infrastructure Volume 1

The remainder of Section 1 further describes the general nature, purpose and function of the Technical Framework. Section 2 introduces the concept of IHE Integration Profiles that make up the Technical Framework.

465 Section 3 and the subsequent sections of this volume provide detailed documentation on each integration profile, including the IT Infrastructure problem it is intended to address and the IHE actors and transactions it comprises.

470 The appendices following the main body of the document provide a summary list of the actors and transactions, detailed discussion of specific issues related to the integration profiles and a glossary of terms and acronyms used.

1.3 Audience

The intended audience of this document is:

- IT departments of healthcare institutions
- Technical staff of vendors participating in the IHE initiative
- 475 • Experts involved in standards development

- Those interested in integrating healthcare information systems and workflows

1.4 Relationship to Standards

480 The IHE Technical Framework identifies functional components of a distributed healthcare environment (referred to as IHE actors), solely from the point of view of their interactions in the healthcare enterprise. At its current level of development, it defines a coordinated set of transactions based on ASTM, DICOM, HL7, IETF, ISO, OASIS and W3C standards. As the scope of the IHE initiative expands, transactions based on other standards may be included as required.

485 In some cases, IHE recommends selection of specific options supported by these standards; however, IHE does not introduce technical choices that contradict conformance to these standards. If errors in or extensions to existing standards are identified, IHE's policy is to report them to the appropriate standards bodies for resolution within their conformance and standards evolution strategy.

490 IHE is therefore an implementation framework, not a standard. Conformance claims for products must still be made in direct reference to specific standards. In addition, vendors who have implemented IHE integration capabilities in their products may publish IHE Integration Statements to communicate their products' capabilities. Vendors publishing IHE Integration Statements accept full responsibility for their content. By comparing the IHE Integration Statements from different products, a user familiar with the IHE concepts of actors and integration profiles can determine the level of integration between them. See ITI TF-1: Appendix C for the format of IHE Integration Statements.

1.5 Relationship to Real-world Architectures

500 The IHE actors and transactions described in the IHE Technical Framework are abstractions of the real-world healthcare information system environment. While some of the transactions are traditionally performed by specific product categories (e.g., HIS, Clinical Data Repository, Radiology Information Systems, Clinical Information Systems or Cardiology Information Systems), the IHE Technical Framework intentionally avoids associating functions or actors with such product categories. For each actor, the IHE Technical Framework defines only those functions associated with integrating information systems. The IHE definition of an actor should therefore not be taken as the complete definition of any product that might implement it, nor should the framework itself be taken to comprehensively describe the architecture of a healthcare information system.

510 The reason for defining actors and transactions is to provide a basis for defining the interactions among functional components of the healthcare information system environment. In situations where a single physical product implements multiple functions, only the interfaces between the product and external functions in the environment are considered to be significant by the IHE initiative. Therefore, the IHE initiative takes no position as to the relative merits of an integrated environment based on a single, all-encompassing information system versus one based on multiple systems that together achieve the same end. IHE demonstrations emphasize the integration of multiple vendors' systems based on the IHE Technical Framework.

1.6 Conventions

This document has adopted the following conventions for representing the framework concepts and specifying how the standards upon which the IHE Technical Framework is based should be applied.

520 1.6.1 IHE Actor and Transaction Diagrams and Tables

Each integration profile is a representation of a real-world capability that is supported by a set of actors that interact through transactions. Actors are information systems or components of information systems that produce, manage, or act on categories of information required by operational activities in the enterprise. Transactions are interactions between actors that
525 communicate the required information through standards-based messages.

The diagrams and tables of actors and transactions in subsequent sections indicate which transactions each actor in a given profile must support.

The transactions shown on the diagrams are identified both by their name and the transaction number as defined in ITI TF-2a and 2b. The transaction numbers are shown on the diagrams as
530 bracketed numbers prefixed with the specific Technical Framework domain.

In some cases, a profile is dependent on a prerequisite profile in order to function properly and be useful. For example, Enterprise User Authentication depends on Consistent Time. These dependencies can be found by locating the desired profile in Table 2-1 to determine which profile(s) are listed as prerequisites. An actor must implement all required transactions in the
535 prerequisite profiles in addition to those in the desired profile.

1.6.2 Process Flow Diagrams

The descriptions of integration profiles that follow include process flow diagrams that illustrate how the profile functions as a sequence of transactions between relevant actors.

These diagrams are intended to provide an overview so the transactions can be seen in the
540 context of an institution's workflow. Certain transactions and activities not defined in detail by IHE are shown in these diagrams in *italics* to provide additional context on where the relevant IHE transactions fit into the broader scheme of healthcare information systems.

These diagrams are not intended to present the only possible scenario. Often other actor groupings are possible, and transactions from other profiles may be interspersed.

545 In some cases the sequence of transactions may be flexible. Where this is the case there will generally be a note pointing out the possibility of variations. Transactions are shown as arrows oriented according to the flow of the primary information handled by the transaction and not necessarily the initiator.

1.6.3 Technical Framework Cross-references

550 When references are made to another section within a Technical Framework volume, a section number is used by itself. When references are made to other volumes or to a Technical Framework in another domain, the following format is used:

<domain designator> TF-<volume number>: <section number>, where

555 <domain designator> is a short designator for the IHE domain (ITI = IT Infrastructure, RAD = Radiology)

<volume number> is the applicable volume within the given Technical Framework (e.g., 1, 2a, 2b, 2x, 3), and

<section number> is the applicable section number.

560 For example: ITI TF-1: 3.1 refers to Section 3.1 in volume 1 of the IHE IT Infrastructure Technical Framework. RAD TF-3: 4.33 refers to Section 4.33 in volume 3 of the IHE Radiology Technical Framework. ITI TF-2x: Appendix B refers to Appendix B in volume 2x of the IHE IT Infrastructure Technical Framework.

When references are made to Transaction numbers in the Technical Framework, the following format is used:

565 [<domain designator>-<transaction number>], where

<transaction number> is the transaction number within the specified domain.

For example: [ITI-1] refers to Transaction 1 from the IHE IT Infrastructure Technical Framework.

1.7 Scope of Changes Introduced in the Current Year

570 The IHE Technical Framework is updated annually to reflect new profiles, corrections and new transactions used in those profiles.

This Version of the IT Infrastructure Technical Framework finalizes the following Integration Profiles.

- 575 1. **Retrieve Information for Display (RID)** – a simple and rapid read-only access to patient information necessary for provision of better care. It supports access to existing persistent documents in well-known presentation formats such as CDA, PDF, JPEG, etc. It also supports access to specific key patient-centric information such as allergies, current medications, summary of reports, etc. for presentation to a clinician.
- 580 2. **Enterprise User Authentication (EUA)** – a means to establish one name per user that can then be used on all of the devices and software that participate in this integration profile, greatly facilitating centralized user authentication management and providing users with the convenience and speed of a single sign-on. This profile leverages Kerberos (RFC 1510) and the HL7 CCOW standard (user subject).
- 585 3. **Patient Identifier Cross-referencing (PIX)** – provides cross-referencing of patient identifiers from multiple Patient Identifier Domains. These patient identifiers can then be used by identity consumer systems to correlate information about a single patient from sources that know the patient by different identifiers. Includes the Pediatric Demographics Option.
- 590 4. **Patient Synchronized Applications (PSA)** – a means for viewing data for a single patient using independent and unlinked applications on a user's workstation, reducing the

repetitive tasks of selecting the same patient in multiple applications. Data can be viewed from different Identifier Domains when used with the Patient Identifier Cross-referencing Integration Profile to resolve multiple identifications for the same patient. This profile leverages the HL7 COW standard specifically for patient subject context management.

- 595
5. **Consistent Time (CT)** – mechanisms to synchronize the time base between multiple actors and computers. Various infrastructure, security, and acquisition profiles require use of a consistent time base on multiple computers. The Consistent Time Profile provides a median synchronization error of less than 1 second.
- 600
6. **Patient Demographics Query (PDQ)** – provides ways for multiple distributed applications to query a patient information server for a list of patients, based on user-defined search criteria, and retrieve a patient’s demographic (and, optionally, visit or visit-related) information directly into the application. Includes the Pediatric Demographics Option.
- 605
7. **Audit Trail and Node Authentication (ATNA)** – establishes the characteristics of a Basic Secure Node:
- It describes the security environment (user identification, authentication, authorization, access control, etc.) assumed for the node so that security reviewers may decide whether this matches their environments.
 - It defines basic auditing requirements for the node
 - 610 • It defines basic security requirements for the communications of the node using TLS or equivalent functionality.
 - It establishes the characteristics of the communication of audit messages between the Basic Secure Nodes and Audit Repository nodes that collect audit information.
 - 615 • This profile has been designed so that specific domain frameworks may extend it through an option defined in the domain specific technical framework. Extensions are used to define additional audit event reporting requirements, especially actor specific requirements. The Radiology Audit Trail Option in the IHE Radiology Technical Framework is an example of such an extension.
- 620
8. **Personnel White Pages (PWP)** – provides access to basic human workforce user directory information. This information has broad use among many clinical and non-clinical applications across the healthcare enterprise.
- 625
9. **Cross-Enterprise Document Sharing (XDS)** – enables a number of healthcare delivery organizations belonging to an XDS Affinity Domain (e.g., a community of care) to cooperate in the care of a patient by sharing clinical records in the form of documents as they proceed with their patients’ care delivery activities. This profile is based upon ebXML Registry standards and SOAP. It describes the configuration of an ebXML Registry in sufficient detail to support Cross Enterprise Document Sharing.
- 630
10. **Cross-Enterprise User Assertion Profile (XUA)** - provides a means to communicate claims about the identity of an authenticated principal (user, application, system...) in transactions that cross enterprise boundaries. To provide accountability in these cross-

- 635 enterprise transactions there is a need to identify the requesting principal in a way that enables the receiver to make access decisions and generate the proper audit entries. The XUA Profile supports enterprises that have chosen to have their own user directory with their own unique method of authenticating the users, as well as others that may have chosen to use a third party to perform the authentication.
- 640 11. **Patient Administration Management (PAM)** - provides patient identity, registration, and encounter management transactions in a healthcare enterprise as well as across enterprises.
- 645 12. **Cross-Enterprise Document Media Interchange (XDM)** - provides document interchange using a common file and directory structure over several standard media. This permits the patient to use physical media to carry medical documents. This also permits the use of person-to-person email to convey medical documents.
- 650 13. **Basic Patient Privacy Consents (BPPC)** - provides a mechanism to record the patient privacy consent(s), and a method for Content Consumers to use to enforce the privacy consent appropriate to the use. This profile complements XDS by describing a mechanism whereby an XDS Affinity Domain can develop and implement multiple privacy policies, and describes how that mechanism can be integrated with the access control mechanisms supported by the XDS Actors (e.g., EHR systems).
- 655 14. **Cross Enterprise Sharing of Scanned Documents (XDS-SD)** – A profile which associates structured, healthcare metadata with non-healthcare specific document format to maintain the integrity of the patient health record as managed by the source system
- 660 15. **Cross-Enterprise Document Reliable Interchange (XDR)** - provides document interchange using a reliable messaging system. This permits direct document interchange between EHRs, PHRs, and other healthcare IT systems in the absence of a document sharing infrastructure such as XDS.
- 665 16. **Multi-Patient Queries (MPQ)** - defines a mechanism to enable aggregated queries to a Document Registry based on certain criteria needed by areas related to data analysis, such as quality accreditation of health care practitioners or health care facilities, clinical research trial data collection or population health monitoring..
- 670 17. **Patient Identifier Cross-referencing HL7 V3 (PIXV3)** – provides cross-referencing of patient identifiers from multiple Patient Identifier Domains. These patient identifiers can then be used by identity consumer systems to correlate information about a single patient from sources that know the patient by different identifiers. This profile uses HL7 V3 as the message format, and SOAP-based web services for transport.
- 675 18. **Patient Demographics Query HL7 V3 (PDQV3)** - provides ways for multiple distributed applications to query a patient information server for a list of patients, based on user-defined search criteria, and retrieve a patient’s demographic information directly into the application. This profile uses HL7 V3 as the message format, and SOAP-based web services for transport.
- 680 19. **Cross-Community Access (XCA)** - supports the means to query and retrieve patient relevant medical data held by other communities. A community is defined as a coupling

of facilities/enterprises that have agreed to work together using a common set of policies for the purpose of sharing clinical information via an established mechanism.

- 675
20. **Retrieve Form for Data Capture (RFD)** - provides a means for the retrieval and submission of forms data between physicians/investigators and electronic data capture systems or other data collection agencies.
21. **Sharing Value Sets (SVS)** - provides two mechanisms for retrieving Value Sets from a Value Set Repository.
- 680
22. **Cross-Community Patient Discovery (XCPD)** - supports the means to locate communities that hold patient relevant health data and the translation of patient identifiers across communities holding the same patient's data.
23. **On Demand Documents (ODD)** - Updates the XCA and XDS Profiles to add support for sharing of document content which is assembled at time of retrieval through use of On-Demand Documents.

685 **1.8 Security Implications**

IHE transactions often contain information that must be protected in conformance with privacy laws and regulations, such as HIPAA or similar requirements in other regions. IHE includes a few security and privacy-focused profiles listed below. Other IHE Profiles generally do not have specific privacy protections, but rather expect a proper grouping with one or more of the security profiles:

690

- The Audit Trail and Node Authentication (ATNA) Profile specifies a means to ensure that nodes in a network are authenticated.
- The ATNA Profile specifies an audit message for reporting security- and privacy-relevant events.
- 695 • The Enterprise User Authentication (EUA) Profile specifies a means to authenticate system users and to share knowledge of the authenticated users among applications.
- The Personnel White Pages (PWP) Profile provides a repository that may be used to hold system users' identification data.

Implementers may follow these IHE profiles to fulfill some of their security needs. It is understood that institutions must implement policy and workflow steps to satisfy enterprise needs and to comply with regulatory requirements.

700

1.9 Comments

IHE International welcomes comments on this document and the IHE initiative. They can be submitted using the Web-based comment form at http://ihe.net/ITI_Public_Comments or by sending an email to the co-chairs and secretary of the IT Infrastructure domain committees at iti@ihe.net.

705

1.10 Copyright Permission

710 Health Level Seven, Inc., has granted permission to the IHE to reproduce tables from the HL7 standard. The HL7 tables in this document are copyrighted by Health Level Seven, Inc. All rights reserved. Material drawn from these documents is credited where used.

1.11 IHE Technical Framework Development and Maintenance Process

715 The IHE IT Infrastructure Technical Framework is continuously maintained and expanded on an annual basis by the IHE IT Infrastructure Technical Committee. The development and maintenance process of the Framework follows a number of principles to ensure stability of the specification so that both vendors and users may use it reliably in specifying, developing and acquiring systems with IHE integration capabilities.

720 The first of these principles is that any extensions, clarifications and corrections to the Technical Framework must maintain backward compatibility with previous versions of the framework in order to maintain interoperability with systems that have implemented IHE Actors and Integration Profiles defined there.

The IHE IT Infrastructure Technical Framework is developed and re-published annually following a three-step process:

- 725 1. The IT Infrastructure Technical Committee develops supplements to the current stable version of the Technical Framework to support new functionality identified by the IHE Strategic and Planning Committees and issues them for public comment.
- 730 2. The Committee addresses all comments received during the public comment period and publishes an updated version of the Technical Framework for “Trial Implementation.” This version contains both the stable body of the Technical Framework from the preceding cycle and the newly developed supplements. It is the version of the Technical Framework used by vendors in developing trial implementation software for the annual IT Infrastructure Connectathon.
- 735 3. The Committee regularly considers change proposals to the Trial Implementation version of the Technical Framework, including those from implementers who participate in the Connectathon. After resolution of all change proposals received within 60 days of the Connectathon, the Technical Framework version is published as “Final Text”.

2 IT Infrastructure Integration Profiles

740 IHE IT Infrastructure Integration Profiles (Figure 2-1), offer a common language that healthcare professionals and vendors can use to discuss integration needs of healthcare enterprises and the integration capabilities of information systems in precise terms. Integration Profiles specify implementations of standards that are designed to meet identified clinical needs. They enable users and vendors to state which IHE capabilities they require or provide, by reference to the detailed specifications of the IHE IT Infrastructure Technical Framework.

745 Integration profiles are defined in terms of IHE Actors and transactions. Actors (see ITI TF-1: Appendix A) are information systems or components of information systems that produce, manage, or act on information associated with clinical and operational activities in the enterprise. Transactions (see ITI TF-1: Appendix B) are interactions between actors that communicate the required information through standards-based messages.

750 Vendor products support an Integration Profile by implementing the appropriate actor(s) and transactions. A given product may implement more than one actor and more than one integration profile.

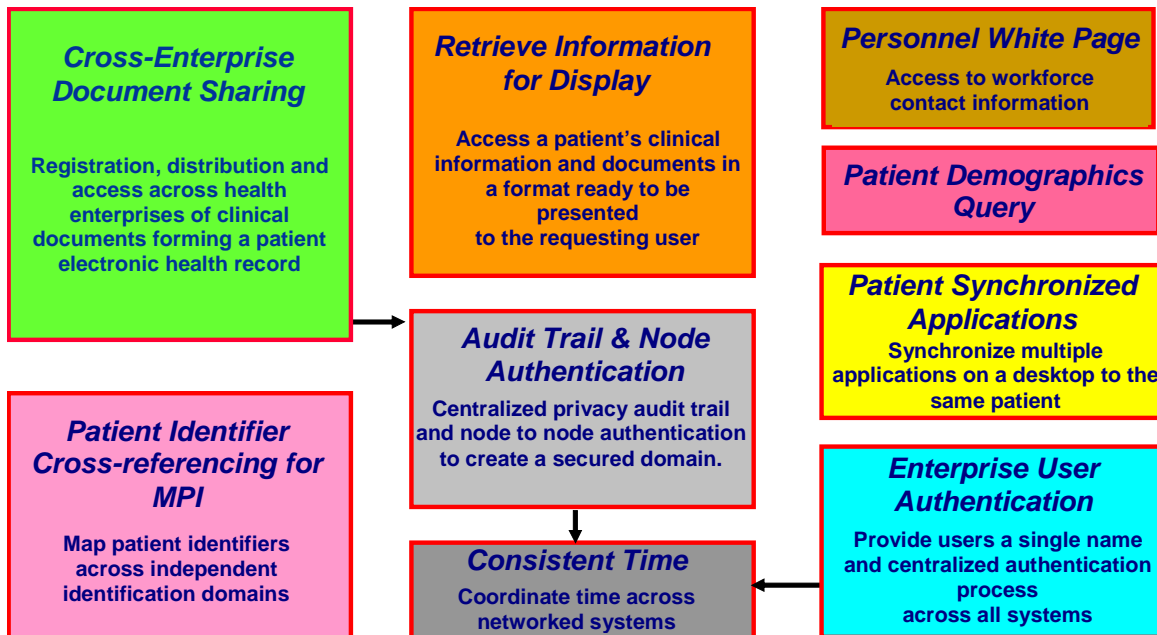


Figure 2-1: IHE IT Infrastructure Integration Profiles

755 2.1 Dependencies among Integration Profiles

Dependencies among IHE Integration Profiles exist when implementation of one integration profile is a prerequisite for achieving the functionality defined in another integration profile. Figure 2-1 provides a graphical view of the dependencies among IHE IT Infrastructure Integration Profiles. The arrows in the figure point from a given integration profile to the

760 integration profile(s) upon which it depends. Table 2-1 defines these dependencies in tabular form.

765 Some dependencies require that an actor supporting one profile be grouped with one or more actors supporting other integration profiles. For example, Enterprise User Authentication (EUA) requires that different participating actors be grouped with the Time Client Actor that participates in the Consistent Time (CT) Integration Profile. The dependency exists because EUA actors must refer to consistent time in order to function properly.

Table 2-1: Integration Profiles Dependencies

Integration Profile	Depends on	Dependency Type	Purpose
Retrieve Information for Display Integration (RID)	<i>None</i>	None	-
Enterprise User Authentication (EUA)	Consistent Time	Each actor implementing EUA shall be grouped with the Time Client Actor	Required to manage expirations of authentication tickets
Patient Identifier Cross-referencing (PIX)	Consistent Time	Each actor implementing PIX shall be grouped with the Time Client Actor	Required to manage and resolve conflicts in multiple updates.
Patient Synchronized Applications (PSA)	<i>None</i>	<i>None</i>	-
Consistent Time (CT)	<i>None</i>	<i>None</i>	-
Patient Demographics Query (PDQ)	<i>None</i>	<i>None</i>	-
Personnel White Pages (PWP)	<i>None</i>	<i>None</i>	-
Audit Trail and Node Authentication (ATNA)	Consistent Time	An ATNA Secure Node Actor shall be grouped with the Time Client Actor	Required for consistent time in audit logs.
Cross-Enterprise Document Sharing (XDS)	Audit Trail and Node Authentication	Each XDS Actor shall be grouped with the ATNA Secure Node or Secure Application Actor.	Required to manage audit trail of exported PHI, node authentication and transport encryption.
Cross-Enterprise Document Sharing (XDS)	Consistent Time	Each XDS Actor shall be grouped with the Time Client Actor	To ensure consistency among document and submission set dates.
Cross-Enterprise User Assertion (XUA)	None	None	
Patient Administration Management (PAM)	None	None	-
Cross-Enterprise Document Media Interchange (XDM)	Audit Trail and Node Authentication	Each XDM Actor shall be grouped with Secure Node or Secure Application Actor	Requires audit trails.
Cross-Enterprise Document Media Interchange (XDM)	Any IHE Content Profile	The Portable Media Importer shall be grouped with a Content Consumer of one or more IHE Content Profile	Enables some form of processing of imported medical data.

Integration Profile	Depends on	Dependency Type	Purpose
Basic Patient Privacy Consent (BPPC)	XDS, XDM or XDR	The BPPC Content Creator shall be grouped with an XDS or XDR Document Source Actor, or an XDM Portable Media Creator. The BPPC Content Consumer shall be grouped with an XDS Document Consumer, or an XDR Document Recipient, or an XDM Portable Media Importer.	The content of a Basic Patient Privacy Consent Acknowledgement document is intended for use in XDS, XDR and XDM.
Basic Patient Privacy Consent (BPPC)	Cross Enterprise Sharing of Scanned Documents	The BPPC Content Consumer shall be grouped with the XDS-SD Content Consumer.	Enables capturing of wet signatures on patients' consent documents.
Cross Enterprise Sharing of Scanned Documents (XDS-SD)	XDS, XDM or XDR	The XDS-SD Content Creator shall be grouped with an XDS or XDR Document Source Actor, or an XDM Portable Media Creator. The XDS-SD Content Consumer shall be grouped with an XDS Document Consumer, or an XDR Document Recipient, or an XDM Portable Media Importer.	The content of this profile is intended for use in XDS, XDR and XDM.
Cross-Enterprise Document Reliable Interchange (XDR)	ATNA	Each XDR Actor shall be grouped with Secure Node or Secure Application Actor	Requires secure communication and audit trails.
Multi-Patient Queries (MPQ)	Audit Trail and Node Authentication	Each Document Registry Actor and each Document Consumer shall be grouped with a Secure Node or a Secure Application Actor	Required to manage audit trail of exported PHI, node authentication and transport encryption
Multi-Patient Queries (MPQ)	Consistent Time	Each Document Registry Actor and each Document Consumer shall be grouped with the Time Client Actor.	To ensure consistency among document and submission set dates
Patient Identifier Cross-Referencing HL7 V3 (PIX v3)	Consistent Time	Each actor implementing PIXv3 shall be grouped with the Time Client Actor	Required to manage and resolve conflicts in multiple updates
Patient Demographics Query HL7 V3 (PDQv3)	None	None	
Cross-Community Access (XCA)	Audit Trail and Node Authentication	Each XCA Actor shall be grouped with Secure Node Actor or Secure Application	Required to manage audit trail of exported PHI, node authentication and transport encryption.
Cross-Community Access (XCA)	Consistent Time	Each XCA Actor shall be grouped with the Time Client Actor.	To ensure consistency among document and submission set dates.

Integration Profile	Depends on	Dependency Type	Purpose
Retrieve Form for Data Capture (RFD)	None	None	-
Sharing Value Sets	Audit Trail and Node Authentication	The Value Set Repository shall be grouped with a Secure Node Actor/Secure Application	Required to manage audit trail of Value Sets sharing and node authentication.
XCPD	Audit Trail and Node Authentication	Each XCPD Actor shall be grouped with Secure Node or Secure Application Actor	- Required to manage audit trail of exported PHI, node authentication and transport security.
XCPD	Consistent Time	Each XCPD Actor shall be grouped with the Time Client Actor	- Required due to ATNA grouping.

770 To support a dependent profile, an actor must implement all required transactions in the prerequisite profiles in addition to those in the dependent profile. In some cases, the prerequisite is that the actor selects any one of a given set of profiles.

2.2 Integration Profiles Overview

In this document, each IHE Integration Profile is defined by:

- 775
- The IHE actors involved
 - The specific set of IHE transactions exchanged by each IHE actor

780 These requirements are presented in the form of a table of transactions required for each actor supporting the Integration Profile. Actors supporting multiple Integration Profiles are required to support all the required transactions of each Integration Profile supported. When an Integration Profile depends upon another Integration Profile, the transactions required for the dependent Integration Profile have not been included in the table.

785 Note that IHE Integration Profiles are not statements of conformance to standards, and IHE is not a certifying body. Users should continue to request that vendors provide statements of their conformance to standards issued by relevant standards bodies, such as HL7 and DICOM. Standards conformance is a prerequisite for vendors adopting IHE Integration Profiles.

790 Also note that there are critical requirements for any successful integration project that IHE cannot address. Successfully integrating systems still requires a project plan that minimizes disruptions and describes fail-safe strategies, specific and mutually understood performance expectations, well-defined user interface requirements, clearly identified systems limitations, detailed cost objectives, plans for maintenance and support, etc.

2.2.1 This section is reserved.

2.2.2 Sharing Value Set Integration Profile (SVS)

Sharing Value Sets (SVS) provides a means through which healthcare systems producing or consuming clinical or administrative data, such as diagnostic imaging equipment, laboratory

795 reporting systems, primary care physician office EMR systems, or national healthcare record systems, can receive a common, uniform nomenclature managed centrally. Shared nomenclatures with specific derived value sets are essential to achieving semantic interoperability.

800 This profile describes transactions for retrieving Value Sets from a Value Set Repository by a Value Set Consumer. A single Value Set Repository can be accessed by many Value Set Consumers, establishing a domain of consistent and uniform set of nomenclatures. It supports automated loading of Value Sets by systems implementing a Value Set Consumer, reducing the burden of manual configuration.

2.2.3 Retrieve Information for Display (RID)

805 *Retrieve Information for Display* enables simple and rapid access to patient information for better care. It supports access to existing persistent documents in well-known presentation formats such as CDA, PDF, JPEG, etc. It also supports access to specific key patient-centric information such as allergies, current medications, summary of reports, etc. for presentation to a clinician. It complements workflows from within the users' on-screen workspace or application.

810 By linking it with two other IHE profiles - Enterprise User Authentication and Patient Identifier Cross-referencing, this profile's reach can extend across organization boundaries within an enterprise. This IHE Integration Profile leverages HTTP, Web Services, IT presentation formats and HL7 CDA Level 1.

2.2.4 Enterprise User Authentication (EUA)

815 *Enterprise User Authentication* defines a means to establish one name per user that can then be used on all of the devices and software that participate in this integration profile. It greatly facilitates centralized user authentication management and provides users with the convenience and speed of a single sign-on. This profile leverages Kerberos (RFC 1510) and the HL7 CCOW standard (user subject). User authentication is a necessary step for most application and data

820 access operations and streamlines workflow for users. Future profiles will deal with other security issues, such as authorization management.

2.2.5 Patient Identifier Cross-referencing (PIX)

825 The *PIX Profile* supports the cross-referencing of patient identifiers from multiple Patient Identifier Domains. These cross-referenced patient identifiers can then be used by "identity consumer" systems to correlate information about a single patient from sources that "know" the patient by different identifiers. This allows a clinician to have more complete view of the patient information.

2.2.6 Patient Synchronized Applications (PSA)

830 *Patient Synchronized Applications* supports viewing data for a single patient among otherwise independent and unlinked applications on a user's workstation. Its implementation reduces the repetitive tasks of selecting the same patient in multiple applications. It also improves patient safety by reducing the chance of medical errors caused by viewing the wrong patient's data. Its ability to work with the Patient Identifier Cross-referencing provides a seamless environment for

835 clinicians and IT staff. This profile leverages the HL7 CCOW standard specifically for patient
subject context management.

2.2.7 Consistent Time (CT)

840 *Consistent Time* defines mechanisms to synchronize the time base between multiple actors and
computers. Various infrastructure, security, and acquisition profiles require use of a consistent
time base on multiple computers. The Consistent Time Profile provides median synchronization
error of less than 1 second. Configuration options can provide better synchronization. The
Consistent Time Profile specifies the use of the Network Time Protocol (NTP) defined in RFC
1305.

2.2.8 Patient Demographics Query (PDQ)

845 *Patient Demographics Query* provides ways for multiple distributed applications to query a
patient information server for a list of patients, based on user-defined search criteria, and retrieve
a patient's demographic (and, optionally, visit or visit-related) information directly into the
application.

2.2.9 Audit Trail and Node Authentication (ATNA)

Audit Trail and Node Authentication establishes the characteristics of a Basic Secure Node:

- 850
1. It describes the security environment (user identification, authentication, authorization,
access control, etc.) assumed for the node so that security reviewers may decide whether
this matches their environments.
 2. It defines basic auditing requirements for the node
 - 855 3. It defines basic security requirements for the communications of the node using TLS or
equivalent functionality.
 4. It establishes the characteristics of the communication of audit messages between the
Basic Secure Nodes and Audit Repository nodes that collect audit information.
 5. It defines a Secure Application Actor for describing product configurations that are not
able to meet all of the requirements of a Secure Node.

860 Note: ATNA security considerations require the use of Secure Nodes. The Secure Application is defined to permit
product configurations to indicate that the product is ready for easy integration into a Secure Node environment
because it performs all of the security related functions that are directly related to the application function. See
ITI TF-1: 9.7 for more details.

865 This profile has been designed so that specific domain frameworks may extend it through an
option defined in the domain specific technical framework. Extensions are used to define
additional audit event reporting requirements, especially actor specific requirements. The
Radiology Audit Trail Option in the IHE Radiology Technical Framework is an example of such
an extension.

2.2.10 Cross-Enterprise Document Sharing (XDS)

870 *Cross-Enterprise Document Sharing* enables a number of healthcare delivery organizations
belonging to an XDS Affinity Domain (e.g., a community of care) to cooperate in the care of a
patient by sharing clinical records in the form of documents as they proceed with their patients'
care delivery activities. Federated document repositories and a document registry create a
875 longitudinal record of information about a patient within a given XDS Affinity Domain. This
profile is based upon ebXML Registry standards and SOAP. It describes the configuration of an
ebXML Registry in sufficient detail to support Cross Enterprise Document Sharing.

2.2.11 Personnel White Pages (PWP)

Personnel White Pages Profile (PWP) provides access to basic human workforce user directory
information. This information has broad use among many clinical and non-clinical applications
880 across the healthcare enterprise. The information can be used to enhance the clinical workflow
(contact information), enhance the user interface (user friendly names and titles), and ensure
identity (digital certificates). This Personnel White Pages directory will be related to the User
Identity provided by the Enterprise User Authentication (EUA) Integration Profile previously
defined by IHE.

885 **2.2.12 This section is intentionally left blank.**

2.2.13 Cross Enterprise User Assertion (XUA)

Cross-Enterprise User Assertion provides a means to communicate claims about the identity of
an authenticated principal (user, application, system...) in transactions that cross-enterprise
boundaries. To provide accountability in these cross enterprise transactions there is a need to
890 identify the requesting principal in a way that enables the receiver to make access decisions and
generate the proper audit entries. The XUA Profile supports enterprises that have chosen to have
their own user directory with their own unique method of authenticating the users, as well as
others that may have chosen to use a third party to perform the authentication.

2.2.14 Patient Administration Management (PAM)

895 *Patient Administration Management* establishes the continuity and integrity of patient data, and
additional information such as related persons (primary caregiver, guarantor, next of kin, etc.). It
coordinates the exchange of patient registration and update information among systems that need
to be able to provide current information regarding a patient's encounter status and location. This
profile supports ambulatory and acute care use cases including patient identity feed, admission
900 and discharge, and transfer and encounter management, as well as explicit and precise error
reporting and application acknowledgment.

The PAM Profile supports two patient encounter management scenarios: either one single central
patient registration system serving the entire institution, or multiple patient registration systems
collaborating as peers serving different clinical settings in an institution.

905 **2.2.15 Cross-Enterprise Document Reliable Interchange (XDR)**

Cross-Enterprise Document Reliable Interchange provides document interchange using a reliable messaging system. This permits direct document interchange between EHRs, PHRs, and other healthcare IT systems in the absence of a document sharing infrastructure such as XDS Registry and Repositories.

910 **2.2.16 Cross-Enterprise Document Media Interchange (XDM)**

Cross-Enterprise Document Media Interchange provides document interchange using a common file and directory structure over several standard media. This permits the patient to use physical media to carry medical documents. This also permits the use of person-to-person email to convey medical documents.

915 **2.2.17 Retrieve Form for Data Capture (RFD)**

Retrieve Form for Data Capture provides a method for gathering data within a user's current application to meet the requirements of an external system. RFD supports the retrieval of forms from a form source, display and completion of a form, and return of instance data from the display application to the source application.

920 **2.2.18 Cross-Community Access (XCA)**

Cross-Community Access supports the means to query and retrieve patient relevant medical data held by other communities. A community is defined as a coupling of facilities/enterprises that have agreed to work together using a common set of policies for the purpose of sharing clinical information via an established mechanism. Facilities/enterprises may host any type of healthcare application such as EHR, PHR, etc. A community is identifiable by a globally unique id called the homeCommunityId. Membership of a facility/enterprise in one community does not preclude it from being a member in another community. Such communities may be XDS Affinity Domains which define document sharing using the XDS Profile or any other communities, no matter what their internal sharing structure.

930 **2.2.19 Basic Patient Privacy Consents (BPPC)**

Basic Patient Privacy Consents provides a mechanism to record the patient privacy consent(s), and a method for Content Consumers to use to enforce the privacy consent appropriate to the use. This profile complements XDS by describing a mechanism whereby an XDS Affinity Domain can develop and implement multiple privacy policies, and describes how that mechanism can be integrated with the access control mechanisms supported by the XDS Actors (e.g., EHR systems).

935 **2.2.20 Scanned Documents Integration Profile (XDS-SD)**

A variety of legacy paper, film, electronic and scanner outputted formats are used to store and exchange clinical documents. These formats are not designed for healthcare documentation, and furthermore, do not have a uniform mechanism to store healthcare metadata associated with the documents, including patient identifiers, demographics, encounter, order or service information. The association of structured, healthcare metadata with this kind of document is important to

945 maintain the integrity of the patient health record as managed by the source system. It is necessary to provide a mechanism that allows such source metadata to be stored with the document.

2.2.21 This section is reserved.

Future.

2.2.22 This section is reserved.

Future.

950 **2.2.23 Patient Identifier Cross-referencing HL7 V3 (PIXV3)**

955 The functionality of this profile is identical to the PIX Profile described in Section 2.2.3. The differences are in the format of the messages, and in the use of SOAP-based web services. These changes make this profile well suited for use within an existing IT infrastructure for cross-enterprise data access and exchange. The PIXV3 Profile supports the cross-referencing of patient identifiers from multiple Patient Identifier Domains. These cross-referenced patient identifiers can then be used by “identity consumer” systems to correlate information about a single patient from sources that “know” the patient by different identifiers. This allows a clinician to have more complete view of the patient information.

2.2.24 Patient Demographics Query HL7 V3 (PDQV3)

960 The functionality of this profile is identical to the PDQ Profile described in Section 2.2.6. The differences are in the format of the messages, and in the use of SOAP-based web services. These changes make this profile well suited for use within an existing IT infrastructure for cross-enterprise data access and exchange. The PDQV3 Profile provides ways for multiple organizations, or multiple distributed applications to query a patient information server for a list
965 of patients, based on user-defined search criteria, and retrieve a patient’s demographic information directly into the application.

2.2.25 Multi-Patient Queries (MPQ)

970 *Multi-Patient Queries* defines a mechanism to enable aggregated queries to a Document Registry based on certain criteria needed by areas related to data analysis, such as quality accreditation of health care practitioners or health care facilities, clinical research trial data collection or population health monitoring.

2.2.26 This section is reserved

Future.

2.2.27 XCPD Integration Profile

975 The Cross-Community Patient Discovery (XCPD) Profile supports the means to locate communities that hold patient relevant health data and the translation of patient identifiers across communities holding the same patient’s data. A community is defined as a group of facilities/enterprises that have agreed to work together using a common set of policies for the

980 purpose of sharing health information within the community via an established mechanism. Facilities/enterprises may host any type of healthcare application such as EHR, PHR, etc. A community is identifiable by a globally unique id called the homeCommunityIdMembership of a facility/enterprise in one community does not preclude it from being a member in another community. Such communities may be XDS Affinity Domains which define document sharing using the XDS Profile or any other communities, no matter what their internal sharing structure.

985 **2.3 Product Implementations**

Developers have a number of options in implementing IHE actors and transactions in product implementations. The decisions cover three classes of optionality:

- For a system, select which actors it will incorporate (multiple actors per system are acceptable).
- 990 • For each actor, select the integration profiles in which it will participate.
- For each actor and profile, select which options will be implemented.

All required transactions must be implemented for the profile to be supported (refer to the transaction descriptions in ITI TF-2a and ITI TF-2b).

995 Implementers should provide a statement describing which IHE actors, IHE integration profiles and options are incorporated in a given product. The recommended form for such a statement is defined in ITI TF-1: Appendix C.

1000 In general, a product implementation may incorporate any single actor or combination of actors. When two or more actors are grouped together, internal communication between actors is assumed to be sufficient to allow the necessary information flow to support their functionality; for example, the Context Manager uses the Patient Identifier Cross-reference Consumer Actor to obtain the necessary patient identifier mapping information from the Patient Identifier Cross-reference Manager. The exact mechanisms of such internal communication are outside the scope of the IHE Technical Framework.

1005 When multiple actors are grouped in a single product implementation, all transactions originating or terminating with each of the supported actors shall be supported (i.e., the IHE transactions shall be offered on an external product interface).

The following examples describe which actors typical systems might be expected to support. This is not intended to be a requirement, but rather to provide illustrative examples.

1010 A departmental system, such as a laboratory information system or a radiology picture archiving and communication system might include an Information Source Actor as well as a Kerberized Server Actor.

A clinical repository might include an Information Source Actor as well as a Kerberized Server Actor and a Patient Identifier Cross-reference Consumer Actor.

1015 A context management server might include a Context Management Actor as well as a Patient Identifier Cross-reference Consumer Actor.

3 Retrieve Information for Display (RID)

1020 The *Retrieve Information for Display Integration Profile (RID)* provides simple and rapid read-only access to patient-centric clinical information that is located outside the user’s current application but is important for better patient care (for example, access to lab reports from radiology department). It supports access to existing persistent documents in well-known presentation formats such as CDA (Level 1), PDF, JPEG, etc. It also supports access to specific key patient-centric information such as allergies, current medications, summary of reports, etc. for presentation to a clinician. It complements workflows with access from within the users’ on-screen workspace or application to a broad range of information.

1025 In this profile, the Information Source is solely responsible to turn the healthcare specific semantics into what this IHE Integration Profile calls a “presentation” format. As a consequence the Display Actor may process and render this “presentation” format with only generic healthcare semantics knowledge. Different formats have specific characteristics in terms of (1) server imposed limitations and (2) flexibility of display on the client side to render within its display constraints (e.g., a generic CDA level 1 style sheet).

1030 The Information Source is entirely responsible for the information returned for display and its clinical accuracy.

1035 This profile offers the capability to leverage industry standards that address both the structure and content of documents that may be returned by information sources. Where this profile references HL7 Clinical Documentation Architecture (CDA), it limits itself to the approved CDA Level 1. Furthermore, it only uses a subset of CDA Level 1 that facilitates making information available for display.

1040 Future extensions to the IHE IT Infrastructure TF will more fully leverage CDA Release 2 and other industry standards, and will incorporate vocabularies such as SNOMED and Clinical LOINC as well as clinical templates.

1045 This profile does not provide specific requirements on the means of assuring access control or security of information in transit. Such measures shall be implemented through appropriate security-related integration profiles, such as Enterprise User Authentication (see ITI TF-1:4). ITI TF-1: Appendix E describes the process flows for usage of the Retrieve Information for Display Integration Profile in conjunction with the Enterprise User Authentication and Patient Identifier Cross-referencing Integration Profiles.

3.1 Actors/ Transactions

1050 Figure 3.1-1 shows the actors directly involved in the Retrieve Information for Display Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in User Authentication and Patient Identifier Cross-referencing are not shown.

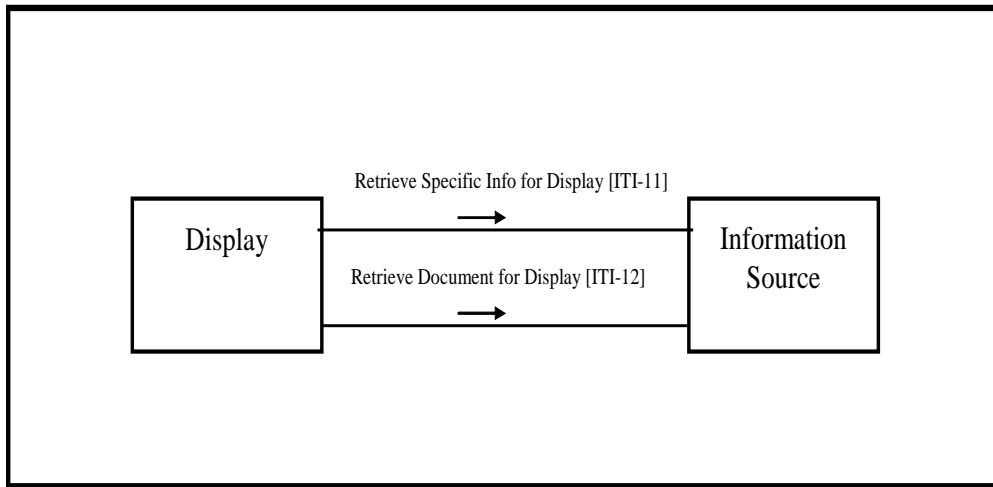


Figure 3.1-1: Retrieve Information for Display Actor Diagram

1055 Table 3.1-1 lists the transactions for each actor directly involved in the Retrieve Information for
 Display Integration Profile. In order to claim support of this Integration Profile, an
 implementation must perform the required transactions (labeled “R”). A complete list of options
 defined by this Integration Profile and that implementations may choose to support is listed in
 ITI TF-1: 3.2.

1060

Table 3.1-1: Retrieve Information for Display Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Display	Retrieve Specific Info for Display [ITI-11]	R	ITI TF-2a: 3.11
	Retrieve Document for Display [ITI-12]	R	ITI TF-2a: 3.12
Information Source	Retrieve Specific Info for Display [ITI-11]	R (see below)	ITI TF-2a: 3.11
	Retrieve Document for Display [ITI-12]	R (see below)	ITI TF-2a: 3.12

Transaction [ITI-11] is required if one of the following Options is selected by the Information Source Actor (see ITI TF-1: 3.2):

1065

Summary of All Reports
Summary of Laboratory Reports
Summary of Radiology Reports
Summary of Cardiology Reports
Summary of Surgery Reports
Summary of Intensive Care Reports
Summary of Emergency Reports
Summary of Discharge Reports
Summary of Prescriptions
List of Allergies and Adverse Reactions
List of Medications

1070

1075 Transaction [ITI-12] is required if the Persistent Document Option is selected by the Information Source Actor (See ITI TF-1: 3.2).

The means for a Display Actor to obtain documents' unique identifiers in order to retrieve them via Transaction [ITI-11] may be either via Transaction [ITI-12] or by other means that are outside the scope of the RID Integration Profile.

3.2 Retrieve Information for Display Integration Profile Options

1080 Options that may be selected for this Integration Profile are listed in the Table 3.2-1 along with the IHE actors to which they apply.

Table 3.2-1: Retrieve Information for Display - Actors and Options

Actor	Options	Vol. & Section
Display	<i>None</i>	- -
Information Source	<i>Persistent Document</i>	ITI TF-2a: 3.12
	<i>Summary of All Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Laboratory Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Radiology Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Cardiology Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Surgery Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Intensive Care Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Emergency Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Discharge Reports (note2)</i>	ITI TF-2a: 3.11
	<i>Summary of Prescriptions (note2)</i>	ITI TF-2a: 3.11
	<i>List of Allergies and Adverse Reactions</i>	ITI TF-2a: 3.11
	<i>List of Medications (note1)</i>	ITI TF-2a: 3.11

1085 Note1: List of Medications includes the list of medications currently known to be administered to the patient. It differs from the Summary of Prescriptions, in that the latter reflects what has been prescribed to the patient, but are not necessarily any longer administered.

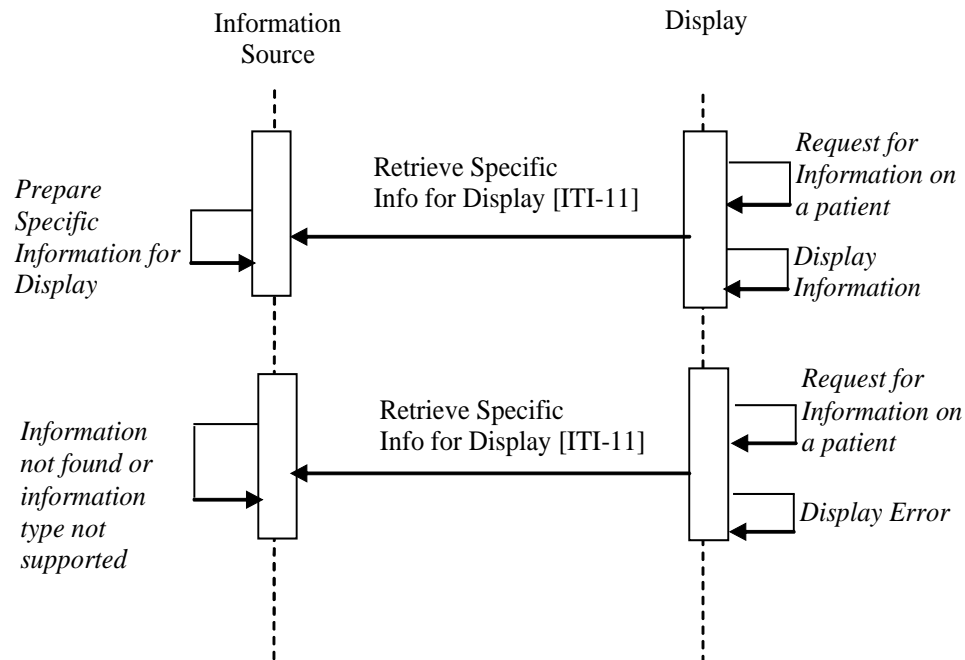
1090 Note2: In all the above options, "summary of reports" means that a general patient context (patient name, etc.) is provided along with a list of entries, where an entry includes key attributes such as date, specialty, and additional information sufficient to allow the viewer to select an entry. An entry may reference a persistent document for RID or other application defined RID summaries. Beyond these general guidelines, the specific content may likely be influenced by the context of use and customer desires. Such summaries are non-persistent in that they are likely to be updated in the course of patient care.

3.3 Retrieve Information for Display Process Flow

1095 This section describes the process and information flow when displayable patient information is retrieved from an information source. Three cases are distinguished.

Case 1-Retrieve *Specific* Information for Display: The first case describes use cases when the Display Actor and the person associated are requesting some information related to a patient. A somewhat specific request for information is issued (e.g., Retrieve a summary of laboratory reports) for a specific Patient ID to an Information Source Actor. The patient ID is assumed to be

1100 unambiguous as fully qualified with the assigning authority. A number of additional filtering
 keys may be used (last N reports, date range, etc.) depending on the specific type of request
 issued. The Information Source Actor responds with presentation-ready information that it
 considers relevant to the request. This Integration Profile leaves entire flexibility to the
 Information Source Actor to organize the content and presentation of the information returned.
 1105 The Display Actor simply displays the information to the person that triggered the request. The
 Information Source Actor shall respond with an error message when it does not support the
 specific type of request or does not hold any records for the requested patient ID.



1110 **Figure 3.3-1: Case 1: Retrieve Specific Information for Display Process Flow**

Case 2 - Retrieve a Document: The second case describes use cases when the Display Actor and the person associated are requesting a uniquely identified document such as a report, an image, an ECG strip, etc. The Information Source Actor responds to the request by using one of the proposed formats to provide the presentation-ready content of the object it manages. The detailed presentation and the clinical integrity of the content of the document are under the control of the Information Source Actor. The Display Actor simply displays the presentation-ready document content to the person that triggered the request. The Information Source Actor shall respond with an error message when the requested document is unknown or when none of the formats acceptable to the Display Actor is suitable to present the requested document.

The main difference between the Retrieve *Specific* Information and the Retrieve *Document* transactions is that the latter applies to a uniquely identifiable persistent object (i.e., retrieving the same document instance at a different point in time will provide the same semantics for its presented content). For the Retrieve Specific Information transaction, this information is always

related to a well-identified patient (Patient ID), but its content, although of a specific type (lab summary, or radiology summary, list of allergies) is generally dynamic (i.e., retrieving the same type of specific information at a different point in time is likely to result in different content; for example, a list of allergies may have been updated between two requests).

1130

Note: This Integration profile is not intended for highly dynamic information such as that used for patient monitoring.

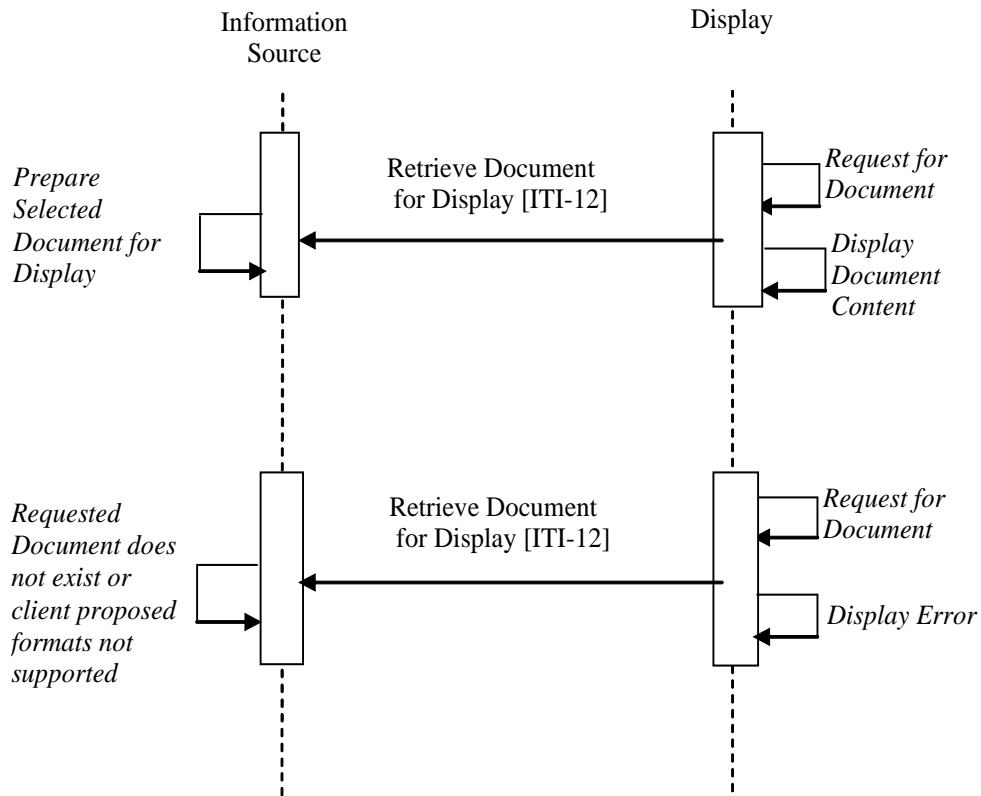


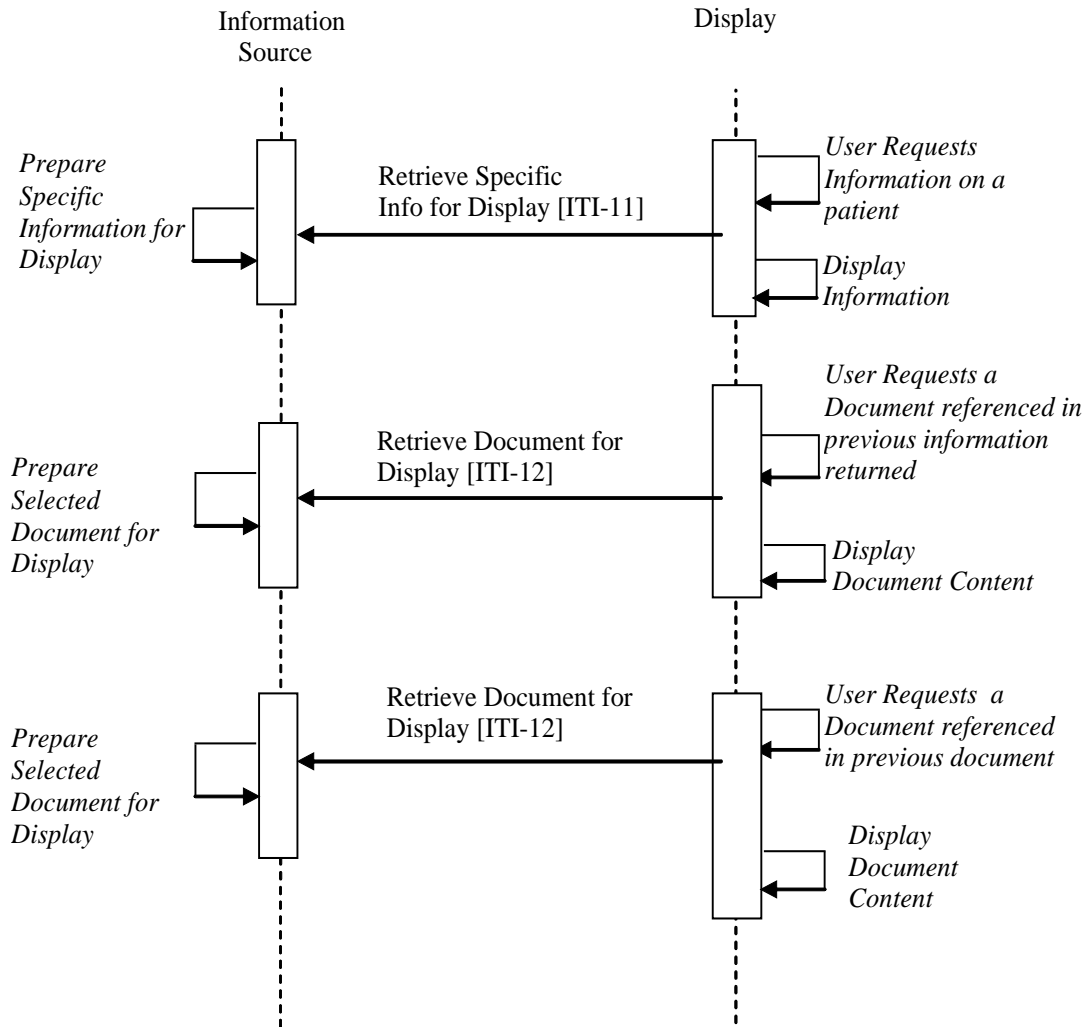
Figure 3.3-2: Case 2: Retrieve a Document Process Flow

Case 3 - Retrieve Specific Information for Display and Retrieve several Documents

1135

Process Flow: The third case combines the two cases above with the capability to associate in sequence the Retrieve Specific Information and the Retrieve Document for Display transactions. This allows for links to persistent documents within the returned specific information or for having persistent documents reference other persistent documents. For example, the user requests a summary of recent discharge reports, and then selects a specific document referenced in that summary list. From the discharge report displayed to the user, the user selects a specific surgery report. This surgery report is retrieved and displayed.

1140



1145 **Figure 3.3-3: Case 3: Retrieve Summary Information for Display and Retrieve several Documents Process Flow**

1150 The same Display Actor may involve more than one Information Source Actor by sequentially issuing different transactions. This Integration Profile assumes that the Display Actors may be configured a priori with one or more remote Information Source Actors along with the type of retrieve transactions/type of requests/specific keys suitable for the application context from which this Retrieve Information for Display requests are issued. Future Integration Profiles may facilitate such site-specific configuration tasks.

4 Enterprise User Authentication (EUA)

1155 *Enterprise User Authentication Profile (EUA)* – This defines a means to establish one name per user that can then be used on all of the devices and software that participate in this integration profile. It greatly facilitates centralized user authentication management and provides users with the convenience and speed of a single sign-on. This profile leverages Kerberos (RFC 1510) and the HL7 CCOW standard (user subject). User authentication is a necessary step for most application and data access operations and it is a workflow improvement for the users. The IHE
1160 EUA Profile adds value to the CCOW specification for the user subject by specifying the user subject and CCOW user subject suffix. This profile does not address security features such as audit trails, access control, authorization management and PKI. Future profiles will be developed to address these security features in a manner complementary to this EUA Profile.

1165 The environment is assumed to be a single enterprise, governed by a single security policy and having a common network domain. Unsecured domains -- in particular, Internet access -- are of interest, but not in the scope of this profile. Considerations for applications such as telemedicine and patient remote access to healthcare data are therefore also not in its scope. See ITI TF-1: Appendix G.

1170 Node and machine authentication is specified in the IHE Audit Trail and Node Authentication (ATNA) Profile and is not part of this profile.

4.1 Actors/ Transactions

1175 A number of transactions used in this profile conform to the Kerberos v5 standard, defined in RFC 1510. This standard has been stable since 1993, is widely implemented on current operating system platforms, has successfully withstood attacks in its 10-year history, and is fully interoperable among platforms. For example, Sun Solaris, Linux, AIX, HPUX, IBM-z/OS, IBM-OS400, Novell, MAC OS X, and Microsoft Windows 2000/XP all implement Kerberos in an interoperable manner. This is not a complete list; many other vendors also support Kerberos.

For additional detailed information on Kerberos, beyond what is specified in this profile, we suggest these references:

- 1180
- RFC 1510 - <http://www.ietf.org/rfc/rfc1510.txt>
 - MIT's Kerberos home page - <http://web.mit.edu/kerberos/www/>
 - The Moron's Guide to Kerberos - <http://www.isi.edu/~brian/security/kerberos.html>
 - Microsoft Kerberos information
<http://www.microsoft.com/TechNet/prodtechnol/windows2000serv/deploy/kerberos.asp>

1185 Kerberos implementations are widely available worldwide. Kerberos does include cryptography that may have restricted use laws in some countries. The US export regulations can be found at <http://www.bxa.doc.gov/Encryption>.

1190 Figure 4.1-1 shows the actors directly involved in the Enterprise User Authentication Profile and the relevant transactions between them. The box labeled "Other IHE Actor" represents actors from other integration profiles that are meant to be grouped with the nearby actor from within

this profile. Other actors that may be indirectly involved due to their use of authentication, etc. are not shown.

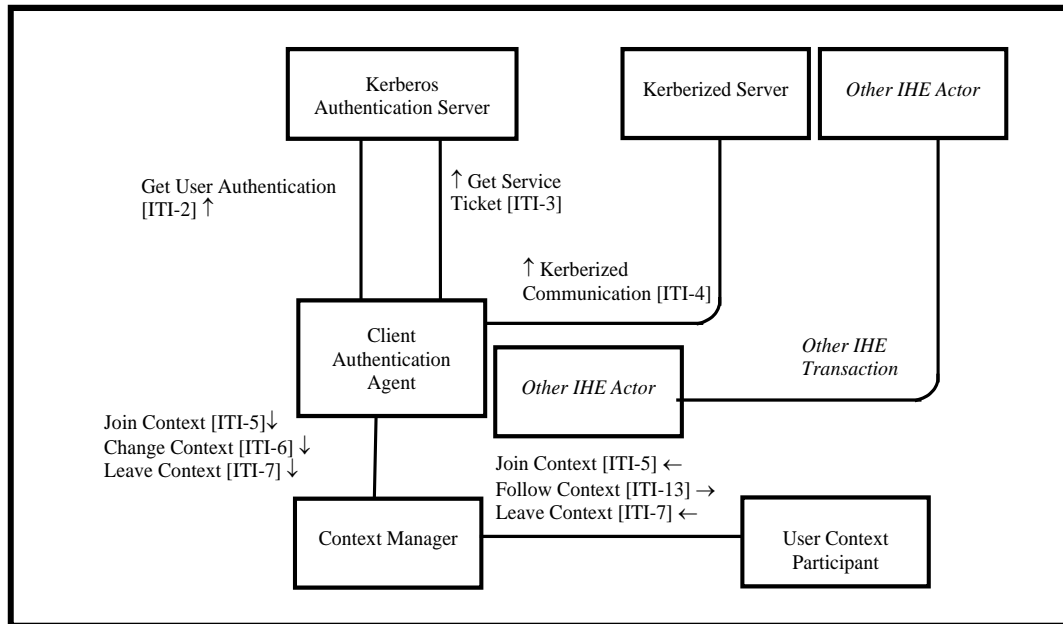


Figure 4.1-1: Enterprise Authentication Actor Diagram

1195 Table 4.1-1 lists the transactions for each actor directly involved in the Enterprise User Authentication Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled "O" are optional. A complete list of options defined in this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 4.2.

1200

Table 4.1-1: Enterprise User Authentication Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Kerberos Authentication Server	Get User Authentication [ITI-2]	R	ITI TF-2a: 3.2
	Get Service Ticket [ITI-3]	R	ITI TF-2a: 3.3
Client Authentication Agent	Get User Authentication [ITI-2]	R	ITI TF-2a: 3.2
	Get Service Ticket [ITI-3]	R	ITI TF-2a: 3.3
	Kerberized Communication [ITI-4]	R	ITI TF-2a: 3.4
	Join Context [ITI-5]	O [Note1]	ITI TF-2a: 3.5
	Change Context [ITI-6]	O [Note1]	ITI TF-2a: 3.6
	Leave Context [ITI-7]	O [Note1]	ITI TF-2a: 3.7
Kerberized Server	Kerberized Communication [ITI-4]	R	ITI TF-2a: 3.4
User Context Participant	Join Context [ITI-5]	R	ITI TF-2a: 3.5
	Follow Context [ITI-13]	R	ITI TF-2a: 3.13
	Leave Context [ITI-7]	R	ITI TF-2a: 3.7
Context Manager	Join Context [ITI-5]	R	ITI TF-2a: 3.5
	Follow Context [ITI-13]	R	ITI TF-2a: 3.13

Actors	Transactions	Optionality	Section
	Leave Context [ITI-7]	R	ITI TF-2a: 3.7
	Change Context [ITI-6]	R	ITI TF-2a: 3.6

Note 1: When the Authentication for User Context Option is supported, then the transaction is required.

1205 CCOW facilitates the sharing of the identity of a EUA authentication user but does not provide for the authentication of users. In order for the Context Manager and User Context Participant to participate in the EUA Profile it is required that the Client Authentication Agent supports the Authentication for User Option. This design provides the User Context Participant with a consistent and enterprise recognized user identity, but does not define access to the Kerberos credentials. Future IHE profiles may address this limitation. Note that the Client Authentication Agent is the key actor when PSA and EUA are combined. See the use case outlined in ITI TF-1: 4.3.2. Applications that implement both the Client Authentication Agent Actor and the User Context Participant Actor shall support configurations where either Actor is disabled.

1210 In any single user environment there shall be only one Client Authentication Agent for one user. In a multi-user environment there shall not be more than one Client Authentication Agent per user.

1215 4.2 Enterprise User Authentication Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 4.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

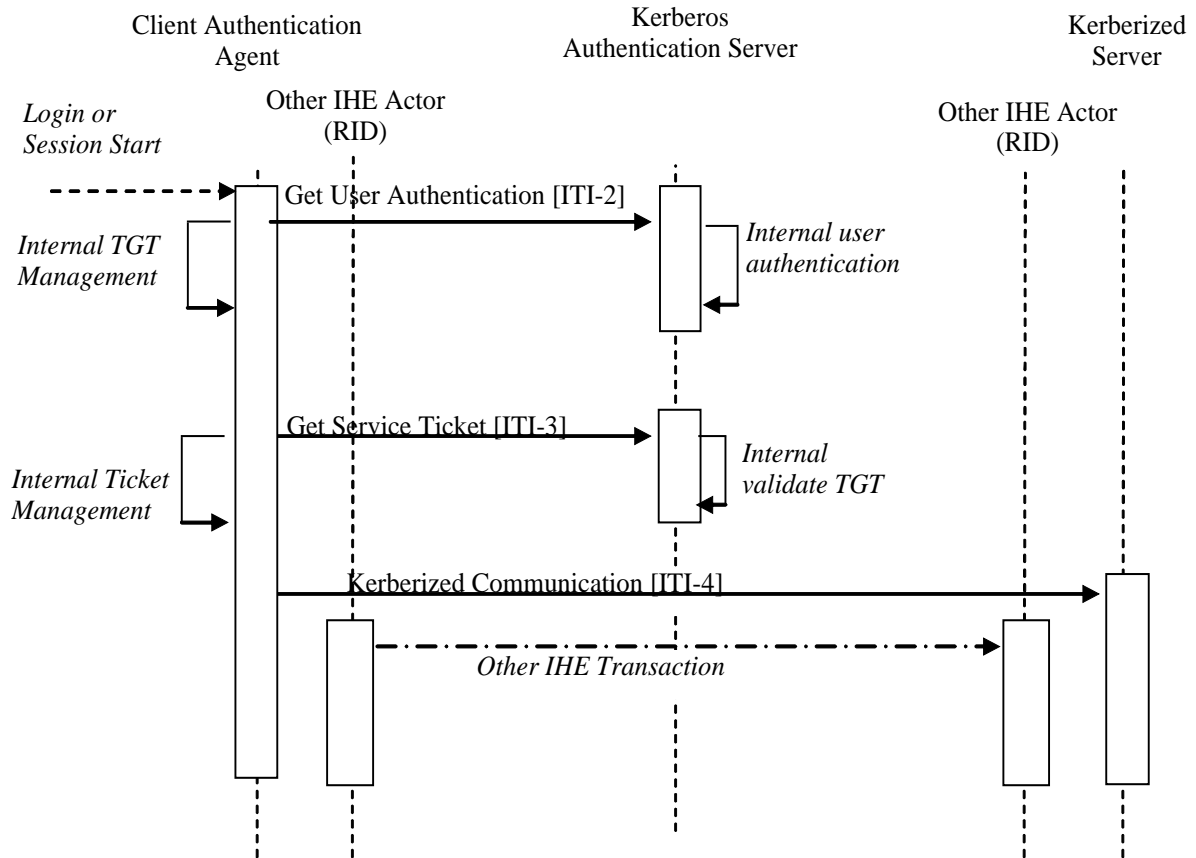
1220 **Table 4.2-1: Enterprise User Authentication - Actors and Options**

Actor	Options	Vol. & Section
Kerberos Authentication Server	<i>No options defined</i>	--
Client Authentication Agent	<i>Authentication for User Context</i>	ITI TF-2a: 3.6
Kerberized Server	<i>No options defined</i>	--
Context Manager	<i>No options defined</i>	--
User Context Participant	<i>No options defined</i>	--

4.3 Enterprise User Authentication Profile Process Flow

4.3.1 Basic User Authentication Process Flow

The following diagram describes the sequence of events in the use of Enterprise User Authentication:



1225

Figure 4.3.1-1: Basic Process Flow in Enterprise User Authentication Profile

The sequence of events in the use of Enterprise User Authentication is:

- 1230 • The user begins the session. This initiates a local username/password authentication that is converted into the challenge/response system used by Kerberos to avoid transmitting the password over the network. This information is used as part of the Get User Authentication Transaction to get a “Ticket Granting Ticket” (TGT).
- The TGT is saved and managed internally by the Client Authentication Agent Actor. The TGT acts as confirmation that the user has been authenticated.
- 1235 • For each service that has been Kerberized, the Client Authentication Agent Actor uses the Get Service Ticket Transaction to obtain a service ticket. The service ticket is then used as part of the Kerberized Communication Transaction.

1240 A Kerberized Communication is a Kerberos data exchange that is integrated into another protocol, such as HL7 or DICOM, which is used in another IHE transaction. The details of Kerberization vary and are described separately for the protocols that have been Kerberized. The Kerberization enables the other IHE Actors involved in the other transaction to use the identity of the authenticated user for purposes such as user authorization or audit messages.

1245 The Client Authentication Agent Actor also maintains an internal cache of credentials such as the TGT and service tickets. It renews the tickets as necessary to deal with ticket expirations, re-uses tickets while they are still valid, and removes credentials from the cache when the user session ends. The Client Authentication Agent shall make the Kerberos credentials available using the local operating system mechanisms. Other IHE Actors that need the Kerberos credentials are strongly encouraged to obtain them using the local operating system mechanisms. Operating system support for ticket management has been implemented and has been defined for various
 1250 operating systems.

4.3.2 User Authentication with User Synchronized Applications Process Flow

In this use case an application supporting user authentication on the same desktop as another application is synchronized to the same user identity, thus giving the user a single-sign-on experience.

1255 The following diagram describes the sequence of events in the use of User Authentication with User Synchronized Applications:

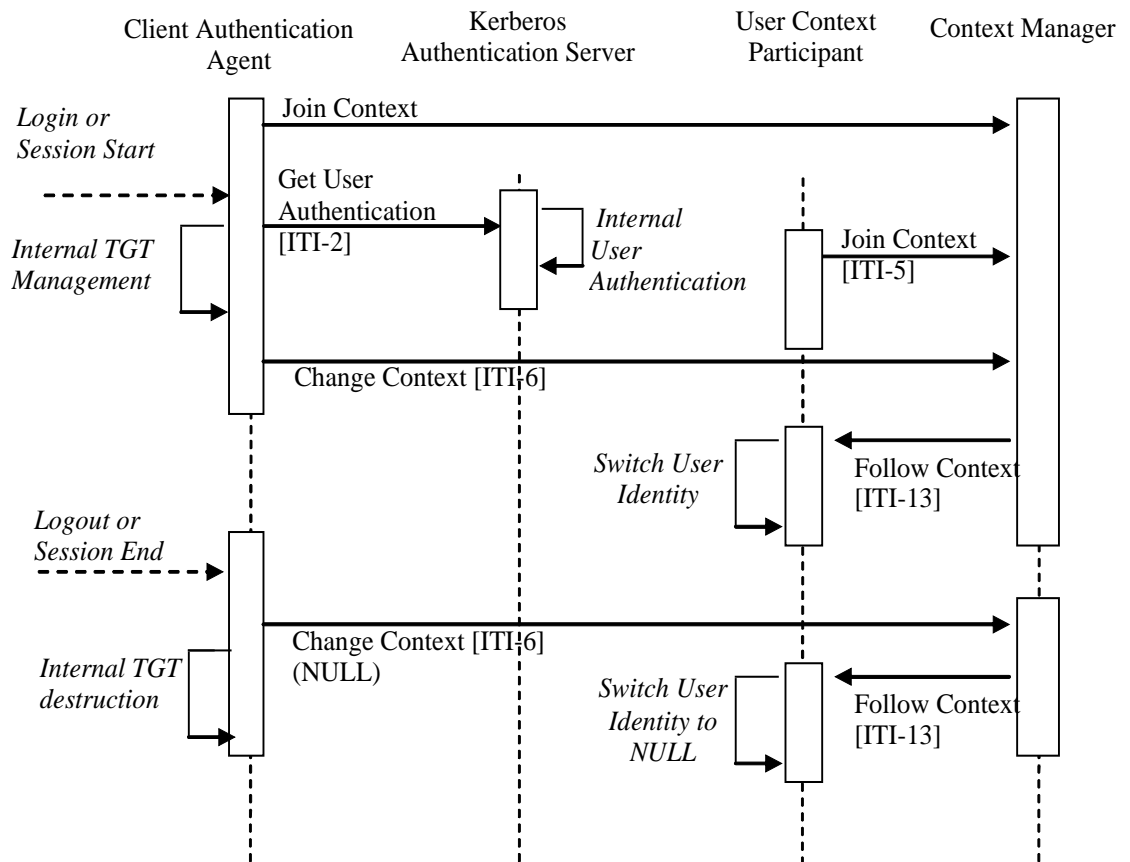


Figure 4.3.2-1: Process Flow with User Synchronized Applications

The sequence of events of the User Authentication with User Synchronized Applications is:

- 1260 • The user initiates a login by starting the Client Authentication Agent.

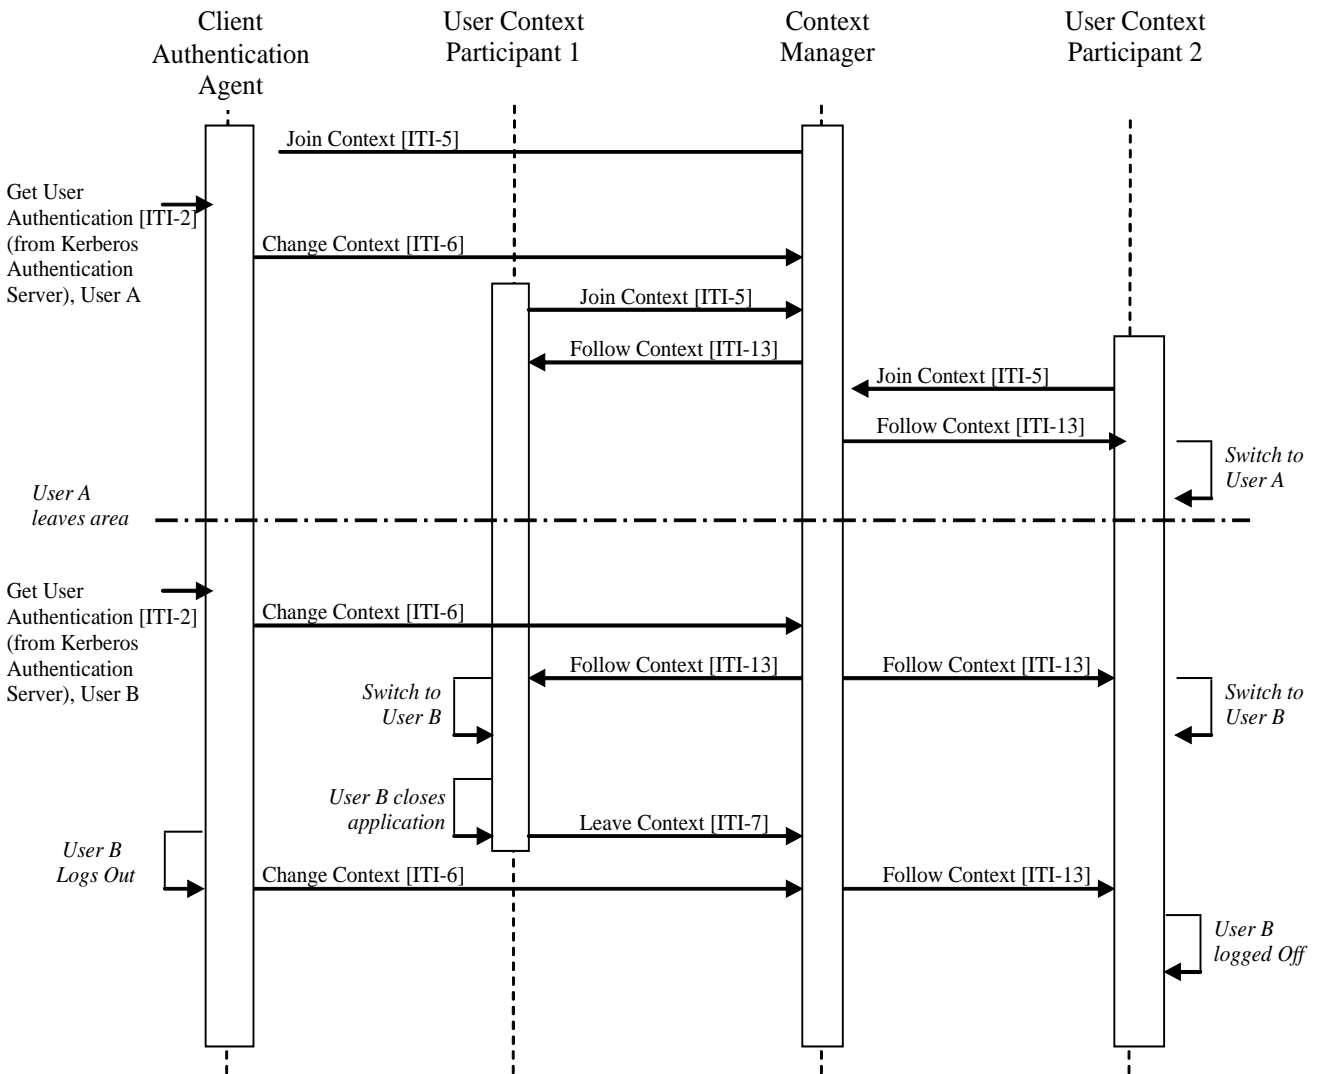
- The Client Authentication Agent joins the CCOW user context by sending a Join Context Transaction to the Context Manager Actor. At this point there is no user identity in the context.
- 1265 • The user provides their username and password to the Client Authentication Agent. This authentication information is converted into the challenge/response system used by Kerberos to avoid transmitting the password over the network. This information is used as part of the Get User Authentication Transaction to get a “Ticket Granting Ticket” (TGT).
- 1270 • The TGT is saved and managed internally by the Client Authentication Agent Actor. The TGT acts as confirmation that the user has been authenticated.
- A Change Context Transaction is sent to the Context Manager Actor with the users fully qualified user name.
- The user is now logged in to the User Context Participant.
- 1275 • When the user ends the session, a Change Context Transaction is sent to the Context Manager Actor with a NULL user name.
- The user is logged out of the User Context Participant.

4.3.3 Fast User Switching with Multiple Applications Process Flow

1280 The use model in the clinical environment can be characterized as multiple clinicians using the same workstation for short intervals of time many times a day. In this shared workstation environment the user requires quick access to the patient data contained in the applications. Traditional methods of logging in and out of the workstation at the operating system or network level can take too long and typically force the applications to terminate. This means that the application clients will potentially need to initialize and establish new database connections, introducing further delay to the Clinician access to patient data. The CCOW standard and more specifically the “user” subject provides a means in combination with the Enterprise Authenticator to allow the user to authenticate at the application level and have all of the other applications tune to the new user.

1285

The following diagram describes the sequence of events in the case of Fast User Switching with Multiple Applications:



1290

Figure 4.3.3-1: Fast User Switching when using Multiple Applications

The process flow would be similar to the following:

1295 Clinician A launches and authenticates via an application containing the Client Authentication Agent (refer to Figure 4.3.3-1 for details). This actor joins the context session and performs a context change to set Clinician A as the user in context.

1300 Clinician A launches the clinical data repository application, containing a User Context Participant Actor, depicted as User Context Participant 1. The actor joins the context session, gets the current user from the Context Manager, and logs clinician A into the application.

Clinician A launches a cardiology application, containing a User Context Participant Actor, depicted as User Context Participant 2. The actor joins the context session, gets the current user from the Context Manager, and logs clinician A into the application.

Clinician A does his job and then gets called away and leaves the workstation.

1305 Clinician B approaches the workstation and authenticates using the Client Authentication Agent. This results in a context change from Clinician A to Clinician B being set in context without the delay typically associated with a logout and login at the operating system level. The clinical data repository and the cardiology application are notified of the context change by the Context Manager resulting in Clinician A being logged out of both applications and Clinician B being
1310 logged into both applications.

Clinician B does his job and then closes the clinical data repository application, which leaves the context prior to terminating the application.

Clinician B is finished reviewing patient data within the cardiology application and logs out using the Client Authentication Agent. This forces a context change to remove the current user
1315 from the context, which results in the user being logged out of the cardiology application.

5 Patient Identifier Cross-referencing (PIX)

1320 The *Patient Identifier Cross-referencing Integration Profile (PIX)* is targeted at healthcare enterprises of a broad range of sizes (hospital, a clinic, a physician office, etc.). It supports the cross-referencing of patient identifiers from multiple Patient Identifier Domains via the following interactions:

- The transmission of patient identity information from an identity source to the Patient Identifier Cross-reference Manager.
 - The ability to access the list(s) of cross-referenced patient identifiers either via a query/response or via update notification.
- 1325

By specifying the above transactions among specific actors, this integration profile does not define any specific enterprise policies or cross-referencing algorithms. By encapsulating these behaviors in a single actor, this integration profile provides the necessary interoperability while maintaining the flexibility to be used with any cross-referencing policy and algorithm as deemed adequate by the enterprise.

1330

The following diagram shows the intended scope of this profile (as described above).

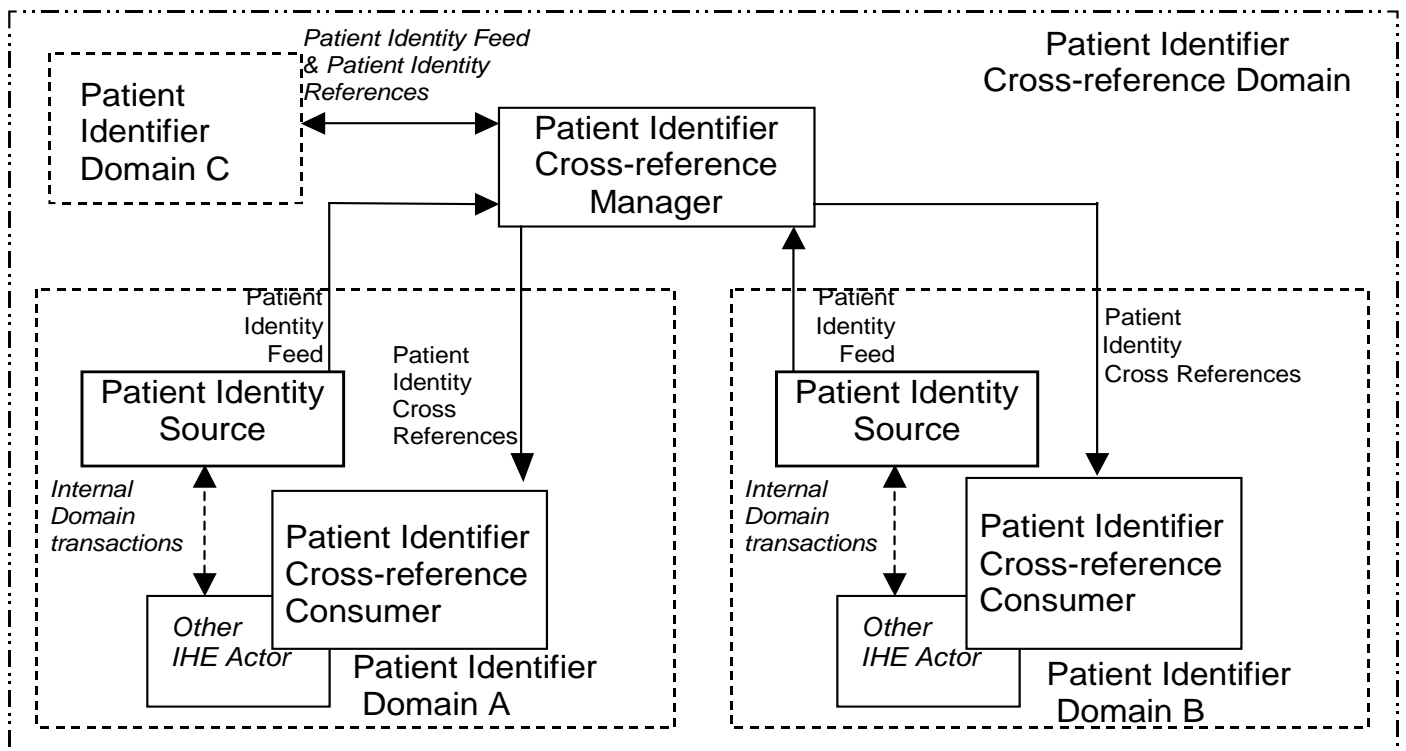


Figure 5-1: Process Flow with Patient Identifier Cross-referencing

1335 The diagram illustrates two types of Identifier Domains: a Patient Identifier Domain and a Patient Identifier Cross-reference Domain.

A Patient Identifier Domain is defined as a single system or a set of interconnected systems that all share a common identification scheme (an identifier and an assignment process to a patient) and issuing authority for patient identifiers. Additionally, a Patient Identifier Domain has the following properties:

- 1340 • A set of policies that describe how identities will be defined and managed according to the specific requirements of the domain.
- An administration authority for administering identity related policies within the domain.
- A **single** system, known as a patient identity source system, that assigns a unique identifier to each instance of a patient-related object as well as maintaining a collection of identity traits.
- 1345 • Ideally, only one identifier is uniquely associated with a single patient within a given Patient Identifier Domain, though a single Patient Identity Source Actor may assign multiple identifiers to the same patient and communicate this fact to the Patient Identifier Cross-reference Manager. For a description of how the Patient Identifier Cross-reference Manager Actor responds to requests for a list of cross-referenced identifiers that include these “duplicates” see ITI TF-2a: 3.9.4.2.2.6).
- 1350 • An “Identifier Domain Identifier” (known as assigning authority) that is unique within a Patient Identifier Cross-reference Domain.
- 1355 • Other systems in the Patient Identifier Domain rely upon the identifiers assigned by the patient identity source system of the domain to which they belong.

A Patient Identifier Cross-reference Domain consists of a set of Patient Identifier Domains known and managed by a Patient Identifier Cross-reference Manager Actor. The Patient Identifier Cross-reference Manager Actor is responsible for creating, maintaining and providing lists of identifiers that are aliases of one another across different Patient Identifier Domains.

1360 The Patient Identifier Cross-reference Domain embodies the following assumptions about agreement within the group of individual Identifier Domains:

- They have agreed to a set of policies that describe how patient identities will be cross-referenced across participating domains;
- They have agreed to a set of processes for administering these policies;
- 1365 • They have agreed to an administration authority for managing these processes and policies.

1370 All these assumptions are critical to the successful implementation of this profile. This integration profile imposes minimal constraints on the participating Patient Identifier Domains and centralizes most of the operational constraints for the overall Patient Identification Cross-reference Domain in the Patient Identifier Cross-reference Manager Actor. If the individual

Identifier Domains cannot agree to the items outlined above, implementation of this profile may not provide the expected results.

1375 The Patient Identifier Cross-reference Manager Actor is not responsible for improving the quality of identification information provided to it by the Identity Source Actors. It is assumed that the Identity Source actors are responsible for providing high quality data to the Patient Identifier Cross-reference Manager. For example, the Patient Identifier Cross-reference Manager Actor is NOT responsible to provide a single reference for patient demographics. The intent is to leave the responsibility for the quality and management of its patient demographics information and the integrity of the identifiers it uses within each Patient Identity Domain (Source actors).
 1380 When receiving reports and displays from multiple PIX domains, it is inevitable that some of those reports and displays will have inconsistent names.

1385 The Patient Identifier Cross-reference Consumer may use either a query for sets of cross-reference patient identifiers or use both a notification about cross-reference changes and a query transaction. In the case of using a notification, the Patient Identifier Cross-reference Consumer may also use the PIX Query Transaction to address situations where the Patient Identifier Cross-reference Consumer may be out of synch with the Patient Identifier Cross-reference Manager. This Integration Profile does not specify the consumer policies in using the PIX Query Transaction (ITI TF-2a: 3.9).

1390 For a discussion of the relationship between this Integration Profile and an enterprise master patient index (eMPI) see ITI TF-1: 5.4.

5.1 Actors/ Transactions

Figure 5.1-1 shows the actors directly involved in the Patient Identifier Cross-referencing Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other related profiles are not shown.

1395

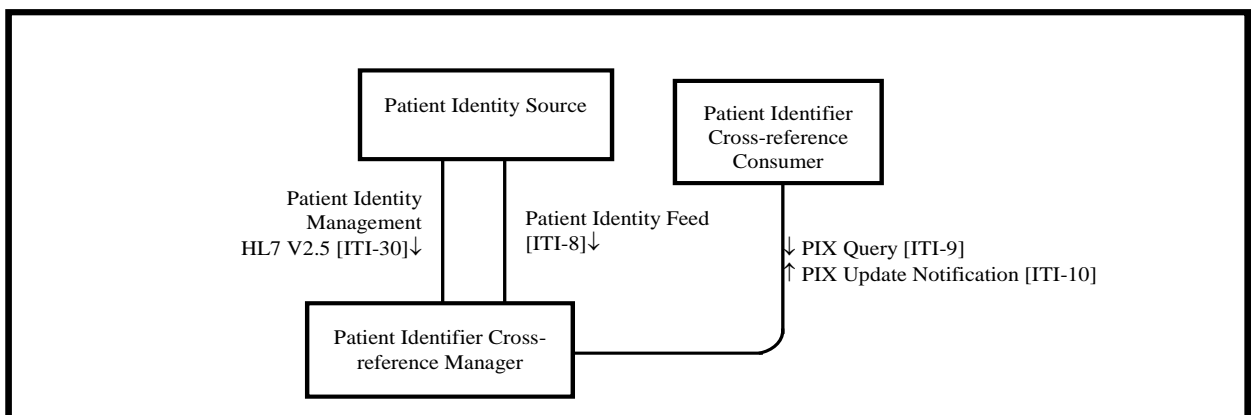


Figure 5.1-1: Patient Identifier Cross-referencing Actor Diagram

1400 Table 5.1-1 lists the transactions for each actor directly involved in the Patient Identifier Cross-referencing Profile. In order to claim support of this Integration Profile, an implementation must

perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in the ITI TF-1: 5.2.

1405 **Table 5.1-1: Patient Identifier Cross-referencing Integration for MPI Profile - Actors and Transactions**

Actors	Transactions	Optionality	Section
Patient Identity Source	Patient Identity Feed [ITI-8]	R	ITI TF-2a: 3.8
	Patient Identity Management [ITI-30]	O	ITI TF-2b: 3.30
Patient Identifier Cross-reference Consumer	PIX Query [ITI-9]	R	ITI TF-2a: 3.9
	PIX Update Notification [ITI-10]	O	ITI TF-2a: 3.10
Patient Identifier Cross-reference Manager	Patient Identity Feed [ITI-8]	R	ITI TF-2a: 3.8
	Patient Identity Management [ITI-30]	O	ITI TF-2b: 3.30
	PIX Query [ITI-9]	R	ITI TF-2a: 3.9
	PIX Update Notification [ITI-10]	R	ITI TF-2a: 3.10

5.2 Patient Identifier Cross-referencing Integration Profile Options

Options that may be selected for this Integration Profile are listed in the Table 5.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

1410

Table 5.2-1: Patient Identifier Cross-referencing - Actors and Options

Actor	Options	Vol. & Section
Patient Identity Source	<i>Pediatric Demographics</i>	ITI TF-1: 5.2.1
Patient Identifier Cross-reference Manager	<i>Pediatric Demographics</i>	ITI TF-1: 5.2.1
Patient Identifier Cross-reference Consumer	<i>PIX Update Notification</i>	ITI TF-2a: 3.10

5.2.1 Pediatric Demographics

1415

The experience of immunization registries and other public health population databases has shown that matching and linking patient records from different sources for the same individual person in environments with large proportions of pediatric records requires additional demographic data.

1420

In particular, distinguishing records for children who are twins, triplets, etc. – that is, avoiding false positive matches - may be difficult because much of the demographic data for the two individuals matches. For instance, twin children may have identical last names, parents, addresses, and dates of birth; their first names may be very similar, possibly differing by only one letter. It can be very difficult for a computer or even a human being to determine in this

1425 situation whether the slight first name difference points to two distinct individuals or just a typographical error in one of the records. Additional information is extremely helpful in making this determination.

Pediatric Demographics makes use of the following six additional demographic fields to aid record matching in databases with many pediatric records.

Table 5.2.1-1: Additional Patient Demographics Fields

Field	Reason for inclusion	Value
Mother's Maiden Name	Any information about the mother is helpful in making a match	Helps create true positive matches
Patient Home Telephone	A telecom helps match into the right household	Helps create true positive matches
Patient Multiple Birth Indicator	Indicates this person is a multiple - twin, triplet, etc.	Helps avoid false positive matches of multiples
Patient Birth Order	Distinguishes among those multiples.	Helps avoid false positive matches of multiples
Last Update Date/Time, Last Update Facility	These fields, although not strictly demographic, can effectively substitute when multiple birth indicator and birth order are not collected. They indirectly provide visit information. Provider visits on the same day may likely indicate two children brought to a doctor together.	Helps avoid false positive matches of multiples

1430 Patient Identity Source actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction and shall provide values, when available, for the fields identified as Pediatric Demographics fields.

1435 Patient Identifier Cross-reference Manager Actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction, and if values for one or more of the Pediatric Demographics fields are specified in the Patient Identity Management [ITI-30], they shall be considered as part of the matching algorithm of the PIX Manager.

Pediatric Demographics are defined as all of the following:

- Mother's Maiden Name
- Patient Home Telephone
- Patient Multiple Birth Indicator
- Patient Birth Order
- Last Update Date/Time

- Last Update Facility

1445 Pediatric Demographic is particularly focused on two data issues:

- Locating a record where the data or the search criterion have differences, but both the data record and the search criterion represent the same person, and
- Avoiding improper linkage of very similar records that do not belong to the same person. This problem is most often encountered in multiple birth situations where twins may be given extremely similar names

1450

5.3 Patient Identifier Cross-referencing Profile Process Flows

The following sections describe use cases that this profile addresses.

5.3.1 Use Case: Multiple Identifier Domains within a Single Facility/ Enterprise

1455 A clinician in the Intensive Care Unit at General Hospital is reviewing a patient chart on the Intensive Care information system and wishes to review or monitor the patient's glucose level, which is included in a laboratory report stored in the hospital's main laboratory system. The Intensive Care system needs to map its own patient ID, which it generates internally, to the patient's medical record number (MRN), which is generated from the hospital's main ADT system and is used as the patient identity by the lab system. In this case the Intensive Care system is essentially in a different identifier domain than the rest of the hospital since it has its own notion of patient identity.

1460

1465 In this scenario, the hospital's main ADT system (acting as a Patient Identity Source) would provide a Patient Identity Feed (using the patient's MRN as the identifier) to the Patient Identifier Cross-reference Manager. Similarly, the Intensive Care system would also provide a Patient Identity Feed to the Patient Identifier Cross-reference Manager using the internally generated patient ID as the patient identifier and providing its own unique identifier domain identifier.

1470 Once the Patient Identifier Cross-reference Manager receives the Patient Identity Feed transactions, it performs its internal logic to determine which, if any, patient identifiers can be "linked together" as being the same patient based on the corroborating information included in the Feed transactions it has received. The cross-referencing process (algorithm, human decisions, etc.) is performed within the Patient Identifier Cross-reference Manager and is outside the scope of IHE. (See ITI TF-2a: 3.9.4.2.2.6 for a more complete description of the scope of the cross-referencing logic boundary).

1475 The Intensive Care system wants to get lab information associated with a patient that the Intensive Care system knows as patient ID = 'MC-123'. It requests the lab report from the lab system using its own patient ID (MC-123) including the domain identifier/ assigning authority. Upon receipt of the request, the lab system determines that the request is for a patient outside of its own identifier domain (ADT Domain). It requests a list of patient ID aliases corresponding to patient ID = 'MC-123' (within the "Intensive Care domain") from the Patient Identifier Cross-reference Manager. Having linked this patient with a patient known by medical record number = '007' in the 'ADT Domain', the Patient Identifier Cross-reference Manger returns this list to the

1480

lab system so that it may retrieve the lab report for the desired patient and return it to the Intensive Care system. Figure 5.3-1 illustrates this process flow.

1485

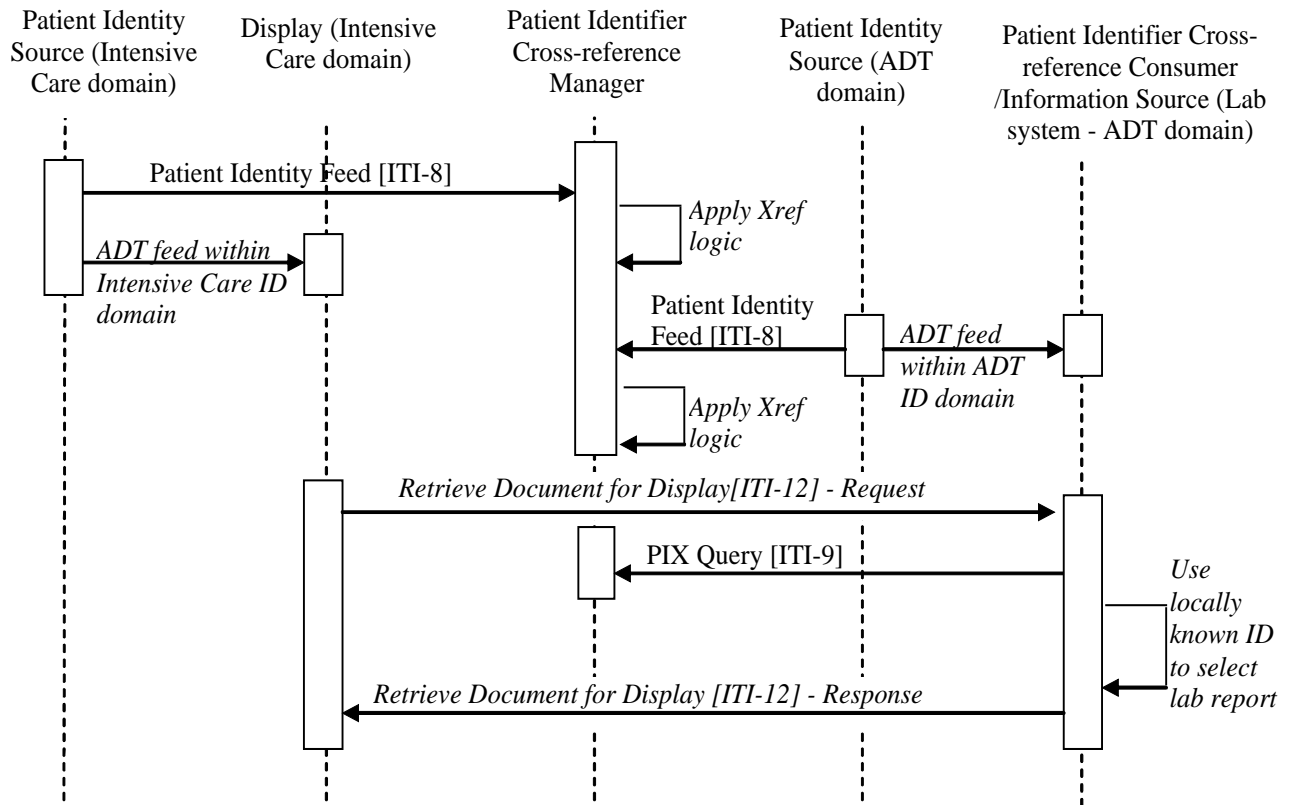


Figure 5.3-1: Multiple ID Domains in a Single Facility Process Flow in PIX Profile

Note: Request and Response portions of the Retrieve Document for Display transaction are not part of this profile and included for illustration purposes only.

1490

5.3.2 Use Case: Multiple ID Domains Across Cooperating Enterprises

A healthcare enterprise is established by the consolidation of two hospitals, each having its own separate patient registration process run by different hospital information systems. When a patient is treated in one hospital, the access to its electronic records managed by the other hospital is necessary. The following use case illustrates this scenario.

1495

Hospitals A and B have been consolidated and have a single Patient Identifier Cross-reference Manager that maintains the ID links between the two hospitals. Each hospital has a different HIS that is responsible for registering patients, but they have consolidated their cardiology information systems. The cardiology system has been configured with a Patient Identifier Cross-reference Consumer to receive patient identity notifications when cross-referencing activity occurs.

1500

A patient is registered and then has some diagnostic stress tests done at hospital A. The cardiology information system queries the Patient Identifier Cross-reference Manager to get a list

1505 of possible ID aliases for the patient to see if any past cardiology reports may be available. No patient ID aliases are found. Sometime later the same patient goes to hospital B to have a second diagnostic stress test done. The patient is registered via the HIS in hospital B which then sends that identity information to the Patient Identifier Cross-reference Manager. The Patient Identifier Cross-reference Manager determines this is in fact the same patient as was registered previously at hospital A. The cardiology information system was previously configured with the Patient Identifier Cross-reference Manager to receive notifications, thus a notification is sent to the cardiology system to inform it of the patient identifier aliases. This notification is done to allow systems that are aware of multiple identifier domains to maintain synchronization with patient identifier changes that occur in any of the identifier domains that they are aware of.

1510

Figure 5.3-2 illustrates the process flow for this use case.

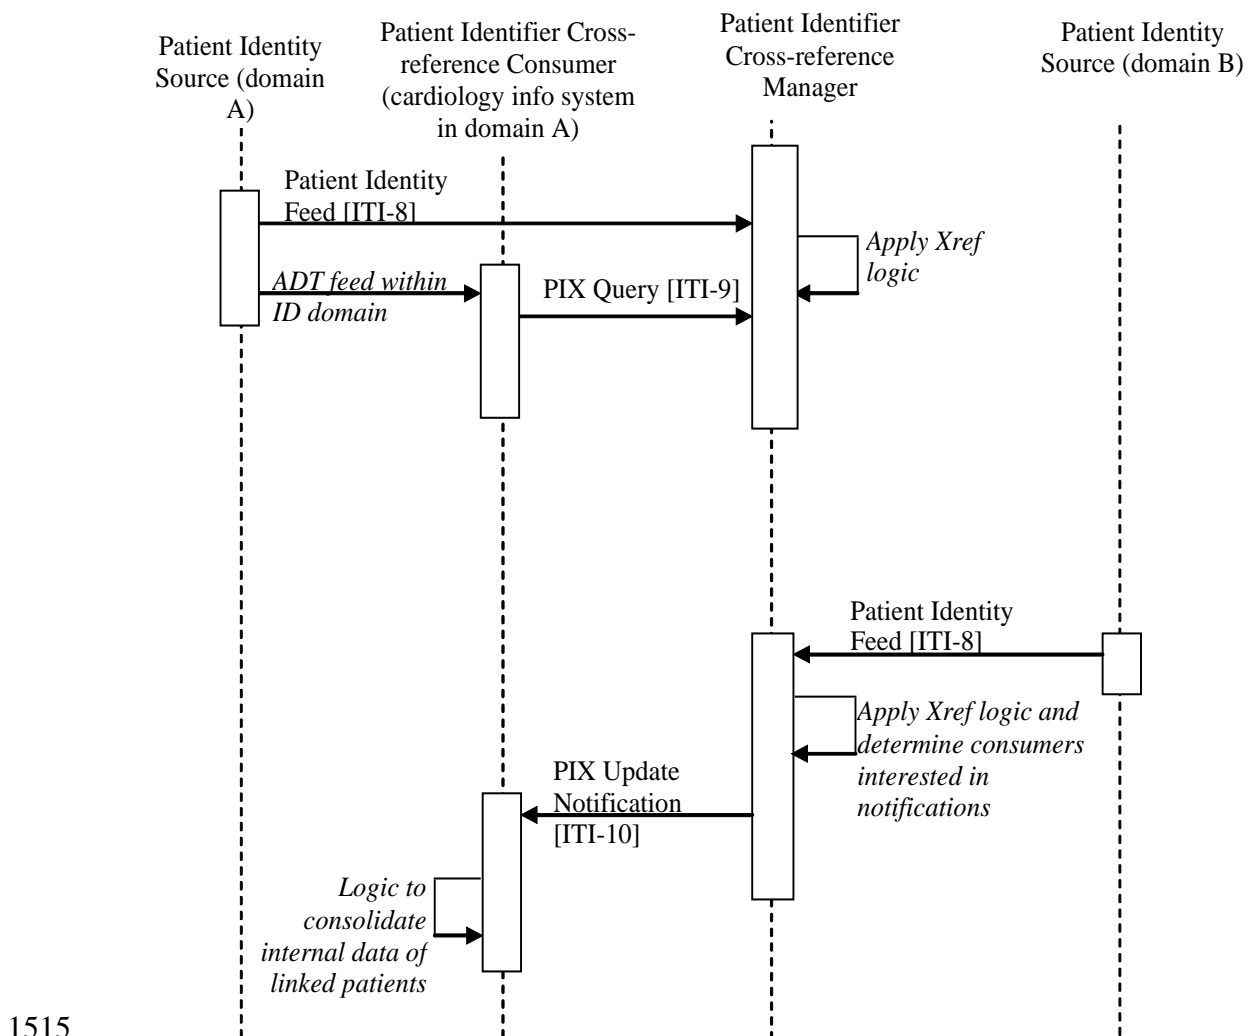


Figure 5.3-2: Multiple ID Domains Across Cooperating Enterprises Process Flow in PIX Profile

Note: PIX Update Notifications are not sent for the first Patient Identity Feed for a patient, since no cross-referencing activity occurred after this first Patient Identity Feed Transaction.

1520 **5.3.3 Pediatric Demographic Option Use Cases**

The following sections describe use cases that the Pediatric Demographics Option addresses.

5.3.3.1 Use Case: High Quality Demographic Feed from a Birth Registry

1525 A regional Immunization Information System (IIS) receives birth registry information about a pair of twins. These twins are named “Lalainne” and “Lalannie” Smith. All of the data elements in the received registration are populated, and they are all identical, except for the Given Name, the Birth Order, and the Birth Certificate #. The IIS cross-referencing system can clearly identify this very similar data as belonging to two separate individuals, because they are both flagged as having been part of a Multiple Birth, their Birth Order numbers are different, and their Birth Certificate #s are different.

1530 **5.3.3.2 Use Case: Normal Demographic Feed from a Point of Care**

A couple years later, the mother of these two twins, who has now divorced and remarried, takes them to Pediatric Healthcare, where they get the immunizations appropriate for 2 year olds. Pediatric Healthcare completes a registration for each of them, and submits the resulting data to the IIS. This data has their new Family Name as “Gomez,” but the clerks had appropriately recorded the Birth Order of each twin. Again, the IIS was able to distinguish the two registration records as belonging to separate individuals, and it was able to match them up to their earlier records because the mother’s Maiden Name was present in both the earlier records and the records submitted from Pediatric Healthcare. Pediatric Healthcare was able to download the full immunization history of each twin.

1540 **5.3.3.3 Use Case: Minimal Demographic Feed from a Health Fair**

1545 The Jackson County Health Department puts on an annual Health Fair in a shopping mall every August, partly to screen school age children for the minimum shots required for admission to the first grade. Mrs. Gomez is now working to pay for her new apartment, but her sister-in-law takes the children to the Health Fair where they are given shots based on the paper “yellow card” the sister-in-law brings with the two twins. Jackson County Health Department staff records the children’s names, and the shots they were given. This information is entered into the computer back at the Clinic the next day, and submitted to the regional IIS.

1550 At this point, even though both children’s names were misspelled as “Lane” and “Lanna”, the Immunization Registry was again able to recognize that the records belonged to twins rather than the same person because, although the demographic data was almost identical, the Last Update Date/Time were very close (Date was the same) and Last Update Facility indicated the same clinic. Unfortunately, they didn’t write down the mother’s information at the Health Fair, but recorded her sister-in-laws name and address instead, so the Immunization Registry was not able to automatically link this new information to the information it already had for “Lalainne” and “Lalannie”.

Other Possibilities:

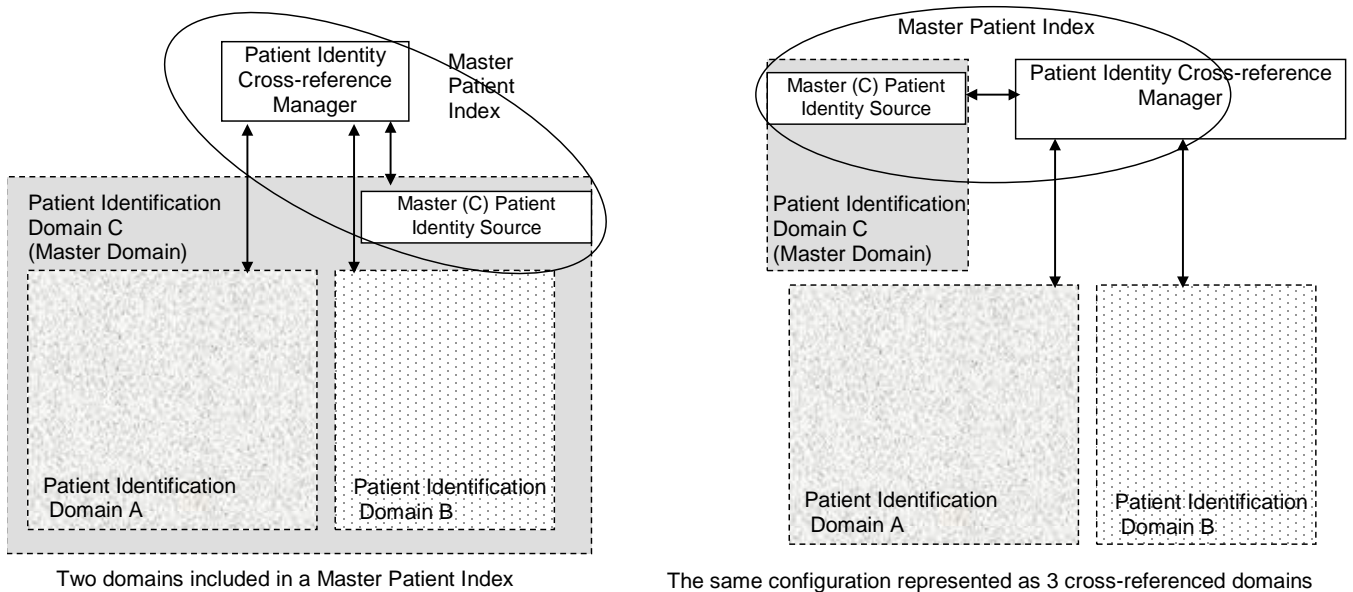
A better outcome could have happened if the clinic had recorded any one of several different data elements that would have helped tie this new data to the previous data. Any one of the

1560 Mothers Maiden Name (even the Mother’s First Name component), the Home Phone Number, or
 1565 the unique identifier for the kids which was printed on the “yellow card” from Pediatric
 Healthcare would have helped.

5.4 Relationship between the PIX Integration Profile and eMPI

1565 The PIX Integration Profile achieves the integration of disparate Patient Identifier Domains by
 using a cross-referencing approach between Patient Identifiers associated with the same patient.
 This section discusses how this approach is compatible with environments that wish to establish
 master patient identifiers (MPI), or enterprise MPI (eMPI) systems. An eMPI may be considered
 a particular variation in implementation of the PIX Integration Profile.

1570 The concept of an MPI is a rather broad concept, yet it is most often associated with the creation
 of a master patient identifier domain. Such a master domain is considered more broadly
 applicable or more “enterprise-level” than the other patient identifier domains it includes. Such a
 hierarchical inclusion of patient identification domains into a “master patient identification
 domain” can be considered a particular case of patient cross-reference, where the patient
 identifiers in the various domains are cross-referenced to the patient identifiers of the master
 domain. Two possible configurations are depicted by Figure 5.4-1.



1575

Figure 5.4-1: PIX Profile Relationship to eMPI

1580 Figure 5.4-1 above shows how the Master Patient Identifier Domain (Domain C), in a typical
 MPI approach, is simply another patient Identification Domain when considered in a Cross-
 referencing approach. The decision to place enterprise-wide systems such as Clinical Data
 Repositories into the so-called master domain is simply a configuration choice. In addition, such
 a configuration sometimes assumes that any system in Patient Domain A not only manages the
 patient Identifiers of Domain A but is also aware of those of Domain C. In the Patient Identifier

1585 Cross-reference Integration Profile, this is a configuration choice where certain systems have been designed and configured to operate across multiple domains. Thus the entity often called an MPI (shown by the oval) is actually the combination of a Patient Identity Source Actor (ADT) along with a Patient Identifier Cross-reference Manager.

1590 The PIX Integration Profile can coexist with environments that have chosen to deploy a distinct MPI, and provides a more scalable approach. Many other configurations can also be deployed, in particular those where the creation of a master domain “including” the other domains is not necessary (i.e., a simple federation of domains where none is actually the master).

6 Patient Synchronized Applications (PSA)

1595 The *Patient Synchronized Applications Profile (PSA)* enables single patient selection for the user working in multiple applications on a workstation desktop. With this Integration Profile patient selection in any of the applications causes all other applications to tune to that same patient. This allows a clinician to use the application they are most familiar with to select the patient and have that selection reflected in the other applications they are using follow along.

1600 This profile leverages the HL7 CCOW standard, specifically for patient subject context management. The scope of this profile is for sharing of the CCOW Patient subject only. The IHE PSA Profile adds value to the CCOW specification for the patient subject by further constraining the patient identifier to ensure consistency across applications supporting PSA, providing guidance for consistent behavior across applications supporting PSA and ensuring consistent interaction with the Patient Identifier Cross-reference Consumer Actor across the enterprise.

1605 For applications that require user authentication, IHE recommends implementation of the Enterprise User Authentication Profile, as opposed to other means, such as a CCOW Authentication Repository. ITI TF-1:4 describes the Enterprise User Authentication Profile and the use of the CCOW user subject.

6.1 Actors/ Transactions

1610 Figure 6.1-1 shows the actors directly involved in the Patient Synchronized Applications Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other profiles are not shown.




Figure 6.1-1: Patient Synchronized Applications Profile Actor Diagram

1615 Table 6.1-1 lists the transactions for each actor directly involved in the PSA Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”).

1620 The Patient Context Participant Actor shall support all four transactions identified in Figure 6.1-1 as defined in ITI TF-2a. The Patient Context Participant Actor shall respond to all patient context changes. This actor shall set the patient context provided the application has patient selection capability.

The IHE Context Manager Actor may encompass more than a CCOW context manager function. It may include a number of other components such as the context management registry and patient mapping agent.

1625 The Context Manager Actor may be grouped with a Patient Identifier Cross-referencing (PIX) Consumer Actor of the Patient Identity Cross-referencing Profile; see ITI TF-2x: Appendix D for a description of the additional responsibilities placed on the Context Manager Actor in this case.

1630 **Table 6.1-1: Patient Synchronized Applications Integration Profile - Actors and Transactions**

Actors	Transactions	Optionality	Section
Patient Context Participant	Join Context [ITI-5]	R	ITI TF-2a: 3.5
	Change Context [ITI-6]	R	ITI TF-2a: 3.6
	Leave Context [ITI-7]	R	ITI TF-2a: 3.7
	Follow Context [ITI-13]	R	ITI TF-2a: 3.13
Context Manager	Join Context [ITI-5]	R	ITI TF-2a: 3.5
	Change Context [ITI-6]	R	ITI TF-2a: 3.6
	Leave Context [ITI-7]	R	ITI TF-2a: 3.7
	Follow Context [ITI-13]	R	ITI TF-2a: 3.13

6.2 Patient Synchronized Applications Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 6.2-1 along with the actors to which they apply. Dependencies between options, when applicable, are specified in notes.

1635

Table 6.2-1: Patient Synchronized Applications - Actors and Options

Actor	Options	Vol. & Section
Patient Context Participant	<i>No options defined</i>	--
Context Manager	<i>No options defined</i>	--

6.3 Patient Synchronized Applications Integration Profile Process Flows

1640 The Patient Synchronized Applications Integration Profile provides maximum value when a user needs to use more than one application simultaneously. The process flow outlined in ITI TF-1: 6.3.1 depicts a use case where the applications only participate in the PSA Profile. The process flow outlined in ITI TF-1: Appendix E illustrates when the PSA and Enterprise User Authentication (EUA) profiles are deployed together.

6.3.1 Use Case: Simple Patient Switching

1645 When the PSA Profile is not grouped with EUA Profile only the patient identity is passed in context. This use case does not explicitly identify the method of user authentication, as it may

not be required by the application or may be accomplished by other means. In this use case both applications share the same patient identifier domain. The process flow for this use case is:

1650 The clinician launches the clinical data repository application, depicted as Patient Context Participant Actor 1. The clinical data repository application joins the context session for the clinician desktop.

The clinician selects patient A in the clinical data repository application. The clinical data repository application sets the identifier for patient A in context.

1655 The clinician launches a cardiology application, depicted as Patient Context Participant Actor 2. The Cardiology application joins the context session, gets the identifier for patient A from context, and tunes its display to patient A.

1660 The clinician selects patient B in the cardiology application. This action results in the initiation of a Change Context transaction by the cardiology application (Patient Context Participant Actor 2). All non-instigating applications participate via the Follow Context transaction, which results in the selected patient being displayed in the clinical data repository application (Patient Context Participant Actor 1).

The clinician closes the clinical data repository application. The clinical data repository application leaves the context prior to terminating the application.

1665 The clinician closes the cardiology application. The cardiology application leaves the context prior to terminating the application.

Figure 6.3-1 illustrates the process flow for this use case.

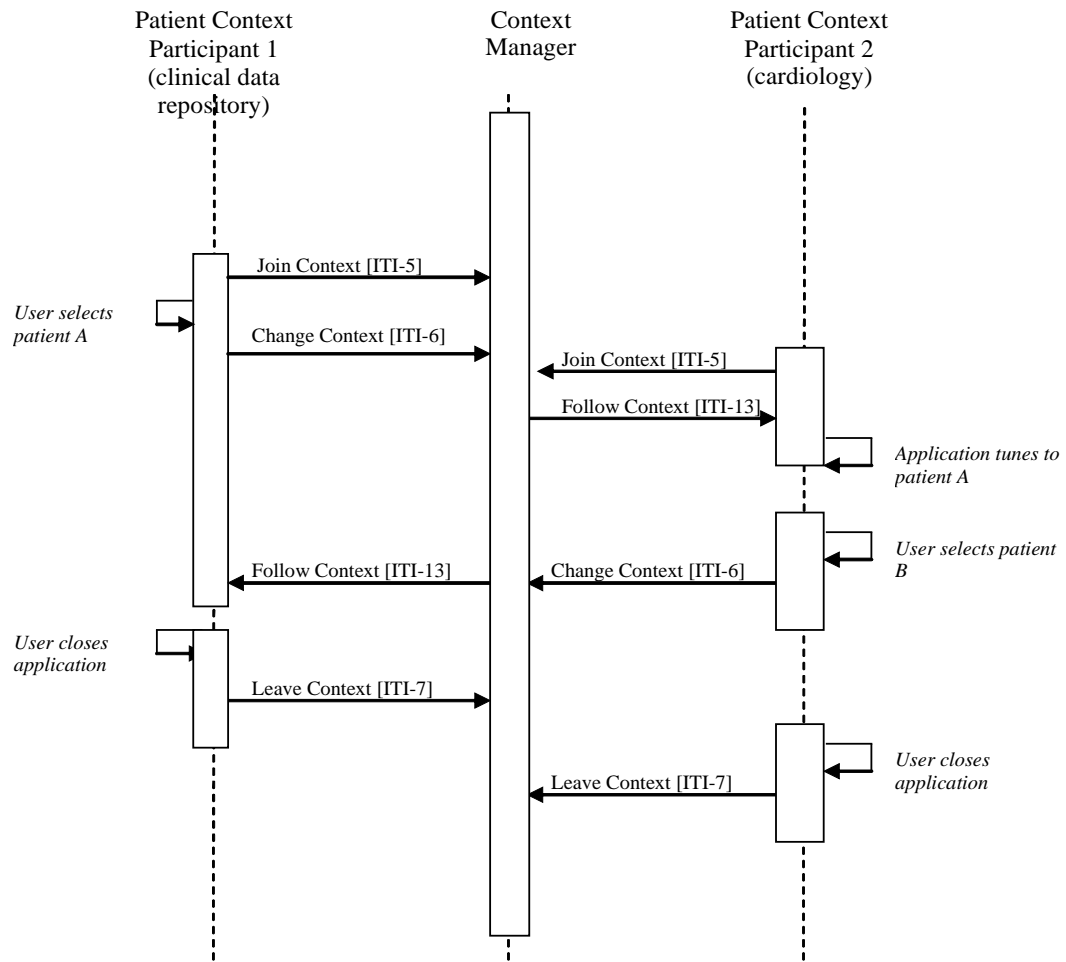


Figure 6.3-1: Simple Patient Switching Process Flow

1670 **7 Consistent Time (CT)**

The *Consistent Time Integration Profile (CT)* provides a means to ensure that the system clocks and time stamps of the many computers in a network are well synchronized. This profile specifies synchronization with a median error less than 1 second. This is sufficient for most purposes.

1675 The Consistent Time Integration Profile defines mechanisms to synchronize the time base between multiple actors and computers. Various infrastructure, security, and acquisition profiles require use of a consistent time base on multiple computers. The Consistent Time Profile requires the use of the Network Time Protocol (NTP) defined in RFC 1305. When the Time Server is grouped with a Time Client to obtain time from a higher tier Time Server, the Time Client shall utilize NTP. For some Time Clients that are not grouped with a Time Server, SNTP may be usable.


1680

This profile was previously a portion of the Radiology Basic Security Profile, but it has a variety of other infrastructure uses.

1685 **Note:** This profile corresponds to a portion of the IHE Radiology Technical Framework, Basic Security Profile. It is required by more than just radiology systems. It is needed by several of the profiles in the IHE IT Infrastructure and will also be needed by Cardiology. It is therefore being re-located from IHE Radiology into IHE IT Infrastructure. There are no changes to the requirements, so actors that supported the Radiology Basic Secure Node or Time Server do not need modification. The Maintain Time [RAD TF-3: 4.33] transaction from Radiology and the Maintain Time [ITI TF-2a: 3.1] transaction for IT Infrastructure are the same.

1690 **7.1 Actors/ Transactions**

Figure 7.1-1 shows the actors directly involved in the Consistent Time Profile and the relevant transactions between them. Other actors that may be indirectly involved because of their participation in profiles that require consistent time are not shown.



1695 **Figure 7.1-1: Consistent Time Profile Actor Diagram**

Table 7.1-1 lists the transactions for each actor directly involved in the Consistent Time Integration Profile. In order to claim support of this integration profile, an implementation must perform the required transactions (labeled “R”).

1700

Table 7.1-1: Consistent Time - Actors and Transactions

Actors	Transactions	Optionality	Section
Time Server	Maintain Time [ITI-1]	R	ITI TF-2a: 3.1
Time Client	Maintain Time [ITI-1]	R	ITI TF-2a: 3.1

7.2 Consistent Time Integration Options

Options that may be selected for this integration profile are listed in the Table 7.2-1 along with the actors to which they apply.

1705

Table 7.2-1: Consistent Time - Actors and Options

Actor	Options	Vol. & Section
Time Server	<i>Secured NTP</i>	ITI TF-2a: 3.1.4-1
Time Client	<i>SNTP, Secured NTP</i>	ITI TF-2a: 3.1.4-1

7.3 Consistent Time Process Flow

This section describes the typical flow related to the Consistent Time Profile. In the process flow Figure 7.3-1, the Time Client B and Time Server B have been grouped. When a Client and Server are grouped they utilize internal communications mechanisms to synchronize their time.

1710

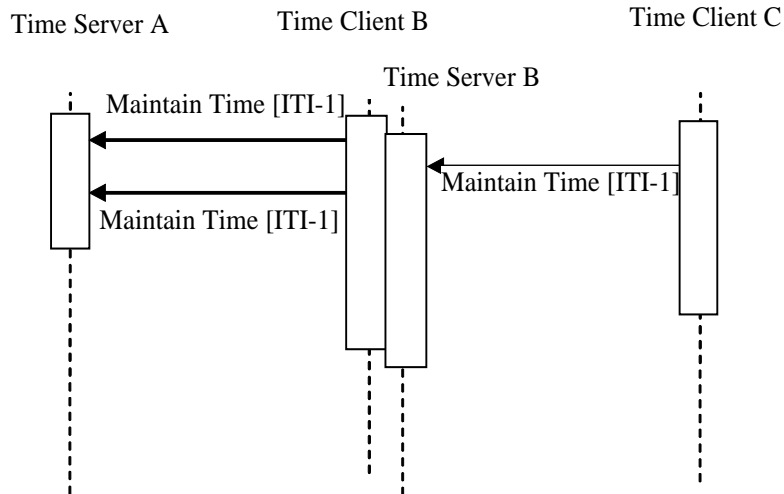


Figure 7.3-1: Basic Process Flow in Consistent Time Profile

1715 The Time Client B maintains time synchronization with the Time Server A. The Time Server B is internally synchronized with Time Client B. The Time Client C maintains time synchronization with Time Server B.

The NTP protocol has been designed to provide network time services for synchronization with this kind of cascaded synchronization. The achievable accuracy is dependent on specific details

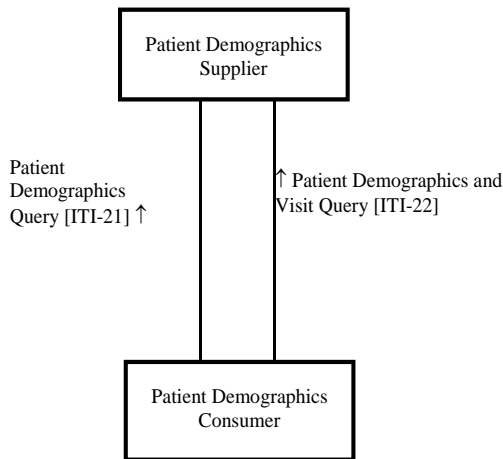
1720 of network hardware and topology, and on details of computer hardware and software implementation. The Time Server and Time Client are grouped to provide synchronization cascading and reduce network traffic.

8 Patient Demographics Query (PDQ)

1725 The Patient Demographics Query Integration Profile (PDQ) provides ways for multiple distributed applications to query a patient information server for a list of patients, based on user-defined search criteria, and retrieve a patient’s demographic (and, optionally, visit or visit-related) information directly into the application.

8.1 Actors/ Transactions

1730 Figure 8.1-1 shows the actors directly involved in the Patient Demographics Query Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in Patient ID Cross-referencing, etc. are not necessarily shown.



1735 **Figure 8.1-1: Patient Demographics Query Profile Actor Diagram**

Table 8.1-1 lists the transactions for each actor directly involved in the Patient Demographics Query Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 8.2.

1740

Table 8.1-1: Patient Demographics Query Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Patient Demographics Consumer	Patient Demographics Query [ITI-	R	ITI TF-2a: 3.21
	Patient Demographics and Visit Query [ITI-22]	O	ITI TF-2a: 3.22
Patient Demographics Supplier	Patient Demographics Query [ITI-	R	ITI TF-2a: 3.21
	Patient Demographics and Visit Query [ITI-22]	O	ITI TF-2a: 3.22

8.2 Patient Demographics Query Integration Profile Options

1745 Options that may be selected for this Integration Profile are listed in Table 8.2-1 along with the actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 8.2-1: Patient Demographics Query - Actors and Options

Actor	Options	Vol. & Section
Patient Demographics Consumer	<i>Patient Demographics and Visit Query</i>	ITI TF-2a: 3.22
	<i>Pediatric Demographics</i>	ITI TF-1: 8.2.2
Patient Demographics Supplier	<i>Patient Demographics and Visit Query</i>	ITI TF-2a: 3.22
	<i>Pediatric Demographics</i>	ITI TF-1: 8.2.2

1750

8.2.2 Pediatric Demographics

The experience of immunization registries and other public health population databases has shown that retrieving patient records for an individual person in environments with large proportions of pediatric records requires additional demographic data.

1755 Information about the mother of the patient or a household telephone number is helpful in retrieving records in large population databases where data quality may be uneven.

Certain other demographics fields are important to include in the query response as they may be used by the Patient Demographics Consumer in verifying the identity of the patient, in particular, they aid in distinguishing records for twins, triplets, and so forth.

1760 Pediatric Demographics makes use of the following six additional demographic fields to aid record matching in databases with many pediatric records.

Field	Reason for inclusion	Value
Mother's Maiden Name	Any information about the mother is helpful in making a match	Helps create true positive matches
Patient Home Telephone	A telecom helps match into the right household	Helps create true positive matches
Patient Multiple Birth Indicator	Indicates this person is a multiple - twin, triplet, etc.	Helps avoid false positive matches of multiples
Patient Birth Order	Distinguishes among those multiples.	Helps avoid false positive matches of multiples

Field	Reason for inclusion	Value
Last Update Date/Time, Last Update Facility	These fields, although not strictly demographic, can effectively substitute when multiple birth indicator and birth order are not collected. They indirectly provide visit information. Provider visits on the same day may likely indicate two children brought to a doctor together.	Helps avoid false positive matches of multiples

1765 Patient Demographics Consumer actors which support the Pediatrics Demographics Option will be able to provide Pediatric Demographics query parameter fields in the Patient Demographics Query transaction [ITI-21], and shall be able to receive and process any values returned for the fields identified as Pediatric Demographics.

1770 Patient Demographics Supplier actors which support the Pediatrics Demographics Option will be able to match on values provided for any Pediatric Demographics fields in the Patient Demographics Query transaction [ITI-21] and shall return values, when available, for the fields identified as Pediatric Demographics.

Pediatric Demographics query parameter fields are:

- Mother's Maiden Name
- Patient Home Telephone

Pediatric Demographics are defined as all of the following:

- 1775
- Mother's Maiden Name
 - Patient Home Telephone
 - Patient Multiple Birth Indicator
 - Patient Birth Order
 - Last Update Date/Time
- 1780
- Last Update Facility

8.3 Patient Demographics Query Process Flow

The Patient Demographics Supplier performs the following functions.

- 1785
- It receives patient registration and update messages from other systems in the enterprise (*e.g.*, ADT Patient Registration systems), which may or may not represent different Patient ID Domains. The method in which the Patient Demographics Supplier obtains the updated patient demographic information is not addressed by this profile.
 - It responds to queries for information.

1790 Specific methods for acquiring demographic information are beyond the scope of this Profile. It is a prerequisite that the Patient Demographics Supplier possess current demographic information. One method by which current demographic information may be obtained is for the

Patient Demographic Supplier to be grouped with another IHE actor, such as Order Filler, that either maintains or receives such information.

1795 In all cases, the Patient Demographics Supplier receives a Patient Demographics Query or Patient Demographics and Visit Query request from the Patient Demographics Consumer, and returns demographics (and, where appropriate, visit) information from the single domain that is associated with the application to which the query message is sent. Identifier information may be returned from multiple or single domains; see the “Using Patient Data Query (PDQ) in a Multi-Domain Environment” section (ITI TF-2x: Appendix M) for a discussion of the architectural issues involved.

1800 **Use Case 1: Patient Information Entering at Bedside**

An admitted patient is assigned to a bed. The patient may or may not be able to provide positive ID information. The nurse needs to enter patient identity information into some bedside equipment to establish the relationship of the assigned bed to the patient. The equipment issues a query for a patient pick list to a patient demographics supplier that provides data for a patient pick list. Search criteria entered by the nurse might include one or more of the following:

- Partial or complete patient name (printed on the patient record or told by the patient)
- Patient ID (this may be obtained from printed barcode, a bed-side chart, etc.)
- Partial ID entry or scan.
- Date of birth / age range
- Bed ID

1810

The system returns a list of patients showing the MRN, full name, age, sex, room/bed, and admit date, and displays the list to the nurse. The nurse then selects the appropriate record to enter the patient identity information into the bedside equipment application.

Use Case 2: Patient Identity Information Entering in Physician Offices

1815 *A patient visits a physician office for the first time. The nurse needs to register the patient; in doing so, it is desired to record the patient’s demographic data in the practice management information system (PMIS). The physician office is connected to a hospital enterprise’s central patient registry. The nurse issues a patient query request to the central patient registry, with some basic patient demographics data as search criteria. In the returned patient list, she picks up an appropriate record for the patient, including the hospital’s patient ID, to enter into the PMIS. (Note that the PMIS uses a different Patient ID domain than that of the central patient registry.)*

1820

1825 *The PMIS uses its own patient identifier, coordinating this identifier with the patient identifier returned in the pick list (sharing the hospital’s Patient ID Domain) to retrieve information from the hospital’s clinical repository.*

Use Case 3: Patient Demographics Query in an Enterprise with Multiple Patient ID Domains

A lab technician enters some basic demographics data (e.g., patient name) into a lab application to query a patient demographics supplier to identify a patient for his lab exams. As the

1830

application also needs the patient identifier in another Patient ID Domain in the enterprise for results delivery, the application is configured to receive patient IDs from other domains in the query response.

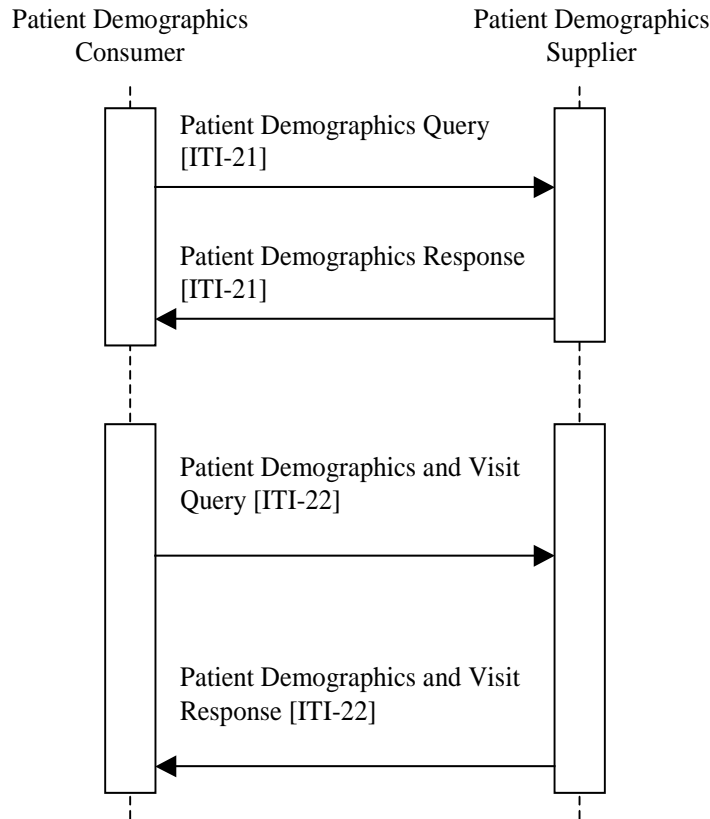


Figure 8.2-1: Basic Process Flow in Patient Demographics Query Profile

1835

8.3.1 Combined Use of PDQ with Other IHE Workflow Profiles

When the Patient Demographics Supplier Actor is grouped with actors in other IHE profiles that perform patient information reconciliation activities (e.g., Radiology PIR), the PDQ Supplier Actor may use the updated information to respond to PDQ Queries. In addition, the Patient Demographics Query Profile may play an integral workflow role in conjunction with other IHE Profiles.

1840

8.3.2 Supplier Data Configuration

A Patient Demographics Supplier Actor that holds demographic information for a single Patient ID domain shall provide matches in that domain.

1845

In the case where the Patient Demographics Supplier Actor holds demographic information for multiple Patient ID domains, the Patient Demographics Supplier Actor shall return information for the domain associated with *MSH-5-Receiving Application* and *MSH-6-Receiving Facility*. See the “Using Patient Data Query (PDQ) in a Multi-Domain Environment” section (ITI TF-2x:

Appendix M) for a further discussion of this case and an illustration of the supporting architecture.

1850

9 Audit Trail and Node Authentication (ATNA)

1855 The Audit Trail and Node Authentication (ATNA) Integration Profile establishes security measures which, together with the Security Policy and Procedures, provide patient information confidentiality, data integrity and user accountability. This environment is considered the Security Domain and can scale from a department, to enterprise or XDS Affinity Domain. The ATNA model considers that within the secure domain the following is true:

- 1860 1. All machines are host authenticated. (There are various means of accomplishing this.) This authentication identifies the machine as being one that is known to the security system of the hospital, with known security characteristics. Unknown machines might be granted access, but with the caveat that they are only granted access to information that is authorized for disclosure to the public or to unknown machines. (A patient might choose to allow information such as appointment schedules to be at risk of machine disclosure by unknown machines while not allowing more sensitive PHI to be disclosed.)
- 1865 2. The host identification is used to determine what (if any) access should be granted to automated processes on that host, and/or persons under the direction of that host's access controls. In practice the automated processes play a critical role, managing issues like pre-fetching, thus person authentication/identification is not sufficient.
- 1870 3. The secure node is responsible for providing reasonable access controls. This typically includes user authentication and authorization. The value of this user authentication needs to be balanced against the possible safety and patient health impacts of delaying delivery of care by the additional authentication steps.
4. The secure node is also responsible for providing security audit logging to track security events. In healthcare this audit log is often more useful than strict access controls and should be relied upon even in emergencies.

1875 This model is partially driven by the underlying assumption that there will be situations where documents are being exchanged between machines and stored on the recipient. This is partly driven by the need for healthcare systems to operate in disasters and overload situations, where the network operation is limited or destroyed. It is not safe to assume that clients are display only. So there will be semi-permanent copies of most information kept. Even in normal operation, healthcare providers may have only 15 minutes per patient. Good healthcare system design recognizes the need to not waste any of those seconds searching and transferring documents over a network. The documents are transferred in advance, and are kept locally until it is determined that they are no longer needed. There are thin client display only applications in healthcare, but they are limited to uses that can fail without introducing risks to safety or patient health, but a complete security/privacy design requires handling situations where data is stored after retrieval.

ATNA Governance Assumptions

The underlying assumptions are:

- 1890 • All systems that are members of the secure domain implement a Secure Node Actor for the ATNA Profile. The ATNA Profile defines transactions between the secure nodes to create a secure domain that is under the management of a domain security officer.
- 1895 • All applications on a secure node will comply with ATNA requirements, regardless of whether they are IHE Actors or not. They apply to all IT assisted activities that directly create, access, update, and delete PHI, not only those specified by IHE and performed by IHE actors.
- 1900 • IHE addresses only those security requirements related to systems within the scope of IHE healthcare applications. It does not address other security requirements such as defending against network attacks, virus infection, etc. The principal objective of the Audit Trail mechanism is to track data access to PHI, not IHE transactions.
- 1905 • Mobile equipment can participate in the Audit Trail and Node Authentication Integration Profile, but special issues related to mobile equipment are not explicitly addressed in this profile.
- ATNA assumes that physical access control, personnel policies and other organizational security considerations necessary to make an enterprise compliant with security and privacy regulations are in place.

9.1 Authentication

1910 ATNA contributes to access control by limiting network access between nodes and limiting access to each node to authorized users. Network communications between secure nodes in a secure domain are restricted to only other secure nodes in that domain. Secure nodes limit access to authorized users as specified by the local authentication and access control policy.

9.1.1 User Authentication

1915 The Audit Trail and Node Authentication Integration Profile requires only local user authentication. The profile allows each secure node to use the access control technology of its choice to authenticate users. The use of Enterprise User Authentication is one such choice, but it is not necessary to use this profile.

9.1.2 Connection Authentication

1920 The Audit Trail and Node Authentication Integration Profile requires the use of bi-directional certificate-based node authentication for connections to and from each node. The DICOM, HL7, and HTML protocols all have certificate-based authentication mechanisms defined. These authenticate the nodes, rather than the user. Connections to these machines that are not bi-directionally node-authenticated shall either be prohibited, or be designed and verified to prevent access to PHI.

1925 **Note:** Communications protocols that are not specified by IHE profiles, e.g., SQL Server, must be bi-directionally authenticated if they will be used for PHI. This profile does not specify how that authentication is to be performed.

1930 This requirement can also be met by ensuring complete physical network security with strict configuration management. This means that no untrusted machine can obtain physical access to any portion of the network. Making the connection authentication configurable enhances performance in physically secured networks. A Secure Node Actor shall be configurable to support both connection authentication and physically secured networks.

1935 IHE does not mandate the use of encryption during transmission. Most hospital networks provide adequate security through physical and procedural mechanisms. The additional performance penalty for encryption is generally not justified for these networks. This profile mandates the use of the TLS security negotiation mechanism for all communications between secure nodes as a means of ensuring that they only communicate with other authorized secure nodes. It permits the negotiation of encryption if both nodes are configured to request and support encryption. This allows installation of IHE secure nodes into environments where the network is not otherwise secured.

9.2 Audit Trails

1940 User Accountability is provided through Audit Trail. The Audit Trail needs to allow a security officer in an institution to audit activities, to assess compliance with a secure domain's policies, to detect instances of non-compliant behavior, and to facilitate detection of improper creation, access, modification and deletion of Protected Health Information (PHI). PHI is considered to be the patient-identifiable information records (e.g., Registration, Order, Study/Procedure, Reports, Images, and Presentation States). PHI may be accessed by users or exchanged between the systems. This includes information exported to and imported from every secured node in the *secure domain*.

1950 The user accountability is further enhanced through a standards based Centralized Audit Record Repository, that provides a central Audit Record repository as the simplest means to implement security requirements. A transfer of Audit Records from all the IHE actors to the Audit Record Repository reduces the opportunities for tampering and makes it easier to audit the department. Disconnected nodes may store audit data for transfer to the Audit Repository upon reconnection to the secure domain network.

The audit trail contains information so that questions can be answered such as:

- 1955
- For some user: which patients' PHI was accessed?
 - For some patient PHI: which users accessed it?
 - What user authentication failures were reported?
 - What node authentication failures were reported?

1960 The Audit Trail and Node Authentication Profile provides tools that are useful for enterprises attempting to become compliant with privacy and security regulations (HIPAA, European, Japanese, etc.), but the profile does not itself make the enterprise compliant. For guidance on proper audit log management enterprises should look to documents such as NIST SP 800-92 – Guide to Computer Security Log Management.

9.2.1 Audit Messages

1965 The use of auditing as part of a security and privacy process is appropriate for situations where the people involved are generally trustworthy and need a wide range of flexibility to respond rapidly to changing situations. This is the typical healthcare provider environment. Auditing tracks what takes place, and the people involved know that their actions are being audited. This means that the audit records must capture event descriptions for the entire process, not just for individual components that correspond to individual IHE actors.

1970

The IHE audit trail is the first of several profiles that correspond to different forms of access control and authentication. Auditing is always needed independent of the access control and authentication method chosen.

The IHE-specified audit flow is illustrated in Figure 9.2-1.

1975

1. Real world activities take place, and some of these activities involve the applications processing of a device that includes support for some IHE profiles. This product has components that may correspond to specific IHE Actors. The product may also have other capabilities that are independent of IHE recommendations.

1980

2. A wide variety of events take place during this process. Some of these events are directly related to IHE Actor activities. Others may be indirectly related, and still others are not related to any IHE specification. The events are both extremely detailed minor events, such as keystrokes, and high level events such as analyzing a diagnostic study. Very few of these events are relevant to security and privacy auditing. Most are too low level to be useful or are otherwise irrelevant.

1985

3. The “Security Audit and Access Accountability Message XML Data Definitions for Healthcare Applications” (RFC-3881) defines an XML schema for reporting events that are relevant to security and privacy auditing. It was defined in cooperation with the ASTM, HL7, and DICOM standards organizations and the NEMA/COCIR/JIRA Security and Privacy Committee. The IHE recommends the use of the RFC-3881 format, and recommends reporting only events that it can describe.

1990

4. DICOM has standardized some of the audit message vocabulary. The DICOM Audit Message Vocabulary extends the basic vocabulary provided with RFC-3881, and also further specifies some optional elements in RFC-3881. An example of vocabulary extension is the addition of a coded value to indicate that a field contains a DICOM Study Instance UID. An example of optional element specification is the requirement that the UserID field in RFC-3881 messages shall be the user ID used by the local device operating system, and that the AlternateID shall be the user ID used by the enterprise authentication system (if it is different).

1995

5. This profile defines other events that do not correspond to events defined in the DICOM vocabulary. These events are describable by RFC-3881, and this profile includes requirements for such descriptions.

2000

6. IHE auditing specifies that when using the RFC-3881, events that can be described using the DICOM vocabulary they shall be reported using the DICOM vocabulary, even if the device is not otherwise a DICOM compliant device. Events that do not match the

- 2005 DICOM vocabulary shall be reported using RFC-3881 vocabulary or other extensions. Events that cannot be reported using RFC-3881 are not candidates for reporting.
7. The local site will then apply its own reporting policies. The IHE profile specifies the capabilities that should be present for audit reporting, and also that there should be controls present to allow the local site security administration to control reporting detail.
- 2010 The IHE profile does not specify any audit reporting functions or formats.
8. IHE specifies events that must be reported in the audit trail. There are other events related to security, which may be reported in the audit trail or by other means. This profile does not describe them and does not require that they use this reporting format or mechanism. Examples of such events are OS login, network routing and firewall logs.

2015




Figure 9.2-1: Flow of Events into Audit Messages

9.2.2 Backwards Compatibility

2020 This profile also defines the continued use of messages that are formatted in accordance with the IHE Provisional Audit Message format from the deprecated Basic Security Profile in IHE Radiology Technical Framework. This older format describes events that are suitable for

reporting in Radiology and other diagnostic and treatment activities. These events are a subset of the kind of events that can be described using RFC-3881 and the DICOM vocabulary.

2025 The IHE ATNA Profile also allows for the reporting of these events using the Provisional format over either of the IHE specified transport mechanisms. The intention is that products will gradually transition from the Provisional message format to RFC-3881 format, but it is recognized that this transition will take time and that there is a significant installed base.

2030 The Provisional format is unlikely to be of interest to other healthcare applications, which should use the RFC-3881 format and DICOM Vocabulary where appropriate.

9.3 Audit Trail Transport

The Audit Trail and Node Authentication Integration Profile specifies the use of Syslog Protocols as the mechanism for logging audit record messages to the audit record repository.

There are two recognized transports specified:

2035 1) Transmission of Syslog Messages over UDP (RFC5426) with The Syslog Protocol (RFC5424) which formalizes and obsoletes Syslog (RFC-3164). There are, several known limitations that seem significant but have not been a problem in practice:

- There is no confirmation to the sender that the audit record message was received at the destination
- 2040 • There is no option to encrypt the audit record messages
- Authentication by means of certificates of the sending nodes and the central audit repository is not possible
- Messages may be truncated or lost.

2045 2) Transmission of Syslog Messages over TLS (RFC5425) with The Syslog Protocol (RFC5424) which formalizes sending syslog messages over a streaming protocol protectable by TLS.

9.4 Actors/Transactions

2050 Table 9.4-1 lists the transactions for each actor directly involved in the Audit Trail and Node Authentication Integration Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile that implementations may choose to support is listed in ITI TF-1: 9.4. Their relationship is shown in Figure 9.4-1.

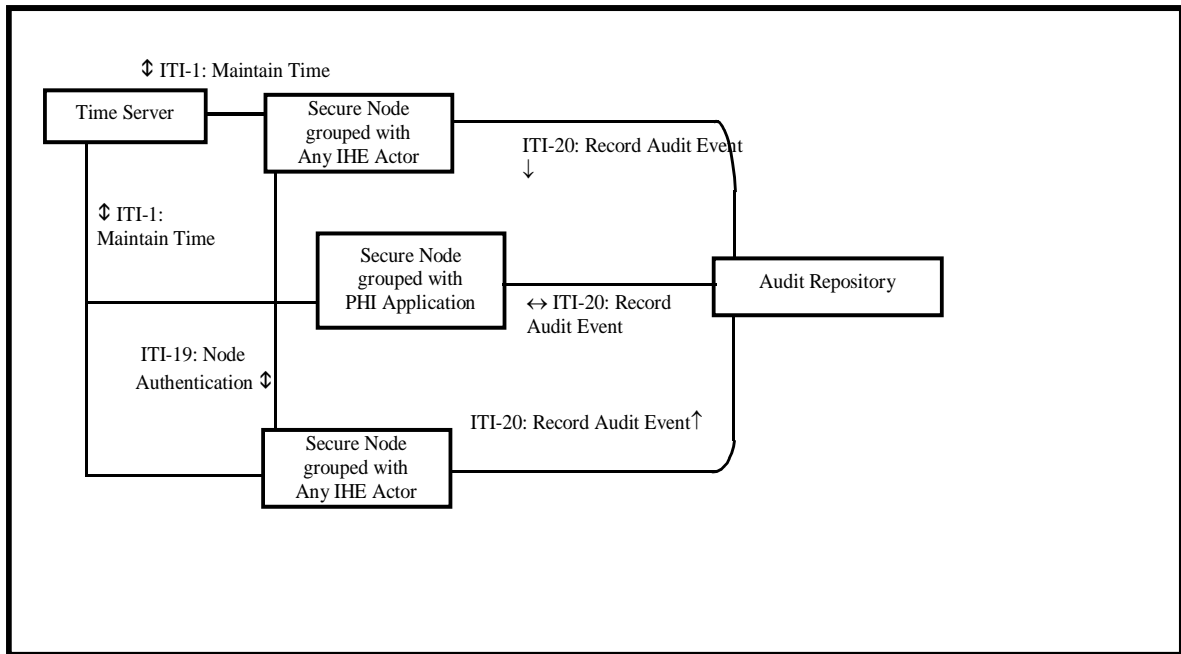


Figure 9.4-1: Audit Trail and Node Authentication Diagram

2055 When an implementation chooses to support this Integration Profile for an actor, that actor shall be grouped with the Secure Node Actor. It is required that all IHE actors and any other activities in this implementation support the Audit Trail and Node Authentication Integration Profile.

A means must be provided to upload the required certificates to the implementation, e.g., via floppy disk or file transfer via network.

2060 Non-IHE applications that process PHI shall detect and report auditable events, and protect access.

Table 9.4-1: Audit Trail and Node Authentication Integration Profile - Actors and Transactions

Actor	Transactions	Optionality	Section
Audit Record Repository	Record Audit Event [ITI-20]	R	ITI-2a: 3.20
Secure Node	Authenticate Node [ITI-19]	R	ITI-2a: 3.19
	Record Audit Event {ITI-20}	R	ITI-2a:3.20
	Maintain Time [ITI-1]	R	ITI-2a: 3.1
Secure Application	Authenticate Node [ITI-19]	O	ITI-2a: 3.19
	Maintain Time [ITI-1]	O	ITI-2a: 3.1
	Record Audit Event [ITI-20]	O	ITI-2a: 3.20

2065 The Secure Node Actor shall include:

1. The Authenticate Node transaction for all network connections that may expose private information as specified in ITI TF-2a: 3.19.
2. All local user activity (login, logout, etc.) protected to ensure only authorized users.
3. The Record Audit Event as specified in ITI TF-2a: 3.20

2070 The difference between the Secure Node and the Secure Application is the extent to which the underlying operating system and other environment are secured. A Secure Node includes all aspects of user authentication, file system protections, and operating environment security. The Secure Application is a product that does not include the operating environment. The Secure Application provides security features only for the application features. See ITI TF-1: 9.7 for the relationships among a Secure Node, Secure Application, and other actors.

2075

1. The Audit Repository shall support:
2. Both audit transport mechanisms.
3. Any IHE-specified audit message format, when sent over one of those transport mechanisms. Note that new applications domains may have their own extended vocabularies in addition to the DICOM and IHE vocabularies. This also means that an ATNA Audit Repository is also automatically a Radiology Basic Security Profile Audit Repository because it must support the IHE Provisional Message format and it must support the BSD syslog protocol.
4. Self protections and user access controls.

2080

2085 This profile does not specify other functions for the Audit Repository, but it is expected that most repositories will perform screening, reporting, archival, etc.

9.5 ATNA Integration Profile Options

Options that may be selected for this Integration Profile are listed in the Table 9.5-1 along with the actors to which they apply. Dependencies between options when applicable are specified in notes.

2090

Table 9.5-1: ATNA - Actors and Options

Actor	Options	Vol. & Section
Audit Record Repository	<i>None</i>	-
Secure Node	<i>Radiology Audit Trail</i>	RAD TF-1: 2.2.1; RAD TF-3: 5.1
Secure Application	<i>Radiology Audit Trail</i>	RAD TF-1: 2.2.1; RAD TF-3:5.1

9.5.1 ATNA Encryption Option (retired)

2095 The ATNA Encryption Option is now retired as the ITI-19: Node Authentication transaction requires support for Encryption.

9.5.2 Radiology Audit Trail Option

2100 The Radiology Audit Trail provides specific requirements as to which audit events Actors in IHE Radiology domain Profiles are required to send. It also details the specific format of certain audit events based on the Radiology Actor.

9.6 Audit Trail and Node Authentication Process Flow

2105 The security measures in the Audit Trail and Node Authentication Integration Profile are user authentication, node authentication, and generation of audit records. Node authentication and user authentication define a number of transactions that establish the concept of a Secure Node and a collection of connected Secure Nodes in a secure domain. Generation of audit records requires a set of audit trigger events and a definition of the content of the audit records. This profile specifies two acceptable message formats:

- 2110 1. Messages formatted in accordance with the IHE Audit Message format. This is a combination of the DICOM Audit Messages format and IHE extensions. The IHE extensions to RFC-3881 add event codes and information needed for uses that are not within the domain of the DICOM Standard.
- 2115 2. The predecessor IHE Provisional Audit Message format. This format was defined as an interim format while the standards work to define the Common Audit Message format and vocabularies progressed through the standards organizations.

Based on the work done in ASTM (E2147-01 Standard Specification for Audit and Disclosure Logs for Use in Health Information Systems) and HL7 (Framework for Audit Messages), IHE defined a detailed set of audit trigger events, a set of general audit messages with the content for the audit record, and a mapping for each event to a general audit message. The content of the audit record has been specified by means of an XML Schema (see ITI TF- 2x: Appendix F).

2120 In the following paragraphs three typical process flows are described for situations in which authorized users, unauthorized users, and unauthorized nodes attempt to gain access to protected health information (PHI).

9.6.1 Normal Node Process Flow

2125 The following scenario shows how the IHE security measures operate for authorized access to PHI from an authorized node in the network:

- 2130 1. Time synchronization occurs independently. These transactions may take place at any time. Correct time is needed to generate Audit Records with a correct timestamp.
2. A user logs on to Image Display/Secure Node Actor.
The user enters valid credentials and is authorized to access the node.
3. The node generates audit records.

4. The user wants to query/retrieve and view some images.
Before image transactions can take place, an authentication process between the Image Display/Secure Node Actor and the Image Manager/Image Archive/Secure Node Actor takes place.
- 2135 5. Following node authentication, the node initiates the query/retrieve transactions.
6. The node generates audit records.




Figure 9.6-1: Authorized Node Process Flow

9.6.2 Unauthorized Node Process Flow

The following scenario shows how the IHE security measures help to prevent unauthorized access to PHI from an unauthorized node in the network:

- 2145 1. An unauthorized node tries to query the Lab Automation Manager/Secure Node Actor for information. This fails because no authentication has taken place, and an audit record is generated.
2. The unauthorized node tries an authentication process with the Lab Automation Manager/Secure Node. This fails because the Lab Automation Manager/Secure Node will not trust the certificate presented by the Malicious Node, and an audit record is generated.

2150 Note that the sequencing of the transactions is just one example; transactions from an unauthorized node are totally unpredictable and may happen in any order.

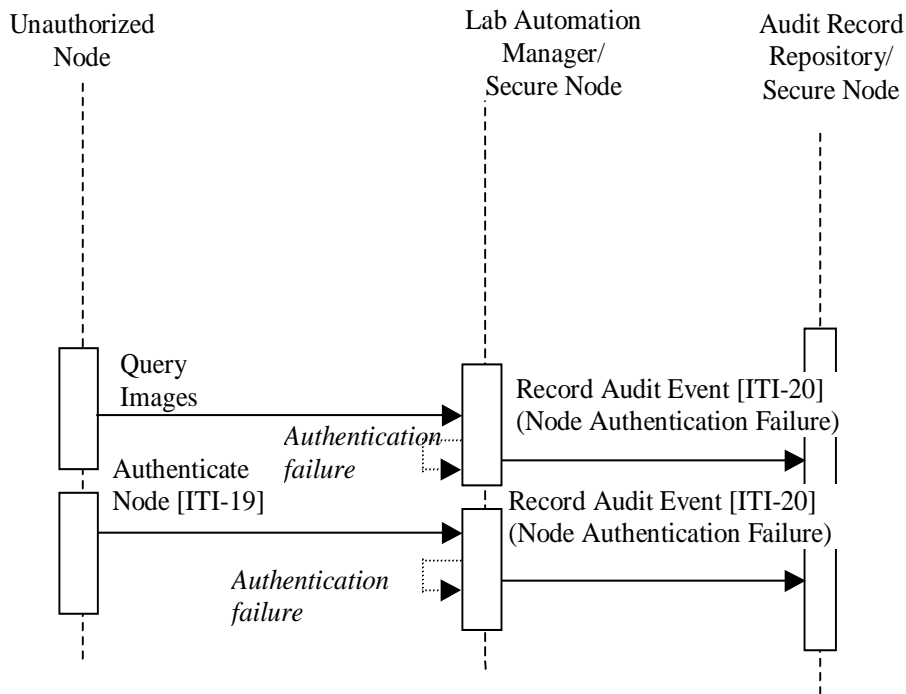


Figure 9.6-2: Unauthorized Node Process Flow

2155

9.6.3 Unauthorized User Process Flow

2160 The following scenario shows how the IHE security measures help to prevent unauthorized access to PHI from an unauthorized user in the healthcare enterprise:

1. An unauthorized user tries an authentication process with the ECG Display/Secure Node Actor. This fails because the ECG Display/Secure Node Actor detects that the user name and credentials presented are not valid at this secure node, and an audit record is generated.

2165

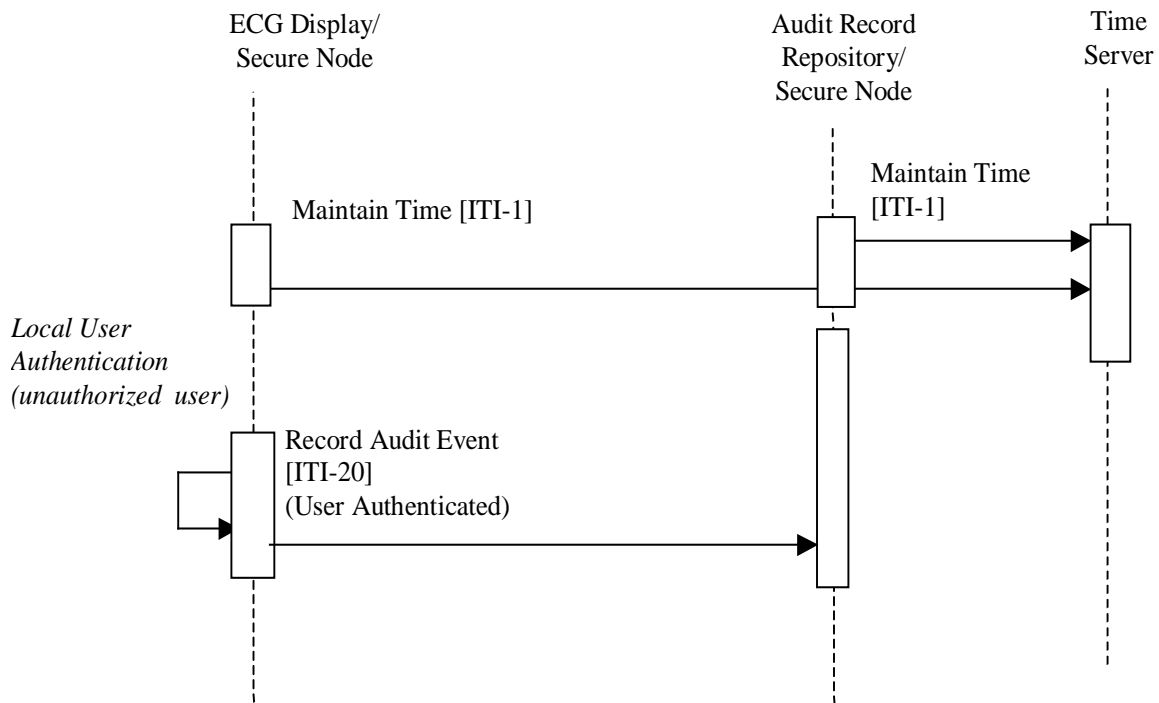


Figure 9.6-3: Unauthorized User Process Flow

9.7 Relationship between Secure Node, Secure Application, and other Actors

2170 The allocation of responsibilities when an actor is grouped with a secure node can be complex when different parties are responsible for different parts of the system. This situation arises frequently in situations like web server applications, where there is an operating system, a web server framework, and individual web applications. These might all be from different vendors. Each of these components has a role in performing security related tasks. There is also a system

2175 integrator who is responsible for assembling these components into the final complete system. It is the responsibility of the system integrator to insure that all of the necessary security functions are implemented by the appropriate system component.

Note: The system integrator might be a product vendor, outside consultant or internal staff. IHE does not specify business relationships. The term is used here to indicate a functional role, not a business relationship.

2180 IHE has split these into two primary categories:

- The healthcare functions. These are identified as IHE actors. IHE does not specify how functional actors are implemented. Multiple actors might be implemented by one web application, and it may take multiple web applications to implement one IHE actor. IHE allocates functions to the actors and it is the implementer's task to allocate these to individual web applications.
- The underlying operating environmental components. The IHE identifies these as the Secure Node Actor. It is the system integrator that determines how the functions of the Secure Node Actor are allocated to individual components.

2190 When a product claims support for the Secure Application Actor, it is claiming that it performs those functions that are appropriate to its IHE task. This will certainly include some audit responsibilities, will probably include some communications security responsibilities, and may include other security responsibilities. The specifics of these responsibilities depend upon the functions and options of that product. For example, a product that includes a user login capability will generate user related audit events and perform the user authentication. In contrast, a single function web application might only generate audit messages related to its function, and will depend upon the external secure node environment for other functions.

2195 This means that product descriptions must be sufficient for the system integrator to determine whether all of the necessary security functions are present. If the single purpose web application is depending on the web server environment to provide node authentication, this must be clear to the system integrator. Not all web server environments provide that authentication, and the integrator will need to ensure that authentication is provided when needed.

2200 When describing what security features have been implemented in a product, the following rules apply:

1. If the product claims to include the Secure Node Actor, the product has been integrated so that all of the operating system and other environmental security features are present.
2. If the product claims only to include the Secure Application Actor, that indicates that only those security features that apply to the application features are provided by the product.

2210

Product selection can then use the IHE conformance claim for a summary view of the security features provided by the product. The system integrator can use this information to determine what additional products or integration work will be needed to establish the functionality provided by a Secure Node if the application products are only Secure Applications.

2215 **10 Cross-Enterprise Document Sharing (XDS.b¹)**

The *Cross-Enterprise Document Sharing* (XDS.b) IHE Integration Profile facilitates the registration, distribution and access across health enterprises of patient electronic health records.

2220 Cross-Enterprise Document Sharing is focused on providing a standards-based specification for managing the sharing of documents between any healthcare enterprise, ranging from a private physician office to a clinic to an acute care in-patient facility.

In the rest of the ITI Technical Framework the term XDS refers generically to any flavor of XDS, currently only XDS.b¹.

2225 The XDS.b IHE Integration Profile assumes that these enterprises belong to one or more XDS Affinity Domains. An XDS Affinity Domain is a group of healthcare enterprises that have agreed to work together using a common set of policies and share a common infrastructure.

Examples of XDS Affinity Domains include:

- Community of Care supported by a regional health information organization in order to serve all patients in a given region.
- Nationwide EHR
- 2230 • Specialized or Disease-oriented Care
 - Cardiology Specialists and an Acute Cardiology Center
 - Oncology network
 - Diabetes network
- Federation of enterprises
- 2235 • A regional federation made up of several local hospitals and healthcare providers
- Government sponsored facilities (e.g., VA or Military)
- Insurance Provider Supported Communities

2240 Within an XDS Affinity Domain, certain common policies and business rules must be defined. They include how patients are identified, consent is obtained, and access is controlled, as well as the format, content, structure, organization and representation of clinical information. This Integration Profile does not define specific policies and business rules; however, it has been designed to accommodate a wide range of such policies to facilitate the deployment of standards-based infrastructures for sharing patient clinical documents. This is managed through federated document repositories and a document registry to create a longitudinal record of information

2245 about a patient within a given XDS Affinity Domain. These are distinct entities with separate responsibilities:

¹ XDS.b is used because in prior versions of the Technical Framework there was an XDS.a. With TF Version 7.0 XDS.a has been deprecated.

- A document repository is responsible for storing documents in a transparent, secure, reliable and persistent manner and responding to document retrieval requests.
- A document registry is responsible for storing information about those documents so that the documents of interest for the care of a patient may be easily found, selected and retrieved irrespective of the repository where they are actually stored.

2255 The concept of a document in XDS is not limited to textual information. As XDS is document content neutral, any type of clinical information without regard to content and representation is supported. This makes the XDS IHE Integration Profile equally able to handle documents containing simple text, formatted text (e.g., HL7 CDA Release 1), images (e.g., DICOM) or structured and vocabulary coded clinical information (e.g., CDA Release 2, CCR, CEN ENV 13606, DICOM SR). In order to ensure the necessary interoperability between the document sources and the document consumers, the XDS Affinity Domain must adopt policies concerning document format, structure and content.

2260 The XDS Integration Profile is not intended to address all cross-enterprise EHR communication needs. Some scenarios may require the use of specific XDS options, the use of other IHE Integration profiles, such as Patient Identifier Cross-Referencing, Audit Trail and Node Authentication, Cross-Enterprise User Authentication, and Retrieve Information for Display. Other scenarios may be only partially supported, while still others may require future IHE Integration profiles, which will be defined by IHE as soon as the necessary base standards are available. Specifically:

1. The sharing and access to dynamic information such as allergy lists, medication lists, problem lists, etc. is addressed in part by XDS through the aggregation of information across structured documents and in part by the XDS On-Demand Documents Option that provides support for requesting most recent computer assembled content. In addition, other IHE profiles such as Retrieve Information for Display Integration Profile provide complementary means to query a-priori known systems that may provide targeted dynamic information for a patient.
2. The placing and tracking of orders (e.g., drug prescriptions, radiology orders, etc.) is not supported by XDS. This does not preclude the use of XDS to store and register orders and corresponding results when such artifacts need to be recorded in the patient's health record. However, XDS provides no facilities for tracking progress of an order through its workflow, and therefore is not intended for order management. A complementary approach to cross-enterprise order workflow (ePrescription, eReferral) may be expected as separate Integration Profiles in the future.
3. The operation of any XDS Affinity Domain will require that a proper security model be put in place. It is expected that a range of security models should be possible. Although the XDS Integration Profile is not intended to include nor require any specific security model, it is expected that XDS implementers will group XDS Actors with actors from the IHE Audit Trail and Node Authentication and will need an Access Control capability that operates in such a cross-enterprise environment. Specific IHE Integration Profiles complementary to XDS are available (e.g., Cross-Enterprise User Authentication, Document Digital Signature, etc.).

- 2290 4. The establishment of independent XDS Affinity Domains will call for their federation, as patients expect their records to follow them as they move from region to region, or country to country. IHE foresees a need for transferring information from one XDS Affinity Domain to another, or to allow access from one XDS Affinity Domain to documents managed in other XDS Affinity Domains. XDS has been designed with this extension in mind. The Cross-Community Access (XCA) Integration Profile that complements XDS provides this function.
- 2295 5. XDS does not address transactions for the management or configuration of an XDS Affinity Domain. For example, the configuration of network addresses or the definition of what type of clinical information is to be shared is specifically left up to the policies established by the XDS Affinity Domain.

2300 **10.1 Actors/Transactions**

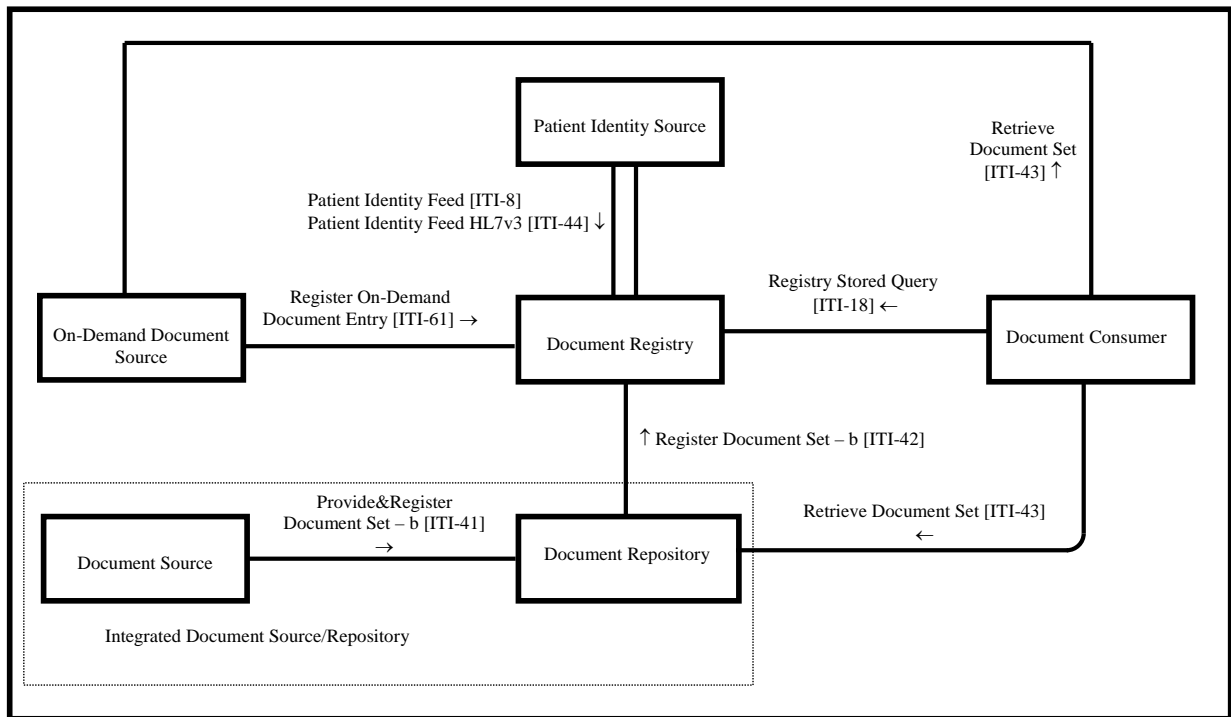


Figure 10.1-1b: Cross-Enterprise Document Sharing – b (XDS.b) Diagram

Table 10.1-1b: XDS.b - Actors and Transactions

Actors	Transactions	Optionality	Section
Document Consumer	Registry Stored Query [ITI-18]	R	ITI TF-2a: 3.18
	Retrieve Document Set [ITI-43]	R	ITI TF-2b: 3.43
Document Source	Provide and Register Document Set-b [ITI-41]	R	ITI TF-2b: 3.41
Document Repository	Provide and Register Document Set-b [ITI-41]	R	ITI TF-2b: 3.41

Actors	Transactions	Optionality	Section
	Register Document Set-b [ITI-42]	R	ITI TF-2b: 3.42
	Retrieve Document Set [ITI-43]	R	ITI TF-2b: 3.43
Document Registry	Register Document Set-b [ITI-42]	R	ITI TF-2b: 3.42
	Registry Stored Query [ITI-18]	R	ITI TF-2a: 3.18
	Patient Identity Feed [ITI-8]	O (Note 2)	ITI TF-2a: 3.8
	Patient Identity Feed HL7v3 [ITI-44]	O (Note 2)	ITI TF-2b: 3.44
	Register On-Demand Document Entry [ITI-61]	O	ITI TF-2b: 3.61
Integrated Document Source/Repository	Register Document Set-b [ITI-42]	R	ITI TF-2b: 3.42
	Retrieve Document Set [ITI-43]	R	ITI TF-2b: 3.43
Patient Identity Source	Patient Identity Feed [ITI-8]	O (Note 1,2)	ITI TF-2a: 3.8
	Patient Identity Feed HL7v3 [ITI-44]	O (Note 1,2)	ITI TF-2b :3.44
On-Demand Document Source	Register On-Demand Document Entry [ITI-61]	R	ITI TF-2b: 3.61
	Retrieve Document Set [ITI-43]	R	ITI TF-2b: 3.43

2305

Note 1: If Assigning Authority of Patient ID presents in the Patient Identity Feed or Patient Identity Feed HL7v3 transaction, the Patient Identity Source is required to use an OID to identify the Assigning Authority. For technical details of the assigning authority information, see ITI TF-2a: 3.8.

Note 2: Document Registry and Patient Identify Source shall implement at least one of Patient Identity Feed or Patient Identity Feed HL7v3.

2310 10.1.1 Actors

10.1.1.1 Document Source

The Document Source Actor is the producer and publisher of documents. It is responsible for sending documents to a Document Repository Actor. It also supplies metadata to the Document Repository Actor for subsequent registration of the documents with the Document Registry Actor.

2315

10.1.1.2 Document Consumer

The Document Consumer Actor queries a Document Registry Actor for documents meeting certain criteria, and retrieves selected documents from one or more Document Repository actors.

10.1.1.3 Document Registry

2320

The Document Registry Actor maintains metadata about each registered document in a document entry. This includes a link to the Document in the Repository where it is stored. The Document Registry responds to queries from Document Consumer actors about documents meeting specific criteria. It also enforces some healthcare specific technical policies at the time of document registration.

2325 **10.1.1.4 Document Repository**

The Document Repository is responsible for both the persistent storage of these documents as well as for their registration with the appropriate Document Registry. It assigns a uniqueId to documents for subsequent retrieval by a Document Consumer.

10.1.1.5 Patient Identity Source

2330 The Patient Identity Source Actor is a provider of unique identifier for each patient and maintains a collection of identity traits. The Patient Identify Source facilitates the validation of patient identifiers by the Registry Actor in its interactions with other actors.

10.1.1.6 Integrated Document Source/Repository

2335 The Integrated Document Source/Repository combines the functionality of the Document Source and Document Repository actors into a single actor that does not initiate nor accept the Provide and Register Document Set transaction. This actor may replace the Document Repository Actor from the perspective of the Register Document Set or Retrieve Document transactions.

10.1.1.7 On-Demand Document Source

2340 The On-Demand Document Source supports On-Demand Document Entries by registering this type of entry with the Document Registry and by responding to Retrieve Document Set transactions for those entries with a document reflecting current information for the entry requested.

10.1.2 Transactions

10.1.2.1 Provide and Register Document Set

2345 A Document Source Actor initiates the Provide and Register Document Set Transaction. For each document in the submitted set, the Document Source Actor provides both the documents as an opaque octet stream and the corresponding metadata to the Document Repository. The Document Repository is responsible to persistently store these documents, and to register them in the Document Registry using the Register Documents transaction by forwarding the document metadata received from the Document Source Actor.

2350

10.1.2.2 Register Document Set

2355 A Document Repository Actor initiates the Register Document Set transaction. This transaction allows a Document Repository Actor to register one or more documents with a Document Registry, by supplying metadata about each document to be registered. This document metadata will be used to create an XDS Document Entry in the registry. The Document Registry Actor ensures that document metadata is valid before allowing documents to be registered. If one or more documents fail the metadata validation, the Register Document Set transaction fails as a whole.

2360 To support composite documents, an XDS Document may be a multipart document. The Document Repository must handle multi-part data sets as an “opaque entity”. The Document

Repository does not need to analyze or process its multi-part structure nor the content of any parts in the context of the XDS Integration Profile.

10.1.2.3 Intentional left blank

10.1.2.4 Registry Stored Query

2365 The Registry Stored Query transaction is issued by the Document Consumer Actor on behalf of a care provider (EHR-CR) to a Document Registry. The Document Registry Actor searches the registry to locate documents that meet the provider's specified query criteria. It will return registry metadata containing a list of document entries found to meet the specified criteria including the locations and identifier of each corresponding document in one or more Document Repositories.

2370

In a Stored Query, the definition of the query is stored on the Registry Actor. To invoke the query, an identifier associated with the query is transmitted along with parameters defined by the query. This has the following benefits:

1. Malicious SQL transactions cannot be introduced
2. Alternate database styles and schemas can be used to implement the Document Registry Actor. This is because the style of SQL query statements is directly related to the table layout in a relational database.

2375

This profile does not define how Stored Queries are loaded into or implemented in the Document Registry Actor.

10.1.2.5 Intentionally Left Blank

2380

10.1.2.6 Patient Identity Feed

2385 The Patient Identity Feed Transaction conveys the patient identifier and corroborating demographic data, captured when a patient's identity is established, modified or merged or in cases where the key corroborating demographic data has been modified. Its purpose in the XDS Integration Profile is to populate the registry with patient identifiers that have been registered for the XDS Affinity Domains.

The Patient Identify Feed Transaction defined in ITI TF-2a:3.8 for HL7v2 and in ITI TF-2b: 3.44 for HL7v3 uses standard HL7 encoding of Patient Identifiers. This is standard encoding for HL7 applications; receiving applications are expected to extract the required data for their use.

2390 When combined with the other XDS transactions, Document Registry actors and other actors that receive HL7 data with Patient Identifiers are required to map the data received in the HL7 message to the format specified in those other XDS transactions. In those transactions, the Patient ID is treated using ebXML encoding rules and not HL7 encoding rules. Specifically, the Patient ID will be treated as a string, and extra components entered in that string shall cause those transactions to fail. XDS actors are required to use the specified encoding for Patient ID values in other transactions and not merely copy the value received in an HL7 transaction.

2395

2400 XDS.b implementations shall support either Patient Identity Feed (ITI TF-2a: 3.8) or Patient Identity Feed HL7v3 (ITI TF-2b: 3.44) or both. It is important to note that the version of HL7 implemented by XDS.b and Patient Identity Feed in a single domain or community need to match in order to allow interoperability. In the case of mixed scenarios, translation between Patient Identity Feed (ITI TF-2a:3.8) and Patient Identity Feed HL7v3 (ITI TF-2b: 3.44) will be required via a bridge or interface engine.

10.1.2.7 Retrieve Document Set

2405 A Document Consumer Actor initiates the Retrieve Document Set transaction. The Document Repository or On-Demand Document Source shall return the document set that was specified by the Document Consumer.

10.1.2.8 Register On-Demand Document Entry

The Register On-Demand Document Entry transaction is used by the On-Demand Document Source to register one or more On-Demand Document Entries in the Document Registry.

2410 **10.1.3 XDS Document Contents Support**

The following table lists a few of the document contents supported in other IHE Integration Profiles, which specify concrete content types for sharing of clinical documents in various domains. These profiles are built on the XDS Profile, and may define additional constraints and semantics for cross-enterprise document sharing in their specific use cases.

2415

Table 10.1-1: List of IHE Integration Profiles and Document Types They Support

IHE Technical Framework Domain	Integration Profile Name	Document Content Supported
IT Infrastructure	An example of an ITI domain content profile defining a document that may be exchanged using XDS is Cross-Enterprise Sharing of Scanned Documents (XDS-SD). Refer to ITI TF-3:5 for other ITI content specifications.	Scanned document, plain text or PDF/A, in HL7 CDA R2 format
Patient Care Coordination	An example of a PCC domain content profile defining a document that may be exchanged using XDS is Cross-Enterprise Sharing of Medical Summaries (XDS-MS). Refer to PCC TF-1 for other document content profiles.	Medical Summary in the HL7 CDA format
Radiology	Cross-Enterprise Document Sharing for Imaging (XDS-I)	Radiology Diagnostic Report in the plain text or PDF formats
		Reference to a collection of DICOM SOP Instances in a manifest document in the DICOM Key Object Selection format

10.2 Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 10.2-1-b along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 10.2-1b: XDS.b - Actors and Options

Actor	Options	Vol. & Section
Document Source	Document Replacement	ITI TF-1: 10.2.1
	Document Addendum	ITI TF-1: 10.2.2
	Document Transformation	ITI TF-1: 10.2.3
	Folder Management	ITI TF-1: 10.2.4
	Basic Patient Privacy Enforcement	ITI TF-2b:3.41.4.1.3.1
	Asynchronous Web Services Exchange	ITI TF-1: 10.2.5
Document Repository	Asynchronous Web Services Exchange	ITI TF-1: 10.2.5
Document Registry (Note 2)	Patient Identity Feed (Note 1)	ITI TF-2a: 3.8
	Patient Identity Feed HL7v3 (Note 1)	ITI TF-2b: 3.44
	Asynchronous Web Services Exchange	ITI TF-1: 10.2.5
	Reference ID	ITI TF-1: 10.2.6
	On-Demand Documents	ITI TF-1: 10.2.7
Integrated Document Source / Repository	Document Replacement	ITI TF-1: 10.2.1
	Document Addendum	ITI TF-1: 10.2.2
	Document Transformation	ITI TF-1: 10.2.3
	Folder Management	ITI TF-1: 10.2.4
	Basic Patient Privacy Enforcement	ITI TF-2b: 3.42.4.1.4.1
	Asynchronous Web Services Exchange	ITI TF-1: 10.2.5
Document Consumer	Basic Patient Privacy Enforcement	ITI TF-2a: 3.18.4.1.3.5 ITI TF-2b: 3.43.4.1.3.1
	Basic Patient Privacy Proof	ITI TF-2a: 3.18.4.1.3.6
	Asynchronous Web Services Exchange	ITI TF-1: 10.2.5
	On-Demand Documents	ITI TF-1: 10.2.7
Patient Identity Source	Patient Identity Feed (Note 1)	ITI TF-2a: 3.8
	Patient Identity Feed HL7v3 (Note 1)	ITI TF-2b: 3.44
On-Demand Document Source	Persistence of Retrieved Documents	ITI TF-1: 10.2.8
	Basic Patient Privacy Enforcement	ITI TF-2b: 3.61.4.1.4.1

Note 1: Document Registry and Patient Identify Source shall implement at least one of Patient Identity Feed or Patient Identity Feed HL7v3.

2425

Note 2: An XDS.b Document Registry has always been required to validate that documents that are registered do contain a confidentialityCode from an XDS Affinity Domain vocabulary. The BPPC Profile is giving some structure to this XDS Affinity Domain defined vocabulary.

10.2.1 Document Replacement Option

2430 In this option the Document Source or Integrated Document Source/Repository shall offer the ability to submit a document as a replacement for another document already in the registry/repository. Grouping with Document Consumer can be used to obtain the most recent metadata and ids to be used in the replace submission.

10.2.2 Document Addendum Option

2435 In this option the Document Source or Integrated Document Source/Repository shall offer the ability to submit a document as an addendum to another document already in the registry/repository.

10.2.3 Document Transformation Option

2440 In this option the Document Source or Integrated Document Source/Repository shall offer the ability to submit a document as a transformation of another document already in the registry/repository.

10.2.4 Folder Management Option

In this option the Document Source offers the ability to perform the following operation:

- Create a folder²
- Add one or more documents to a folder

Note: In order to support document addition to an existing folder, grouping with the Document Consumer may be necessary in order to Query the registry (e.g., for UUIDs of existing folder).

10.2.5 Asynchronous Web Services Exchange Option

Actors that support this option shall support the following:

- 2450 • Document Source Actor shall support Asynchronous Web Services Exchange for the Provide & Register Document Set – b [ITI-41] transaction
- Document Consumer Actor shall support Asynchronous Web Services Exchange for the Registry Stored Query [ITI-18] and Retrieve Document Set [ITI-43] transactions
- 2455 • Document Repository Actor shall support Asynchronous Web Services Exchange for the Provide & Register Document Set – b [ITI-41] and Register Document Set – b [ITI-42], and Retrieve Document Set [ITI-43] transactions

² The term “folder” comes from the medical community which commonly places patient records in folders for specific purposes. In computer science terminology this concept is most consistent with the UNIX directory format, where a file can be simultaneously within multiple directories.

- Document Registry Actor shall support Asynchronous Web Services Exchange for the Registry Stored Query [ITI-18] and Register Document Set – b [ITI-42] transactions

2460 Use of Synchronous or Asynchronous Web Services Exchange is dictated by the individual install environment and affinity domain policy. Refer to section ITI TF-2x: V.5 Synchronous and Asynchronous Web Services Exchange for an explanation of Asynchronous Web Services Exchange.

10.2.6 Reference ID

2465 A Document Registry that supports this option shall store the referenceIdList value provided in a Register Document Set-b [ITI-42] transaction and shall implement the FindDocumentsByReferenceId query as specified in ITI TF-2a:3.18.4.1.2.3.7.14. Document Consumers may use this query to discover documents that have specified values in the document entry metadata referenceIdList.

2470 Document Consumers that use this query will need to be robust to Registries that don't support the query, for example, falling back to the FindDocuments query and filtering the results locally. The Document Consumers will need to also be robust to missing reference Id values due to Document Source actors that don't include the referenceIdList metadata values, or don't include all the reference identifier values expected; or where other inconsistent results may be returned (e.g., XCA gateways which may have different or changing availability of relevant reference
2475 ids).

Document Consumers may fall back to issuing an alternative query, such as FindDocuments, and filtering locally, if the availability or comprehensiveness of a FindDocumentsByReferenceId result is not assured. For example, if an Affinity Domain cannot guarantee that a
2480 FindDocumentsByReferenceId result of a single Document Entry matching an order id is in fact the only Document Entry in the registry for that order id (as opposed to the only Document Entry with that order id in its referenceIdList), then a Document Consumer might issue FindDocuments queries instead.

10.2.7 On-Demand Documents Option

2485 The XDS Integration Profile offers a basic service where any document consumer may discover and retrieve any of the documents for a specific patient that document sources have chosen to make available. Time of release and attestation of content are under the full control/responsibility of the document source. The On-Demand Documents Option offers a complementary service for document consumers to discover one or more document sources that have the capability to
2490 produce, for a specific patient, an on-demand document with content assembled at the time of processing the document consumer retrieve request.

A Document Consumer declares the On-Demand Documents Option when it is able to specify in a query a request for On-Demand Document Entries and is capable of retrieving such entries.

A Document Registry declares the On-Demand Documents Option when it:

- supports registration of On-Demand Document Entries via a Register On-Demand Document Entry Transaction
- 2495

- is also able to respond to Registry Stored Query requests which request On-Demand Document Entries
- stores IsSnapShotOf associations received on a Register Document Set-b transaction

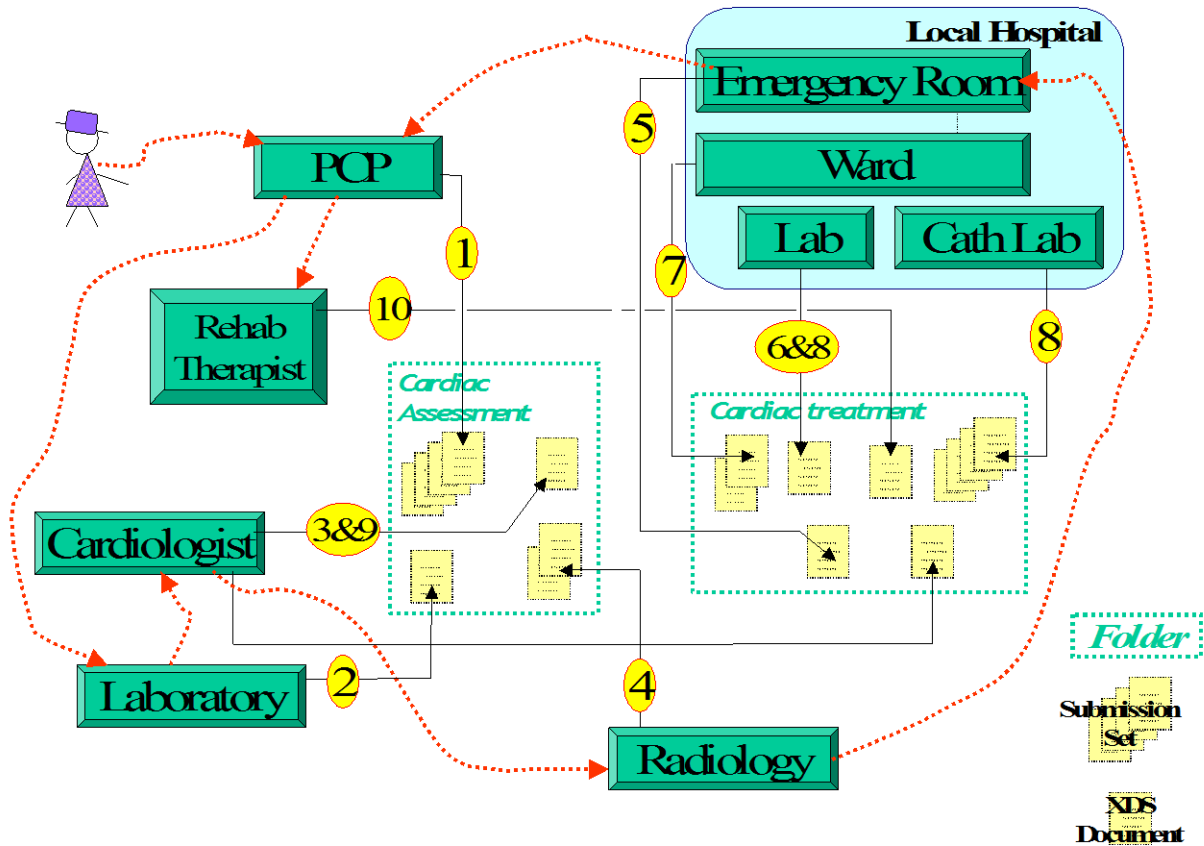
10.2.8 Persistence of Retrieved Documents Option

- 2500 An On-Demand Document Source which supports the Persistence of Retrieved Documents Option shall be grouped with either a Document Source Actor or an Integrated Document Source/Repository Actor in order to register in the Document Registry a new Stable Document Entry which represents each document created as a result of receipt of a Retrieve Document Set which specified the uniqueID of an On-Demand Document Entry. This grouped Document
- 2505 Source or Integrated Source/Repository shall include in the registration an association indicating that the Document Entry begin registered is a snapshot of the On-Demand Document Entry.

10.3 Integration Profile Process Flow

- 2510 A typical patient goes through a sequence of encounters in different care settings. In each care setting, the resulting patient information is created and managed by multiple care delivery information systems (EHR-CRs). Through a sequence of care delivery activities, a number of clinical documents are created. The EHR-LR provides the means to share the relevant subset of these documents, as they are contributed by the various EHR-CRs that are part of the same XDS Affinity Domain.

Example: Cardiac Patient Management Scenario



2515

Figure 10.3-1: Cardiac Patient Management Scenario Transaction Process Flow

2520 This scenario spans about 3 weeks of a patient’s cardiac episode. The patient presents to her
 primary care provider (PCP) with complaints of shortness of breath, nausea, tiredness and chest
 pains. This doctor works closely with a local hospital that has recently established a cardiac care
 network that allows PCPs, cardiologists, laboratories and two local hospitals to share clinical
 documents to improve patient care. This cardiac network is part of a local care data exchange
 community that has been set-up in this community and to which the care plan to which this
 patient belong has encouraged patients to subscribe. Our patient has been provided a health
 2525 record account number.

1. During the patient examination, the PCP records the complaint, and determines that he
 should perform an ECG. He queries the cardiac care network to see if there are prior
 ECG reports (step 1 in Figure 10.3-2), using a coded document class “report” and a coded
 practice setting “cardiology” established by the cardiac care network for ECG reports.
 2530 Among the matching Documents, he locates a prior ECG report that is then retrieved
 (step 2 in Figure 10.3-2). He compares the two results and determines that the patient
 should be referred to a cardiologist. He searches for additional reports in the cardiac care
 network (step 3 in Figure 10.3-2) for this patient, but finds none.

2535

Using the ambulatory EHR system, he creates a submission request onto the patients' health record account number for a "PCP office visit" that includes a submission set consisting of three new documents (visit note, referral letter, new ECG report) and of one reference to the prior ECG report (step 4 in Figure 10.3-2). Following the Cardiology Network XDS Affinity Domain policy, he creates a "cardiac assessment" Folder to contain all four documents in order to facilitate collaboration with the cardiologist.

2540

The repository used by the ambulatory EHR system will then register the documents that are part of this submission request (step 5 in Figure 10.3-2).

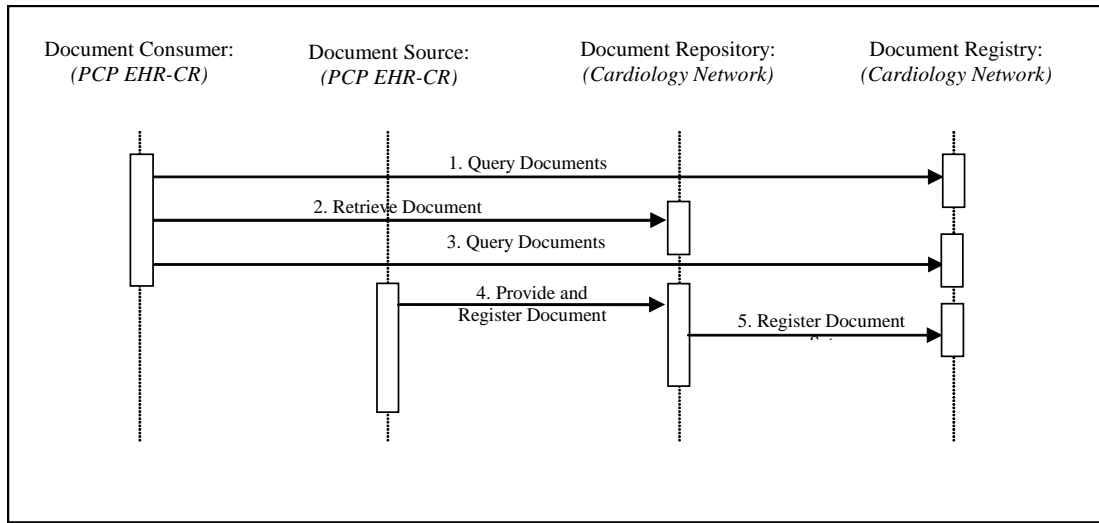


Figure 10.3-2: PCP Query Transactions Process Flow

2545

The PCP EHR system implements the Document Consumer and Document Source actors to issue the Query, Retrieve and Provide & Register transactions as shown in Figure 10.3-2. The transactions are processed by the Document Repository and the Document Registry provided by the cardiology care network.

2550

2. The patient appointment with the cardiologist is scheduled. The patient goes to the lab for the lab tests required before appointment. The lab creates a submission set with a clinical code of "laboratory tests" containing the lab results. The lab is not aware of the "cardiology assessment" folder.

2555

3. The cardiologist sees the patient. He queries the repository for any patient's records in a "cardiac assessment" folder (step 1 in Figure 10.3-3). Available are the visit note from the PCP, the ECG and prior ECG, and the referral letter, which he retrieves and reviews (steps 2-5 in Figure 10.3-3). He also queries for recent lab reports, and finds the lab results (step 6 in Figure 10.3-3). This is also retrieved and reviewed (step 7 in Figure 10.3-3).

2560

4. The cardiologist performs an ultrasound, dictates a visit note, and orders a nuclear stress test. The visit note and ultrasound images and report are registered as a "cardiologist office visit" submission set and placed in the "cardiac assessment" Folder. In addition, the lab report is added to the "cardiac assessment" Folder (step 8 in Figure 10.3-3).

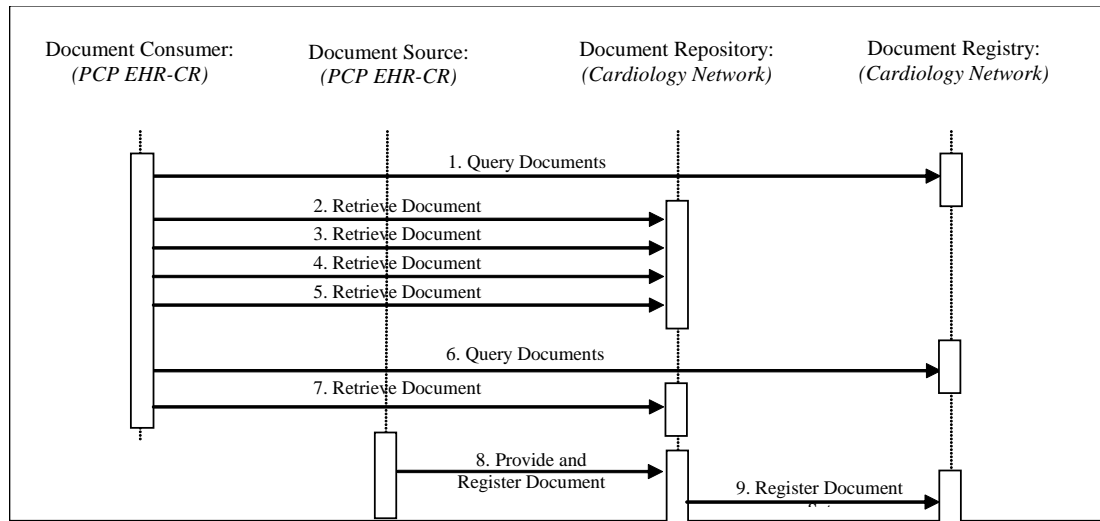
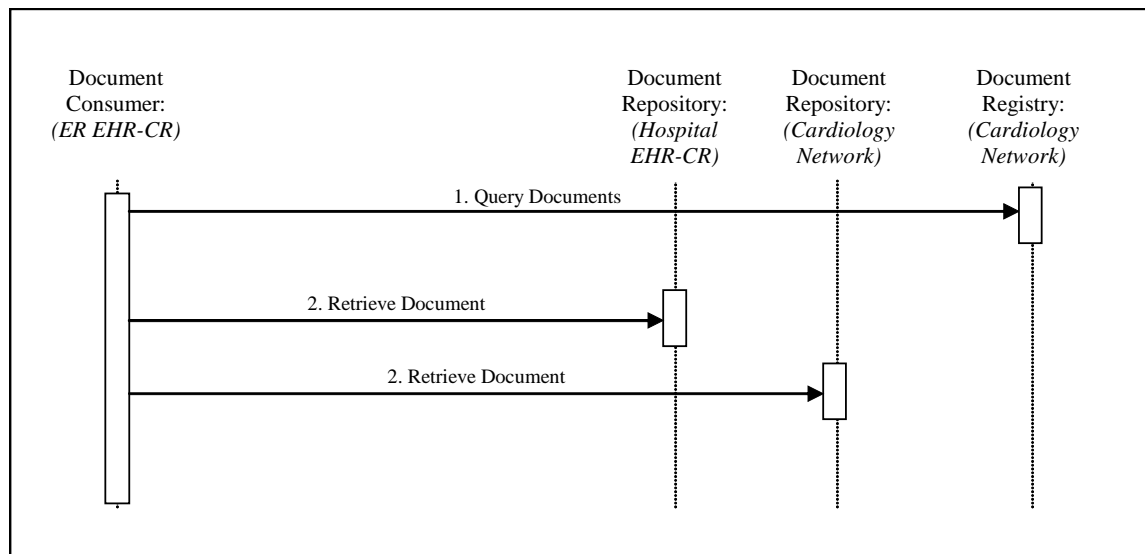


Figure 10.3-3: PCP Query Transactions Process Flow

- 2565 5. The patient is seen at a radiology facility for the nuclear stress test. The test is performed, and the radiologist dictates the report. The nuclear stress test report is registered in a “radiology examination” submission set and associated with the “cardiac assessment” Folder
- 2570 6. Although she has a scheduled appointment with her cardiologist in two days, she wakes up with severe chest pain. On the way to work, she decides to go to the emergency room (ER) of her local hospital. The ER doctor uses the hospital EHR system to query the cardiac care network registry and repositories for documents related to the patient in reverse chronological order (step 1 in Figure 10.3-4). Available documents from latest cardiac related Folder are the visit notes from the PCP and cardiologist, the recent and prior ECGs, the lab results, and the ultrasound images and report, and the nuclear stress test images and report.
- 2575
- The ER doctor retrieves and reviews the two most relevant reports (step 2 and 3 in Figure 10.3-4).
- 2580 The ER doctor orders lab tests, ECG, and places the patient under monitoring. The lab tests and ECG are placed in the hospital EHR that acts as a Document Repository Actor for the cardiac network. Abnormal cardiac activity requires a catheterization, diagnostics and possibly intervention. The ER doctor admits the patient to the cardiology service and contacts the cardiologist.



2585

Figure 10.3-4: ER Query Transactions Process Flow

1. While talking to the ER physician, the cardiologist accesses the cardiac care network from his home office. He queries for all documents related to the patient since the last visit in his office. The nuclear stress test report that he did not previously review is available, along with lab results and ECG results from the ER. The two physicians determine a plan of care and the cardiologist makes arrangements to see the patient in the hospital.
 2. As the patient is transferred from the ER, the ER visit notes are submitted as an “emergency department visit” submission set and placed in a newly created “cardiology treatment” Folder along with the earlier lab and ECG results.
 3. The patient is transferred to an inpatient bed with the following sequence of events.
 - The patient is scheduled for a catheterization procedure in cath lab.
 - Additional lab tests are ordered and performed.
 - A diagnostics procedure is performed in cath lab.
 - An intervention with the placement of a stent is performed.
 - A cath intervention report is dictated.
 - Patient is returned to monitored care for recovery.
 - Education given to patient and family.
 - Discharge Summary dictated by cardiologist.
 - Cardiologist orders lab tests to be completed prior to scheduled follow-up visit.
- The admission assessment, lab results, cath intervention report and key images, and discharge summary form a “cardiology intervention” submission set, which is

2605

registered with the cardiac care network registry in the “cardiac treatment” Folder started by the ER.

- 2610
4. The patient returns to the cardiologist for the post discharge follow-up visit. The resulting visit note, cardiac rehab and summary letters are placed in a “cardiology office visit” submission set and in the “cardiac treatment” Folder.
 5. The patient goes to rehab sessions as scheduled by the cardiologist. The patient recovers and is seen by the PCP and cardiologist for routine visits.

10.4 General Principles

2615 10.4.1 EDR-CR Concept

An EHR-CR or Care-delivery Record abstracts the information system or systems of a care delivery organization, which may support a broad variety of healthcare facilities: private practice, nursing home, ambulatory clinic, acute care in-patient facility, etc.

2620 Typically a patient goes through a sequence of encounters in different care settings as depicted in the figure below.




Figure 10.4.1-1: Sequence of encounters across care delivery organizations

It is out of the scope of this IHE Integration Profile to define or restrict the type of care provided, nor the internal workflow of a care delivery organization. The EHR-CR system participates only to the cross-enterprise clinical document sharing as Document Source and Document Consumer Actors according to the following principles:

2625

1. EHR-CR as Document Source contributes documents in any one of the document formats that are supported by the XDS Affinity Domain (e.g., CDA Release 1, CDA Release 2 with specific templates, DICOM Composite SOP Classes, ASTM-CCR, CEN ENV 13606 etc.).
2. This Profile does not require that the EHR-CR as Document Sources and Consumers store and manage their internal information in the form of documents as they are shared throughout the XDS Affinity Domain.
3. By grouping a Document Source with a Document Repository, an EHR-CR may leverage existing storage provide a unified access mechanism without needing to duplicate storage.

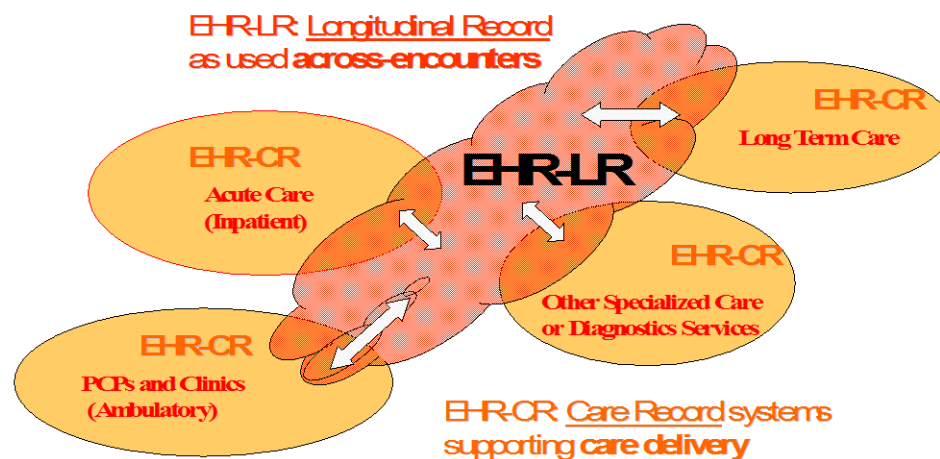
2630

2635

4. EHR-CRs as Document Sources and Consumers are responsible to map their local codes into the XDS Affinity Domain codes if necessary.

2640

The XDS Documents shared by the EHR-CR and tracked by the XDS Registry form a Longitudinal Record for the patients that received care among the EHR-CRs of the XDS Affinity Domain.



2645 **Figure 10.4.1-2: Contributing and sharing to a patients' longitudinal health record**

This shared clinical record is called an EHR-LR in this Integration Profile.

10.4.2 XDS Document Concept

An XDS Document is the smallest unit of information that may be provided to a Document Repository Actor and be registered as an entry in the Document Registry Actor.

2650 An XDS Document is a composition of clinical information that contains observations and services for the purpose of exchange with the following characteristics: Persistence, Stewardship, Potential for Authentication, and Wholeness. These characteristics are defined in the HL7 Clinical Document Architecture specification. An XDS Document may be human readable (with the appropriate application). In any case, it should comply with a published standard defining its structure, content and encoding. IHE intends to define content-oriented Integration Profiles

2655 relying on such content standards to be used in conjunction with XDS.

The XDS Integration Profile manages XDS Documents as a single unit of information; it does not provide mechanisms to access portions of an XDS Document. Only the Document Sources or Document Consumers have access to the internal information of the XDS Document. When

2660 submitted for sharing, an XDS Document is provided to the Document Repository Actor as an octet stream. When retrieved through the Retrieve Document Set transaction, it shall be unchanged from the octet stream that was submitted.

2665 The Document Source Actor is responsible for producing the metadata that will be submitted to the Document Registry Actor to form the XDS Document Entry that will be used for query purposes by XDS Consumer Actors. The Document Source maintains responsibilities over the XDS Documents it has registered. It shall replace XDS Documents that may have been submitted in error. See ITI TF-1: Appendix K for a more detailed discussion of the concept of XDS Document.

2670 XDS Documents are required to be globally uniquely identified. See ITI TF-2x: Appendix B for a definition of globally unique identifiers.

10.4.3 Submission Request

An XDS Submission Request is a means to share XDS Documents. It may be conveyed:

- by a Document Source Actor in a *Provide and Register Document Set Transaction* to the Document Repository Actor, or
- 2675 • by a Document Repository Actor in a *Register Document Set Transaction* to the Document Registry Actor

An XDS Submission Request contains elements of information that will ensure the proper registration of XDS Documents. These are:

1. Metadata to be placed in Document Entries for new XDS Documents being submitted,
- 2680 2. A Submission Set that includes the list of all new XDS Documents and Folders being submitted and optionally a list of previously submitted XDS Documents,
3. If desired, Folders to be created with the list of included XDS Documents (new document being submitted as well as previously submitted),
4. If desired, addition to previously created Folders of lists of XDS Documents (new document being submitted as well as previously submitted), and
- 2685 5. Zero or more XDS Document octet streams for the new XDS Documents being submitted.

2690 Following a successful Submission Request, new XDS Documents, Submission Set, and Folders included in the Submission Request are available for sharing in an XDS Affinity Domain. In case of failure to process a Submission Request, the Submission Set and any XDS Documents and Folders shall not be registered.

10.4.4 Submission Set Concept

2695 An XDS Submission Set is related to care event(s) of a single patient provided by the care delivery organization EHR-CR performing the submission request. It creates a permanent record of new XDS Documents as well as pre-existing (i.e., already registered) XDS Documents that

have a relationship with the same care event(s). It also includes the record of new XDS Folders creation.

2700 An XDS Submission Set shall be created for each submission request. It is related to a single Document Source Actor and is conveyed by a single Provide & Register Document Set Transaction or a Register Document Set Transaction.

The Document Registry may be queried to find all documents registered in the same XDS Submission Set.

2705 The same XDS Document, initially registered as part of a Submission Set, may also be referenced by later XDS Submission Set. This allows older documents relevant to the present care of a patient to be associated with more recent Submission Sets.

XDS provides complete flexibility to EHR-CRs to relate Documents and Submission Sets to an encounter, a visit, an episode of care, or various workflow processes within EHR-CRs.

10.4.5 Concept of Folder

2710 The purpose of an XDS Folder is to provide a collaborative mechanism for several XDS Document Sources to group XDS Documents for a variety of reasons (e.g., a period of care, a problem, immunizations, etc.) and to offer the Document Consumers a means to find all Document Entries placed in the same Folder. The following principles apply to an XDS Folder:

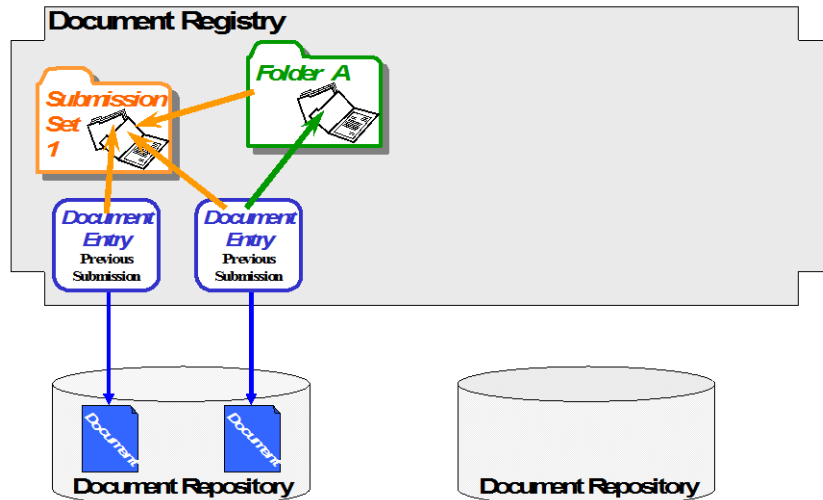
1. A Folder groups a set of XDS Documents related to the care of a single patient,
2. One or more Document Source Actors may submit documents in a given Folder,
- 2715 3. A Folder may be created by a Document Source,
4. The content of a Folder is qualified by a list of codes/meaning,
5. Document Source Actors may find existing Folders by querying the Document Registry or by means outside the scope of XDS (e.g., Cross-enterprise workflow, such ePrescription, eReferral, etc.),
- 2720 6. Once created a Folder is permanently known by the Document Registry,
7. Placing previously existing Documents in Folders is recorded by the updating Submission Set,
8. Folders in XDS may not be nested,
9. The same documents can appear in more than one Folder, and
- 2725 10. Folders have a globally unique identifier.

10.4.6 Example of use of Submission Request, Submission Set and Folder

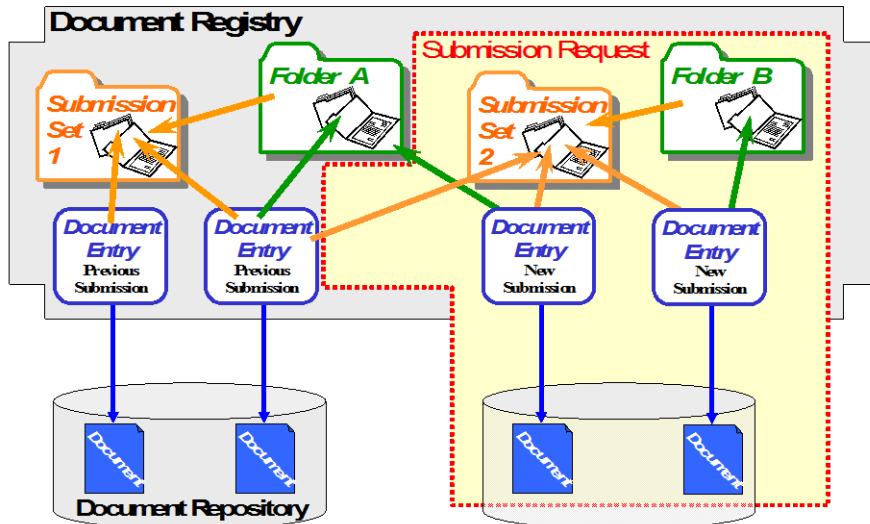
2730 The sequence of figures below shows an example of a submission request that includes two new documents, a reference to a pre-existing document and the use of two folders. The first figure depicts the initial state of a Document Registry in which two Documents have been submitted where one is associated with a Folder A. The second figure depicts a submission request that

adds two new documents, placing one of them into a pre-existing folder and the other one into a new Folder B.

Document Repository and Registry – Initial State



Document Repository and Registry – Submission Request



2735

Figure 10.4.6-1: Example of a submission flow to an XDS Registry

From the above example, the contents of a Submission Set are shown by the figure below. The Document Entries associated with the Submission Set are logical part of the Submission Set.

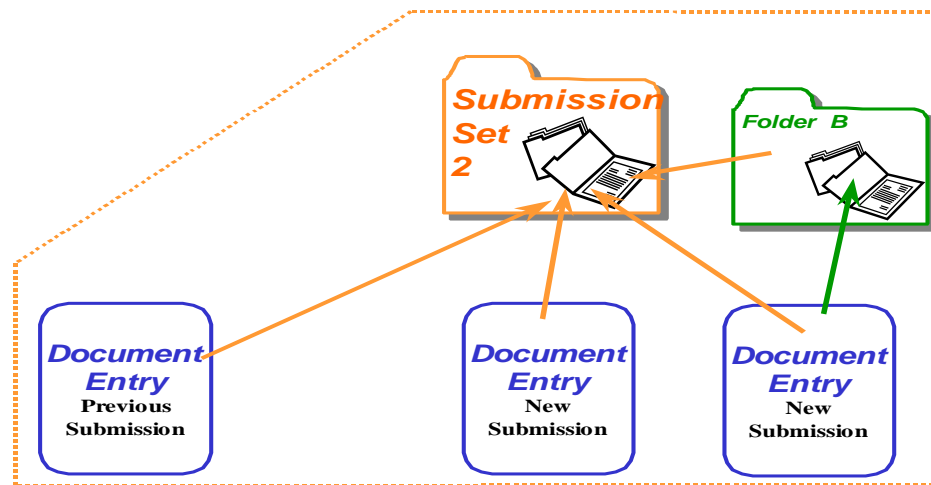


Figure 10.4.6-2: The logical content of a Submission Set

2740 10.4.7 XDS Registry Data Model and Attributes

The XDS Integration Profile provides a means to place documents in a repository chosen by the Document Source, and also to place information about this document (or metadata) in an entry of the Document Registry that manages the XDS Affinity Domain.

2745 The term metadata reflects that this information is “about” the documents. The purpose of well-specified document metadata is to enable a uniform mechanism for Document Consumers to locate clinical documents of interest much in the way a card catalog in a library helps readers find the book they want.

2750 This section addresses the high-level data model in which the metadata is registered and against which queries of the XDS Document Registry are performed. Then it presents the specific attributes that may be registered and used to filter the document entries of the registry.

10.4.7.1 XDS Document Registry Data Model

The following entities are used in the XDS Document Registry Data Model:

2755 **XDS Stable Document Entry:** Information entity managed by a Document Registry Actor that contains a set of metadata describing the major characteristics of an XDS Document along with a link to the Document Repository Actor where the actual XDS Document may be retrieved.

XDS On-Demand Document Entry: Information entity managed by a Document Registry Actor that contains a set of metadata describing the characteristics of a set of on-demand content which can be instantiated through retrieval from the On-Demand Document Source. On-demand content is document content that is generated at the time of the retrieve request and contains the most recent information available related to the metadata which describes it.

2760

XDS Document Entry: Refers generically to both XDS Stable Document Entry and XDS On-Demand Document Entry. See Section 10.4.13 for more details about the use of both kinds of XDS Document Entries.

2765 **XDS Document:** A stream of bytes stored in a Document Repository Actor and pointed to by an XDS Document Entry.

2770 **XDS Folder:** A logical container that groups one or more XDS Document Entries in any way required (e.g., by source care delivery activities, by episode, care team, clinical specialty or clinical condition). This kind of organizing structure is used variably: in some centers and systems the Folder is treated as an informal compartmentalization of the overall health record; in others it might represent a significant legal portion of the EHR relating to the originating enterprise or team. The Folder is a means of providing organization of XDS Documents (or Composition in EHRCOM). The same XDS Document Entry may belong to zero or more Folders.

2775 **XDS Submission Set:** When XDS Documents are registered by a Document Source Actor, they shall be included in one and exactly one Submission Set. An XDS Submission Set groups zero or more new XDS Documents and references to already registered XDS Documents to ensure a persistent record of their submission.

2780 **XDS Submission Request:** A Submission Request includes one and only one Submission Set, zero or more new XDS Folders and assignment of XDS Documents into new or existing Folders. A Submission Request is processed in an atomic manner by the Document Repository and the Document Registry (i.e., all XDS Documents included or referenced in a Submission Set as well as the Folders and inclusion of Folders references are registered or none will). This ensures that they are all made available to Document Consumer Actors at the same time.

2785

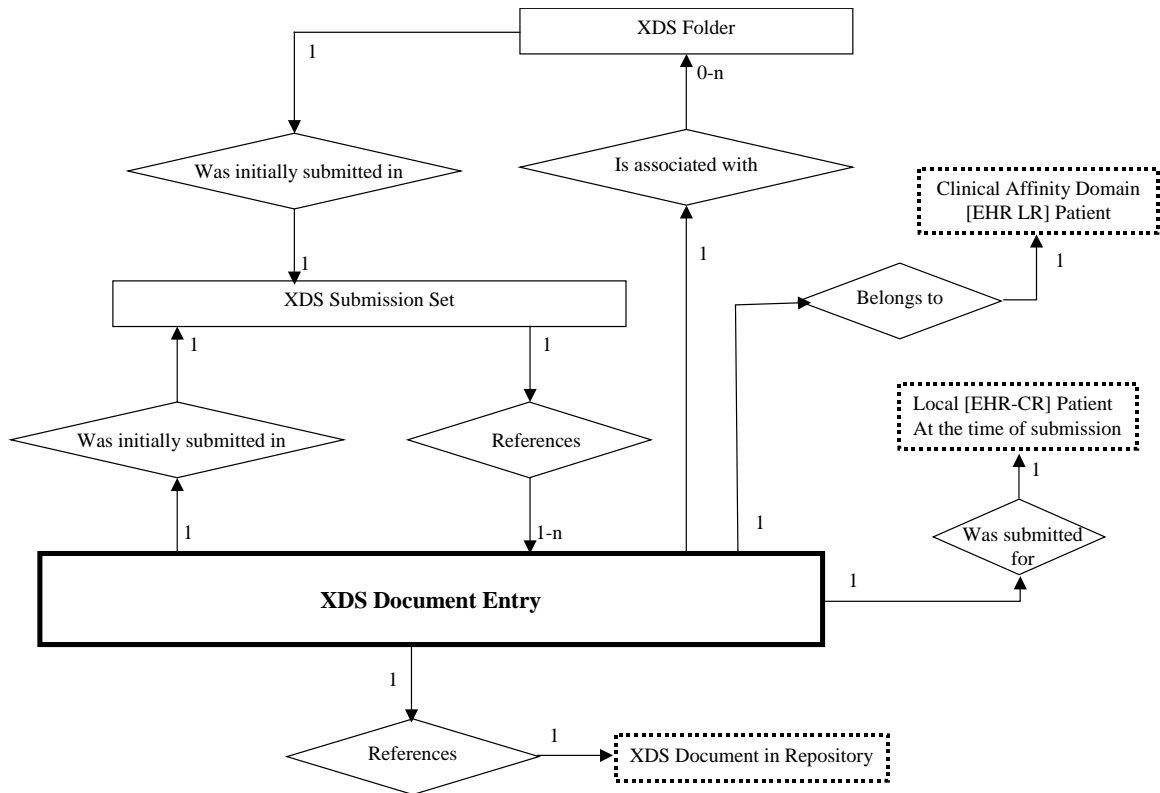


Figure 10.4.7-1: XDS Document Registry Data Model

10.4.7.2 Attributes of the XDS Document Entries

2790 The specific attributes of each entity in the above registry data model have been selected from document header attributes from several standards (see ITI TF-2x: Appendix L), including:

- ANSI/HL7 CDA R1-2000
- HL7 CDA Release 2 (draft) Document header definition (Dec 2003 Committee Ballot)
- Composition attributes from EHR ENV 13606 (draft)

2795 XDS defines a well-focused set of primary attributes that support the most common use cases to search the most relevant documents. These include:

Patient Id
Service Start and Stop Time
Document Creation Time
Document Class Code and Display Name
Practice Setting Code and Display Name
Healthcare Facility Type Code and Display Name
Availability Status (Available, Deprecated)
Document Unique Id

2800 The three codes (Document Class, Practice Setting and Healthcare facility Type) are code set that are expected to generally include a limited number of values (between 10 and 100), thus ensuring a reasonably easy search capability.

2805 A number of additional query attributes or attributes used to perform a secondary selection in order to decide to retrieve a specific document are also defined by this Integration Profile. At the Document Level, these include a fine grained Document Type (e.g., LOINC classification), a list of Event Code that can be used as key word, the document author and associated institution, the document relationship to manage replacement addendum and a variety of transformations, a confidentiality code, language code, etc.

The complete list of attributes and their definition is documented in ITI TF-3: 4.1.

10.4.8 Concept of an XDS Affinity Domain

2810 An XDS Affinity Domain is an administrative structure made of a well-defined set of Document Source Actors, set of Document Repositories, set of Document Consumers organized around a single Document Registry Actor that have agreed to share clinical documents.

Note: Document Sources, Repositories and Consumers may belong to more than one XDS Affinity Domain and share the same or different documents. This is an implementation strategy and will not be further described.

2815 Note: the XDS Integration Profile does not support the federation of XDS Affinity Domains directly, but the Cross-Community Access (XCA) Profile addresses the cooperation of multiple Document Registry Actors serving different XDS Affinity Domains.

2820 A number of policies will need to be established in an XDS Affinity Domain in order to ensure effective interoperability between Document Sources and Consumers. Some of the key technical policies include (A more extensive list of policy agreements that need to be made by XDS Affinity Domains is discussed in ITI TF-1: Appendix L):

1. The document formats that will be accepted for registration
2. The various vocabulary value sets and coding schemes to be used for the submission of metadata of document, submission set and folders registration.
3. The Patient Identification Domain (Assigning Authority) used by the Document Registry.

2825 See ITI TF-1: Appendix K for a detailed discussion of the concepts of XDS Affinity Domain.

10.4.9 Patient Identification Management

Since the central focus of the DS Integration Profile is “sharing documents”, it is critical that each document be reliably associated with the corresponding patient (Patient Id).

2830 The XDS Document Registry is not intended to be an authority for patient identification and demographics information. This Integration Profile uses a Patient Identity Source Actor as the authoritative source of Patient Identifiers (master patient ID) for the XDS Affinity Domain.

Note: This Integration Profile can be easily extended to support a scenario where no master patient ID is defined (i.e., no Patient Identity Source for the XDS Affinity Domain). Such an option would require the use of federated patient identities at the time of query of the XDS Document Registry.

2835 The following principles are defined:

- 2840 1. The Patient Identifier Domain managed by the Patient Identity Source Actor in the XDS Affinity Domain is the source of patient identifiers (and merge operations) used by the XDS Document Registry to link Documents to a specific Patient. This Patient Identifier Domain is called the XDS Affinity Domain Patient Identification Domain (XAD-PID Domain).
2. Submission Requests for Documents related to Patients with IDs not registered in the XDS Affinity Domain Patient Identifier Domain shall be rejected by the XDS Document Registry.
- 2845 3. The XDS Document Registry will contain certain patient information (e.g., source patient ID, Surname, Given Name, Sex, Birthdate) for the purpose of audits and potential verification by Document Consumers. As this Integration Profile does not make any assumptions about the referential integrity and update of this information, these fields³ shall not be used as query matching keys.
- 2850 4. As XDS Document Sources and Consumers may belong to different Patient Identification Domains, these systems need to cross-reference their own local Patient ID to the corresponding patient ID in the XAD-PID Domain of the Registry. Preferably, these systems may choose to use the IHE Patient Identifier Cross-referencing Integration Profile (See ITI TF-1: Appendix E.3) for this purpose.
- 2855 5. The XDS Document Registry is responsible for validating Document metadata in accordance with the XDS Affinity Domain's policies. The Document Registry should reject submissions Requests that do not conform to these policies.

³ It is possible to submit a new document to replace a previously submitted one, with a new document entry created in the registry to correct for errors in the submitted document in the original submission request. However this is not a mechanism that updates only the metadata, as the replaced document is only deprecated and remains pointed by the original metadata.

2860

The figure below depicts an example of an XDS Affinity Domain with its Patient Identifier Domain (called XAD) and two EHR-CRs where the cross-referencing is performed internally to the Document Source and the Document Consumer Domains (Domain C and Domain D2 respectively).

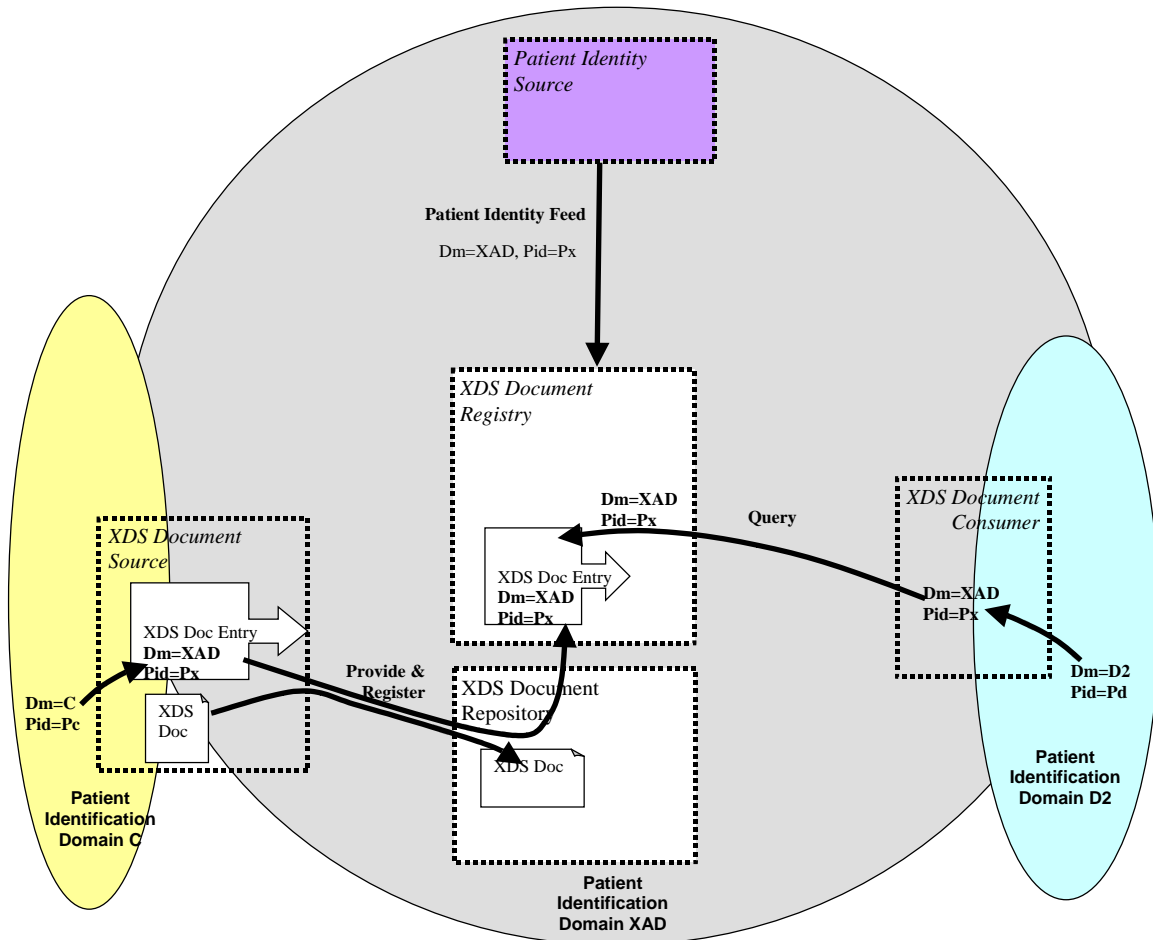


Figure 10.4.9-1: XDS Affinity Domain with patient ID cross-referencing internal to the EHR-CRs

10.4.10 Document Lifecycle

2865 **10.4.10.1 Document Availability Status**

Each XDS Document contained in a XDS Document Registry will be assigned one of the following Availability Status codes:

Approved: Available for patient care (assumes that it is authenticated, if applicable)

Deprecated: Obsolete, but may still be queried and retrieved

2870 The XDS Document availability status is set to “approved” after the XDS Document Repository and the XDS Document Registry have successfully processed a submission request.

Note: ebXML Registry Services defines a Status of Submitted, which is used in a transient manner to provide an atomic submission. It is not significant to make this specific status externally visible.

2875 An “approved” XDS Document may be changed to “deprecated” under the primary responsibility of its original Document Source with possible patient supervision. It is part of security policies that are beyond the scope of the XDS Integration Profile to have the XDS Repository/Registry enforce this ownership. The reason and responsible party for deprecating a document are tracked as part of the XDS Document Registry audit trail, which is a required capability. A “deprecated” Document remains available for Document Consumer queries. Except for the status change, a “deprecated” Document Entry metadata remains the same as when it was in the “approved” status.

2880 An “approved” or “deprecated” XDS Document Entry may be deleted. This change is associated with the decision to completely remove a Document from an XDS Document Repository and the corresponding Document Entry from the XDS Document Registry. The XDS Affinity Domain shall establish the security policies associated with Document deletion. There are no transactions defined by this Integration Profile to support such operation.

2885

See ITI TF-1: Appendix K for a detailed discussion of the concepts of XDS Document life cycle.

10.4.10.2 Document Relationships

XDS Documents may be related to predecessor documents by one of three methods:

2890

- Replacement,
- Addendum
- Transformation
- Transformation-Replacement

2895

These relationships between XDS Documents are tracked in the XDS Document Registry. The parent relationship attribute contained in the metadata of such Documents is a coded value that describes the type of relationship. An original Document has no parent and consequently its parent Id and parent relationship are absent. XDS Document Registry shall reject submissions that contain relationships to documents that are not registered or have been “deprecated”.

2900

Document stubs are supported by XDS to allow for a valid relationship to a known but not registered Document.

2905

A replacement document is a new version of an existing document. The replacement document has a new document Id; its parent Id attribute contains the document Id of the Document Entry associated with the previous version of the XDS Document, and parent relationship contains the code “RPLC”. The Document Entry for the previous version shall have its Availability Status changed to “deprecated”.

Replacement of On-Demand Document Entries works exactly the same as replacing Stable Document Entries in that the new entry is marked Approved and the replaced entry is marked Deprecated. But an On-Demand Document Entry which has been replaced cannot be expected to

2910 provide content that is current, so the On-Demand Document Source may return an error if a Retrieve Document Set request is received with a uniqueID for a deprecated entry.

2915 An addendum is a separate XDS Document that references a prior document, and may extend or alter the observations in the prior document. It modifies the parent document, but the parent document remains a valid component of the patient record and shall remain in the state “approved” or available for care. The addendum XDS Document metadata contains the identifier of the previous XDS Document version in parent Id, and its parent relationship contains the code “APND”.

2920 A transformed document is derived by a machine translation from some other format. Examples of transformed documents could be CDA documents converted from DICOM Structured Reporting (SR) reports, or a rendering of a report into a presentation format such as PDF. The transform XDS Document contains the document Id of the previous version in parentId, and its parent relationship contains the code “XFRM”. XDS Affinity Domains may define rules that determine whether or not a transformed XDS Document replaces the source, but typically this would not be the case. If it is, an additional parent relationship of type “RPLC” is to be used.

10.4.11 Document Query

2925 Query return info shall be either:

- a list of Registry Objects Values (e.g., XDS Document Entries)
- a list of Registry Objects UUIDs. This allows an XDS Document Consumer to receive a potentially long list of matching entries and to request them by subsets.

10.4.12 Intentionally left blank.

2930 10.4.13 XDS Document Entry Types

Of all the XDS Entity types (see 10.4.7.1 for a list of all XDS Entities) there are two that support retrieval of content and are generically referred to as XDS Document Entries.

2935 **XDS Stable Document Entry:** A Stable Document Entry contains metadata about an already created document available for retrieval. This document is stable because the contents have been effectively combined in the exact representation that will be returned in a Retrieve Document Set. Support for this entry type is required of Document Source, Document Consumer, Document Repository, Document Registry, and Integrated Document Repository/Source. Only a Document Repository or Integrated Document Repository/Source will register a Stable Document Entry with a Document Registry, and so only those actors support retrieval of Stable Document Entries.

2945 **XDS On-Demand Document Entry:** An On-Demand Document Entry contains metadata describing the characteristics of a set of on-demand content and a unique identifier which can be used to create a document which collects the latest, most recent available information at the time of retrieval. On-Demand Document Entries get registered in a Register On-Demand Document Entry transaction by the On-Demand Document Source.

10.4.13.1 Use Cases Summary

Each type of XDS Document Entry is designed for a different environment of document sharing. To determine which type of XDS Document Entry should be used, an analysis of the environment is needed. Stable Document Entries and On-Demand Document Entries may co-exist in an XDS Affinity Domain. In particular, queries can be formulated to return both document entry types for a specific patient.

Case 1: Sharing of Stable, Source Attested Document

In this model, health data is published by a well-identified and responsible source (clinician, care delivery organization, individual consumer, etc.). The source determines when a meaningful collection of data should be published based on clinical events or other activities understood by the source and potential remote entities. The source publishes stable collections of data in the form of one or more documents, for potential use by other entities. Other entities access the documents by querying for a list of documents that have been published, and retrieving those that are of interest. Documents are source attested; consistency and wholeness is the responsibility of the source, which provides explicit context for each document (legal record keeping requirements). Complete documents are stored in repositories. If errors need to be corrected or updates are needed, they are the responsibility of the source. Entities accessing these records are offered trust guaranties equivalent or superior to paper records shared today.

Case 2: Sharing of On-Demand Document

As part of an extract from a health record, some documents, containing specific types of content, may be automatically generated with non-stable or dynamic content. These documents are “dynamic” in that each retrieval of the document may result in different content. Documents whose content is assembled at the time of retrieval have no inherent stable properties like persistence or stewardship. The wholeness of a dynamically created document will not be based on any clinician attestation and may require careful clinical interpretation depending on the content and the span of aggregation performed by the document authoring system. If the document authoring system originates from a single care delivery organization, context and wholeness may be quite good. If the aggregation was performed across multiple health delivery organizations, the aggregator may have difficulties assuming legal and clinical responsibility for the aggregated content. An example of a dynamically created document is a summary that collects information related to multiple healthcare events or on-going healthcare events.

On-Demand Documents Process Flow

The use of On-Demand Documents supports registration of the availability of dynamically assembled content, content that is expected to change over time, and in response to a retrieve request always returns the most current content available to the responder. The use of On-Demand Documents is intended for an application architecture where the supplier of data wishes to provide, through a single request mechanism, the most current knowledge available at the time of the request. Dynamically created content are services that present data in the form of documents using XDS and XCA transactions and infrastructure. The dynamic nature of the data means this environment is more complicated to support but allows easy access to, for instance,

summary data for a specific patient. However, it does not provide for robust source attestation of the overall document content because the content is selected through automation rather than overseen and attested in whole by a clinician.

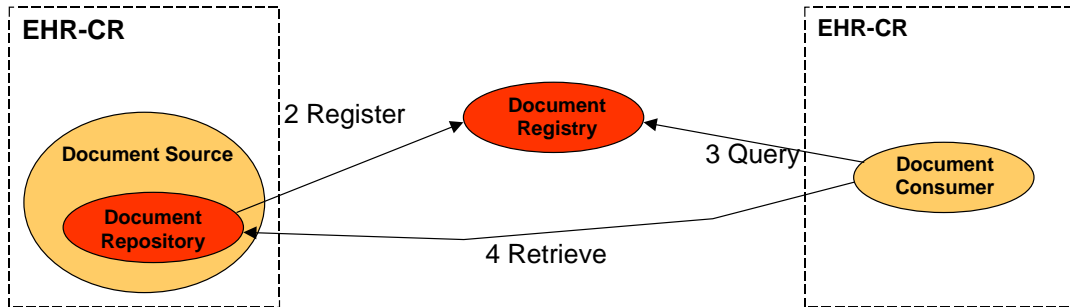
- 2990 The following describes a high level XDS workflow for On-Demand Documents where the content for the data to be shared is in a clinical database.
- The On-Demand Document Source registers an On-Demand Document Entry for a patient for which it maintains healthcare data. The On-Demand Document Entry describes what type of content will be returned and what format the content will be delivered in (e.g., XDS-MS patient summary). This On-Demand Document Entry is registered with the XDS Document Registry.
- 2995
- A Document Consumer which supports access to On-Demand Document Entries queries the Document Registry requesting that On-Demand Document Entries (possibly as well as other types of entries) be returned. The Document Registry returns the entries created by the On-Demand Document Source that match the query parameters.
- 3000
- The Document Consumer uses the metadata from the On-Demand Document Entry to request the most recent content from the On-Demand Document Source. The On-Demand Document Source searches its clinical database for content of the type defined by the On-Demand Document Entry, formulates this into a document, and returns the document. If the Document Consumer repeats the request for data defined by the On-Demand Document Entry, the content returned on the second request may be different than the first due to updates to the underlying clinical database. The On-Demand Document Source may choose to register the document created as a consequence of the Document Consumer's retrieve request, in which case it will group with a Document Source to register the newly created document. A special association is used to link the newly created document to the On-Demand Document Entry.
- 3005
- 3010

10.5 Implementation Strategies

3015 The XDS Integration Profile addresses the requirements of three major implementation strategies reflecting different groupings of actors within an EHR-CR as well as different configurations of the EHR-LR. This range of implementation strategies reflects the need to accommodate a variety of workflows and configurations. These implementation strategies may coexist in some environments. Other implementation strategies are possible.

- Strategy 1: Repository at the Source. A single information system acts as both the Document Source and Document Repository for the documents it creates and registers with the Document Registry
- 3020
- Upon completion of a phase of care, an EHR-CR will register a submission-set of documents in a Document Repository Actor with which it is grouped (same system). Then it registers this set of documents (newly created and priors documents of interest) with the Document Registry Actor [2].
- 3025
- Any other Document Consumer Actor in the XDS Affinity Domain may query the Document Registry Actor to find documents related to all phases of care for the patient

[3]. It may choose to retrieve some of these documents from any Document Repository Actor [4].

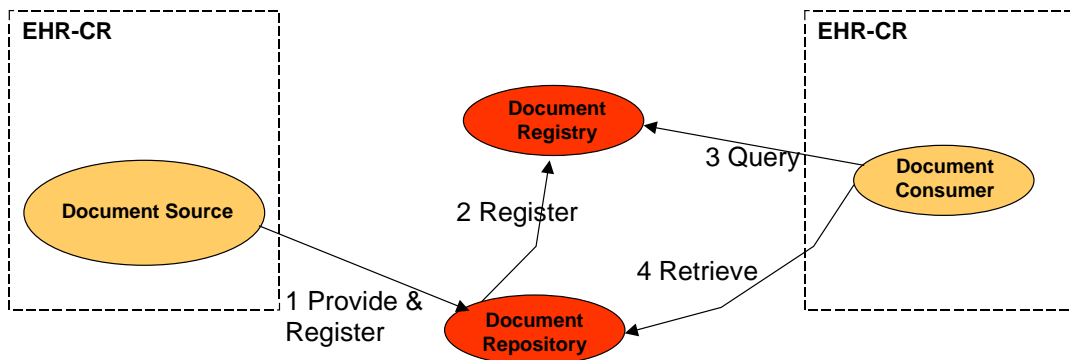


3030

Figure 10.5-1: Implementation Strategy with Repository at the Source

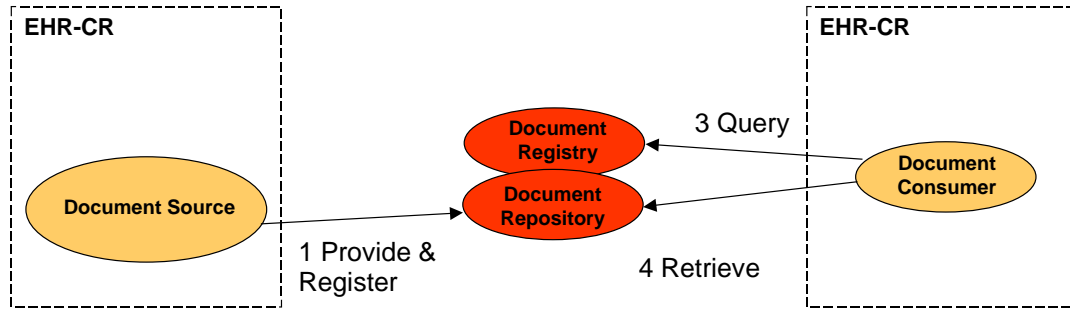
3035

- Strategy 2: Third Party Repository. The EHR-CR does not wish to be a Document Repository Actor, but rather uses the services of a third party Document Repository Actor to which it entrusts the documents it creates. First it provides both the metadata and the set of documents to this Document Repository Actor [1], which in turn forwards the registration request for the set of documents (newly created and prior documents of interest) to the Document Registry Actor [2].
- Any other Document Consumer Actor may query the Document Registry Actor to find out about documents related to all phases of care for the patient [3]. It may choose to retrieve some of these documents from any Document Repository Actor [4].



3040

Figure 10.5-2: Implementation Strategy with 3rd party repository



3045

Figure 10.5-3: Implementation Strategy with 3rd party central repository and registry

3050

3055

- Strategy 3: Direct Patient Transfer-Referral. The Document Source Actor completes a phase of care for a patient. It decides to directly provide and register [1] the set of documents (newly created and prior documents of interest) with a Document Repository [2] that has been grouped along with the Document Registry with the EHR-CR Document Consumer (Grouped Actors).
- In this case the span of the XDS Affinity Domain may be quite limited as it could be defined to cover only the two EHR-CRs. However the same transaction [1] applies. Note that, in this implementation strategy the other transactions, although supported by the actors, are not used by the Document Consumer since the Document Registry and Document Repository reside within the Document Consumer.



3060

Figure 10.5-4: Direct patient referral with registry and repository at consumer

Patient access to an EHR-LR may be supported by a specialized EHR-CR (i.e., a portal) implementing the Document Source and Document Consumer Actors.

10.6 Patient Identifier Communication Requirements

3065

When using ITI Transaction 8 as the patient identity feed, ITI TF-2a: 3.8 defines the format requirements for the patient identifier in PID-3. Specifically, the value for PID-3.4, Assigning Authority can be omitted, expressed using the first subcomponent (namespace ID) or the second and third subcomponents (universal ID and universal ID type). These rules shall apply in this profile:

- 3070 1. If the Patient Identity Source does not include a value for PID-3.4, Assigning Authority, then
- a. PID-3, Patient Identifier List, is constrained to include one entry referring to one identifier.
 - b. The Patient Identity Source and Document Registry shall agree that all messages from this source shall refer to a single assigning authority.
- 3075 2. If PID-3.4 does contain a value for PID-3.4, Assigning Authority, then
- a. The Patient Identifier Source may send multiple patient identifiers with properly formatted components. The Document Registry shall be responsible for selecting the one identifier from the Patient Identifier List (not necessarily in the first position) that is too used to register the selected patient.
 - b. As specified in ITI TF-2a: 3.8, the value for PID-3.4, Assigning Authority, can be expressed using the first subcomponent (namespace ID) or the second and third subcomponents (universal ID and universal ID type). Both methods shall be accepted by the Document Registry and shall be considered as equivalent.

When using Transaction ITI-44, the Assigning Authority is required.

3085 Transactions ITI-18, ITI-41 and ITI-42 express patient ID as a string that is not parsed using typical HL7 parsing logic; please refer to requirements for Patient ID in those transactions. Document Registry actors will have to map between the Patient ID feed provided in ITI-8 or ITI-44 as described above and the PID provided by those transactions in this profile.

3090 XDS.b implementations shall support either Patient Identity Feed (ITI TF-2a: 3.8) or Patient Identity Feed HL7v3 (ITI TF-2b: 3.44) or both. It is important to note that the version of HL7 implemented by XDS.b and Patient Identity Feed in a single domain or community need to match in order to allow interoperability. In the case of mixed scenarios, translation between Patient Identity Feed (ITI TF-2a: 3.8) and Patient Identity Feed HL7v3 (ITI TF-2b: 3.44) will be required via a bridge or interface engine.

3095 **10.7 Security Considerations**

Coordinating the security and privacy policies of all the care delivery organizations in an XDS Affinity Domain may be a challenge. An agreement is needed on security procedures, goals, auditing, record keeping, etc. This can result in changes to other enterprise policies, such as human resources procedures. XDS Affinity Domain members are trusting to some extent the access of their published data by other members of the XDS Affinity Domain. The level of control is dependent on Policies and application of other security and privacy profiles offered by IHE. This relationship requires a close ongoing partnership that ensures ongoing maintenance of policies, procedures, and activities. If laws change, relevant policies must be adjusted throughout the group. Corporate changes to group members affect the policies. Security events must be managed as a group. This must be managed as a long-term activity, not a one-time event.

Particular problem areas are likely to be:

- 3110
 - Authorized access and modification policies. The details of access policies are likely to have enterprise differences and conflicts that must be resolved. The XDS Affinity Domain relationships also introduce new policy requirements. For example, changes to employment (e.g., employee hiring and firing) must now include suitably rapid notifications to other XDS Affinity Domain members. (See ATNA and XUA)
 - Changes to privacy restrictions (e.g., divorces) now require full XDS Affinity Domain notifications, not merely enterprise notifications. (See BPPC)
- 3115
 - Audit trail and access record keeping are often quite sensitive internal enterprise activities that must now be appropriately coordinated with the full XDS Affinity Domain. (See ATNA and Section 10.8.1)
 - Changes to laws and regulations now affect not only the policies of the individual enterprises; they also must be reflected in the XDS Affinity Domain relationship contracts, policies, and procedures.
- 3120
 - Patient identity management. (See PIX/PDQ/XCPD)
 - Patients may have access through an authorized Document Consumer or Document Source implemented in an application such as a PHR.
 - Trans-border communication of Personal Health Information (PHI) often presents legal and regulatory issues.
- 3125 ITI TF-2x: Appendix K goes into more detail listing many of the threats, objectives, policies, and mitigations that need to be coordinated among XDS Affinity Domain members.

The XDS Integration Profile for two main reasons does not prescribe such Security and Privacy policies. First, it is clear that the broad range of possible solutions to these policies that will depend on the legal framework and the types of healthcare system, calls for XDS to be offer such
- 3130 flexibility. Decisions in this domain will have some impact on the implementations of XDS Actors, but it is expected that these will be minimal.

10.7.1 Use of ATNA to address Basic Security

3135 The XDS Profile requires all actors be grouped with a Secure Node Actor as defined in the IHE Audit Trail and Node Authentication Integration Profile. This use of the ATNA Profile in an XDS Affinity Domain does not require a centralized XDS Affinity Domain Audit Repository Actor.

The use of ATNA along with XDS does require that each member of the XDS Affinity Domain does have audit and security mechanisms in place. See ITI TF-2x: Appendix K.

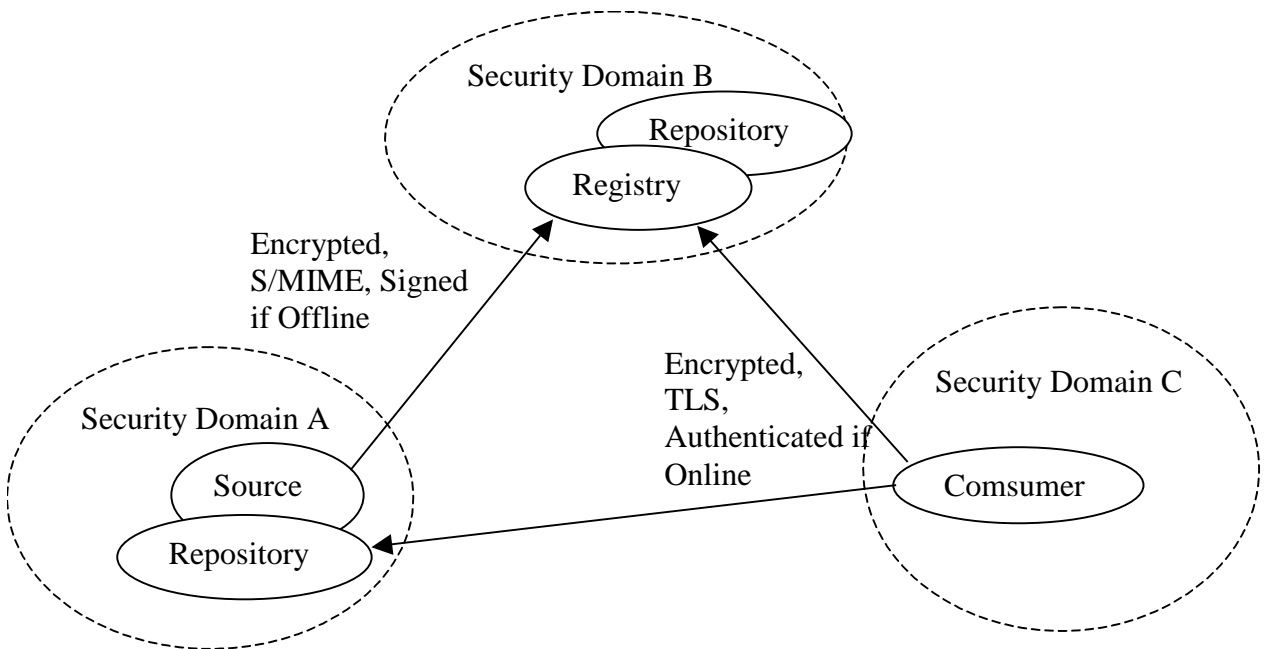
3140 The individual actors involved are often members of different secure domains, as illustrated in Figure 3.14.5.1-2. The data transfers between different secure domains may need different protection than transfers within a secure domain. The use of encryption and other security mechanisms will depend upon the policies of the domains involved.

Transfers within a single secure domain may choose to omit encryption if it is unnecessary, so it is recommended that the online transfer security mechanisms be configurable. Certificate

3145 management and exchange is defined as part of the XDS Affinity Domain business relationships and no IHE Integration Profile is specified at this time, see ITI TF-1: Appendix L.

Each transaction will result in audit records describing the transaction. Each secure domain has its own audit server to capture the records for the actors that are within that domain. Access to audit records by other enterprises within the XDS Affinity Domain is managed and controlled by the business relationship terms of the XDS Affinity Domain. There is no automatic IHE transaction for such access.

The audit records that shall be generated (references IHE ATNA Integration Profile) by normal XDS activities are defined in the appropriate Security Considerations section of each transaction:



3155 All Actors are part of the same Clinical Affinity Domain

Figure 10.7-1: Example Security Domain Relationships

3160 Security and Privacy can be further addressed through the application of IHE-BPPC, IHE-XUA. See these profiles for their impact and use.

10.8 Intentionally Left Blank

11 Personnel White Pages (PWP)

3165 The Personnel White Pages (PWP) Profile provides access to basic directory information on human workforce members to other workforce members within the enterprise. This information has broad use among many clinical and non-clinical applications across the healthcare enterprise. The information will be used to

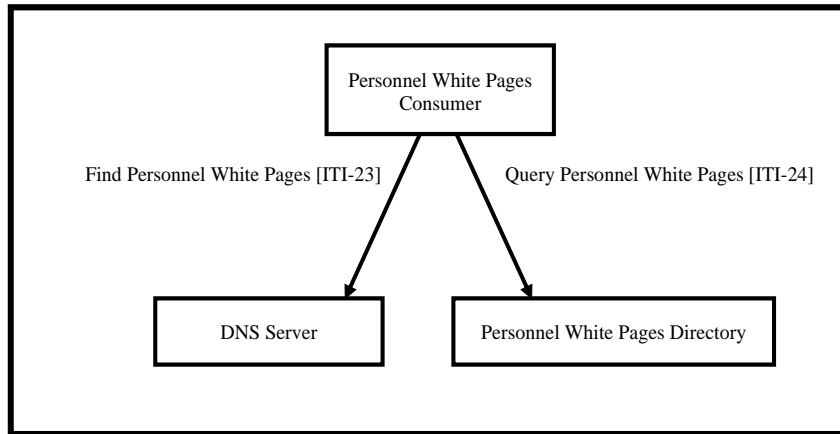
1. enhance the clinical workflow
 - a) contact information,
 - b) phone numbers,
 - 3170 c) email address
2. enhance the user interface
 - a) displayable names,
 - b) titles

3175 This Personnel White Pages Profile specifies a method of finding directory information on the User Identities (user@realm) supplied by the Enterprise User Authentication (EUA) Integration Profile. This Profile assumes but does not define access controls, and audit trails. The use of the PWP Profile is intended for use within a healthcare enterprise. Extension to support sharing of the PWP between healthcare enterprises is possible but not fully addressed by this profile. The PWP Profile is the first step on an IHE roadmap that includes Digital Certificates, Encryption, 3180 Digital Signatures, Medical Credentials, and Roles.

The directory need not support use cases beyond healthcare operations (e.g., Human Resource Operations), but does not forbid a properly designed overlap with other use cases. This profile does not intend for patients or other individuals that are not acting as part of the human healthcare workforce.

3185 11.1 Actors/ Transactions

Figure 11.1-1 shows the actors directly involved in the PWP Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in EUA Profile are not necessarily shown.



3190

Figure 11.1-1: Personnel White Pages Profile Actor Diagram

Table 11.1-1 lists the transaction for each actor directly involved in the PWP Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 11.2.

3195

Table 11.1-1: PWP Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Personnel White Pages Consumer	Find Personnel White Pages [ITI-23]	O	ITI TF-2a: 3.23
	Query Personnel White Pages [ITI-24]	R	ITI TF-2a: 3.24
DNS Server	Find Personnel White Pages [ITI-23]	R	ITI TF-2a: 3.23
Personnel White Pages Directory	Query Personnel White Pages [ITI-24]	R	ITI TF-2a: 3.24

11.2 PWP Integration Profile Options

3200

Options that may be selected for this Integration Profile are listed in Table 11.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 11.2-1: PWP Integration Profile - Actors and Options

Actor	Options	Vol. & Section
Personnel White Pages Consumer	<i>no option</i>	-
DNS Server	<i>no option</i>	-
Personnel White Pages Directory	<i>no option</i>	-

3205 **11.3 PWP Integration Profile Process Flow**

The Personnel White Pages Profile addresses the following use cases:

- 3210 • A Clinical user logs into an acquisition device that is acting as a Personnel White Pages Consumer. The clinical application queries the DNS Server Actor using [ITI-23] to find the Personnel White Pages Directory. The clinical application then queries [ITI-24] the Personnel White Pages Directory using the user’s username and displays the user’s full name with First Name, Middle, and Last. There are information fields to support both European and Asian naming conventions.
- 3215 • The Clinical user acquires clinical data. The application queries [ITI-24] the Personnel White Pages Directory for the user’s demographics to include the user’s organization identification to embed in the data record.
- The User then needs to send this report by means of email to a colleague. The application allows the user to search [ITI-24] the Personnel White Pages Directory for the destination user, and selects the destination user’s email address.
- 3220 • The User reviews an existing clinical report and finds initials have been recorded in the report. The user system does a query [ITI-24] of the Personnel White Pages Directory for the initials found in the report and the system displays the displayable name(s).

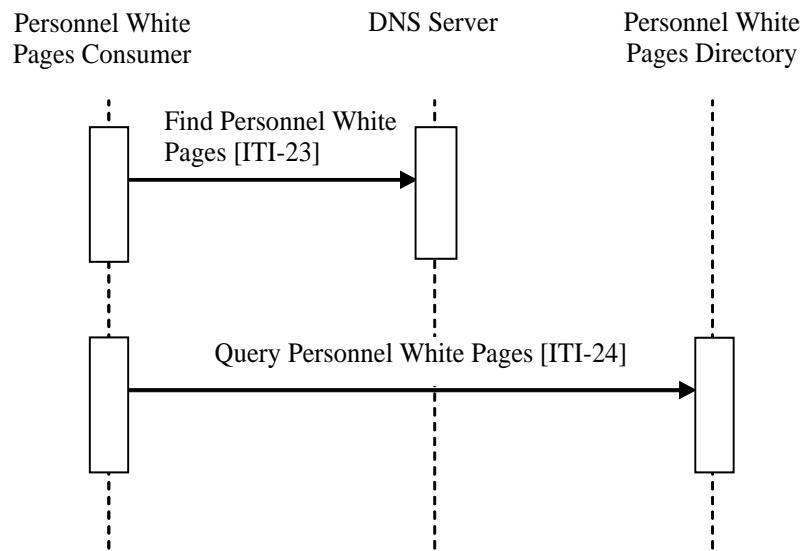


Figure 11.2-1: Basic Process Flow in PWP Profile

3225 **12 This is intentionally left blank.**

13 Cross Enterprise User Assertion (XUA) Integration Profile

3230 The Cross-Enterprise User Assertion Profile (XUA) provides a means to communicate claims about an authenticated principal (user, application, system...) in transactions that cross enterprise boundaries. To provide accountability in these cross-enterprise transactions there is a need to identify the requesting user in a way that enables the receiver to make access decisions and proper audit entries. The XUA Profile supports many solutions including enterprises that have chosen to have their own user directory with their own unique method of authenticating the users, and others that have chosen to use a third party to perform the authentication.

3235 There are transactions defined by IHE that cross enterprise boundaries and are web-services based on ITI TF-2x: Appendix V. The existing IHE profiles for an authenticated user identity (IHE Enterprise User Authentication Profile [EUA]) are not intended to function in cross-enterprise transactions. In a cross-enterprise environment it is more likely that the transactions will be going between two enterprises that maintain their own independent user directories (IHE Personnel White Pages [PWP]). This type of requirement is the focus of the Identity Federation standards. Identity Federation has received much attention by the security and the platforms industry. Identity Federation is agnostic to the type of user directory; it allows for a centralized user directory, but also supports the more powerful federation of user directories. Identity Federation supports:

- 3245
 - A Country that delegates the provisioning of all users into a single assigning authority domain (e.g., France) and provides a common service that handles all user authentication requests
 - Support for centralized user directories
 - A Region that knits together a network of cooperating hospitals and clinics where each hospital/clinic manages its own users.
- 3250
 - Support for distributed user directories
 - Patients who wish to use an identity provider of their choosing (e.g., ISP, email provider).
 - Support for non-healthcare specific user directories
 - A Hospital that provisions users by issuing identity badges with picture and name printed, RFID for building access, and smart-card for strong authentication
- 3255
 - Support for claims about the method used to authenticate the user (e.g., strong authentication methods such as smart-cards)
 - A Small clinic in a rural setting that supports a dozen users.
 - Support for small scale systems (e.g., user at a kiosk, system using simple passwords)
- 3260
 - A General practice doctor retrieving results of a test performed by an outpatient clinic, where the outpatient clinic wants to have an audit trail specific to the user requesting the information.
 - Support for the service provider to get a user identity for audit log purposes

- An automated System, based on a scheduled procedure, that is capable of being a delegate for a doctor pre-fetches the available documents so that it can determine a relevant few documents to offer to the doctor when the patient arrives

3265

The XUA Profile leverages Web-Services Security, SAML 2.0 Token Profile and the various profiles from [W3C](#), and [OASIS](#) to support identity federation. In this way we will be able to take advantage of the vast experience of the communities outside of healthcare standards. This profile leverages the experience of programs around the globe that have started work with SAML in healthcare.

3270

13.1 Use Cases

The XUA Profile supports complex environments, for example one where two different trust domains are operating under different technology, procedures, role-models, etc. They are cooperating in the XDS Affinity domain under an overarching trust relationship policy (See ITI TF-2x: Appendix L) that indicates that these differences can be rationalized. The XDS transactions are transferring control from one entity to another, for example, when using XDS to exchange data between a single doctor practice and large multi-site hospital. It is not likely that they will all agree to the same access control model (organizational roles, functional roles, workflows, permissions, etc.). It is not necessary to have the same access control across these entities, but it is reasonable that at the policy level they will agree to a set of processing rules. This illustrates an important fact that the XUA is useful for security audit logging, but is to a lesser extent useful for access controls.

3275

3280

The following is a list of use-cases that have been proposed for XUA. Some of these use-cases will not be supported due to lack of standards or sufficient guidance on the proper solution.

3285

1. Country that provisions users into a single assigning authority domain (e.g., Germany) and handles all user authentication requests
 - Support for centralized user directories

3290

2. Region that knits together many competing hospitals and clinics where each hospital/clinic manages its own users.

- Support for distributed user directories
3. Patients who wish to use their email provider as their authentication authority uses a PHR-like application to access their own information in an XDS Affinity Domain.
 - Support for non-healthcare specific user directories

3295

4. Hospital that issues identity badges with picture and name printed, RFID for building access, and smart-card for strong authentication
 - Support for claims about the method used to authenticate the user (e.g., strong authentication methods such as smart-cards)
5. Small clinic in a rural setting that supports a dozen users.
 - Support for small scale systems (e.g., user at a kiosk, system using simple passwords)

- 3300 6. General practice doctor who retrieving results of a test performed by an outpatient clinic, where the outpatient clinic wants to have an audit trail specific to the user requesting.
- Support for the service provider to get a user identity for audit log purposes
7. System, based on a scheduled procedure, pre-fetches the available documents so that it can determine a relevant few documents to offer to the doctor when the patient arrives.
- 3305
- Support for identifying the user as the system for tasks that are not initiated by a human user
8. User using Registry or Repository where the service provider wants to be assured that the user has been authenticated to a specific assurance level. This is not a case of not trusting the system, but recognition that the requester supports different levels of authentication. For example the system supports a proximity card as a form of authentication, as well as Smart-Card with PIN. This is not a replacement for ATNA access controls which give distributed access controls.
- 3310
- User Identity with level of assurance of that identity is needed.
9. Specialized XDS Affinity Domain for Emergency Dataset. In this case the transfer of information to the XDS Consumer is not critical to fully control, and thus the administration is willing to accept requests from any system as long as they can provide a user-assertion from a trusted source. This trusted-source may be a specialized identity provider for First Responders. (See RSA Pilot)
- 3315
- In this case only a user identity with proper linkage to a trusted identity provider is needed. No specific attributes are needed.
- 3320
10. User acting in an identified clinical role accesses the Registry where the Registry wants to know the user identity and the role they are acting in to record the identity and role in the audit log.
- Support inclusion of functional roles as named vocabulary
- 3325
- The Role of the user as the data subject (patient)
11. Service provider wants to enforce some form of access controls based on the user identity and/or functional role.
- Support for the service provider to augment access controls based on some non-specified rules that are applied to the user and/or functional role
- 3330
12. Access to a document by an individual that can't be identified because the Assertion Provider is not accessible

13.2 XUA Development

- 3335 The vast majority of the use-cases (items 1-11) rely on claims about an authenticated identity, which a SAML 2.0 Identity Assertion can provide. This is a mature standard produced by OASIS. XUA Profile is focused on Web-Services transactions that follow ITI TF-2x: Appendix V. XUA specifies that when a Cross-Enterprise User Assertion is needed, these Web-Services

3340 transactions will additionally use the Web-Services Security header with a SAML 2.0 Token containing the identity Assertion. As with any IHE profile, the applications are not forbidden to use other methods of providing the principal (user) identity, providing that interoperability has been assured through some policy.

A very clear need on all the use-cases is the recording of the user identity in any security audit logs. The XUA Profile does not define these auditable events. The need to record a security audit event is driven by the grouped transactions (e.g., Registry Stored Query, and Retrieve Document Set). XUA does specify how to reference the Identity Assertion in an ATNA Audit Message.

3345 The method of authenticating the principal (user) and the method that the X-Service User Actor (e.g., XDS.b Document Consumer) uses to get the Identity Assertion are outside the scope of this profile.

3350 There are principal (user) attributes that appear to be needed in the use-cases: Doctor, Patient, Guardian, Emergency-Access. The Identity Assertion can contain attributes about the principal (user). At this time it is not clear what standards to use to identify these attributes and their values, so this is left to specific implementations that have defined a local vocabulary or vocabulary translation.

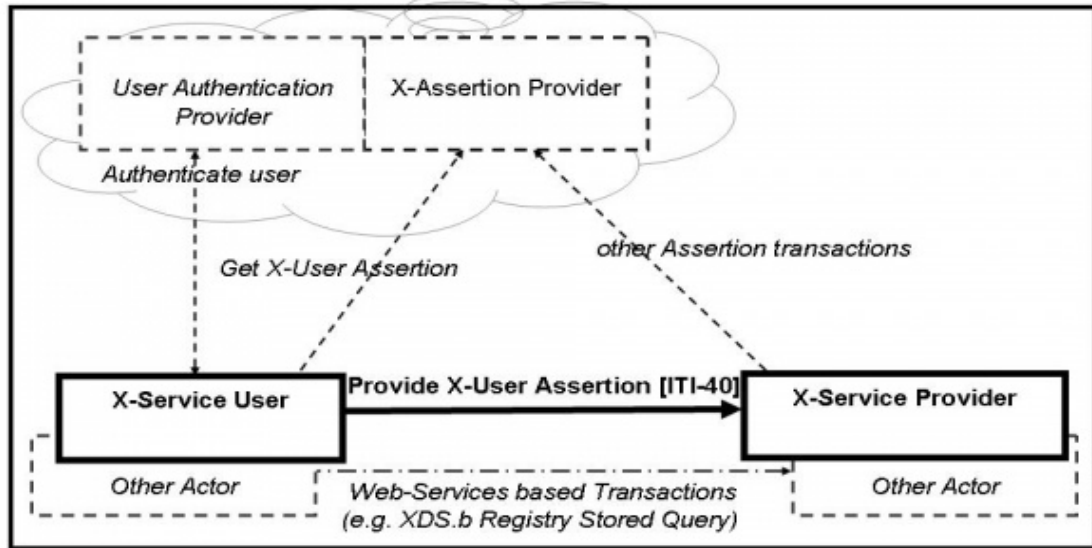
3355 The method used by the X-Service User (e.g., XDS.b Document Consumer) Actor to determine the contents of the Identity Assertion is outside the scope of this profile. This might be accomplished using the SAML Metadata and WS-Policy.

It is expected that extending this solution to HL7 and DICOM will be supported in the future.

13.3 Intentionally Left Blank

13.4 Actors/Transaction

3360 Figure 13.4-1 shows the actors directly (Bold and Solid Boxes) involved in the XUA Integration Profile and the relevant transactions between them (Bold and Solid Line). The diagram also shows ancillary actors (Dashed and Grey Boxes) that are not profiled but include interactions (Dashed and Grey Lines). Actors grouped with are shown as the dashed line between the X-Service User and the X-Service Provider.



3365

Figure 13.4-1: Cross-Enterprise User Assertion Actor Diagram

Table 13.4-1 lists the transactions for each actor directly involved in the XUA Profile. The ancillary actors and associated transactions may be supported by various technologies and system configurations varying from internal shared services to infrastructures for identity management.

3370

In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 13.5.

3375

Table 13.4-1: XUA - Actors and Transactions

Actor	Transaction	Optionality	Section
X-Service User	Provide X-User Assertion [ITI-40]	R	ITI TF-2b: 3.40
X-Service Provider	Provide X-User Assertion [ITI-40]	R	ITI TF-2b: 3.40

13.5 Options

Options that may be selected for this Integration Profile are listed in Table 13.5-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

3380

Table 13.5-1: XUA - Actors and Options

Actor	Option	Section
X-Service User	<i>Subject-Role</i>	13.5.1
	<i>Authz-Consent</i>	13.5.2
	<i>PurposeOfUse</i>	13.5.3
X-Service Provider	<i>Subject-Role</i>	13.5.1
	<i>Authz-Consent</i>	13.5.2
	<i>PurposeOfUse</i>	13.5.3

13.5.1 Subject-Role

3385 Role-Based-Access Control is a common architecture for managing and enforcing
 3390 authorizations. In this model users are assigned to roles, and permissions are assigned to these
 roles. In this way the administration of permissions to users is grouped through roles. In an
 RBAC model these roles are well known and agreed to between the system that is managing
 users and the access control system. Given that a Federated Identity environment allows for a
 loose coupling of the Identity-Manager and the Access Control point, there is a need to have a
 well-defined vocabulary for the roles. This set of roles will likely expand over time and will be
 extended with local codes into a Value-Set within any specific Security-Domain.

3395 This option recommends that the Value-Set be derived from the role codes found in SNOMED-
 CT, ISO 21298, or ASTM E1986. The Value-Set used would bridge between different policy
 domain roles used in a client domain to those used in the service domain. In this way it is
 possible for local role definitions to be used as long as they can be bridged to the roles found in
 the selected Value-Set. Implementations should expect that the Value-Set used may be using
 locally defined values. The use of the IHE Sharing of Value-Sets (SVS) Profile may assist with
 this.

See section ITI TF-2b:3.40.4.1.2.1 and ITI TF-2b:3.40.4.1.3.1 for transaction requirements.

3400 13.5.2 Authz-Consent

There are transactions where the requester of the transaction knows of specific
 Consent/Authorization evidence that would enable that transaction. The identification could be
 used by the relying party Access Control engine as a hint. The Access Control engine could
 explicitly retrieve that specific object, validate that it is indeed a properly formed
 3405 Consent/Authorization, and determine if that Consent/Authorization does indeed affect the
 Access Control decision.

3410 This can be used where the requester had previously published Consent/Authorization evidence
 Document but where this new knowledge had not yet propagated fully to the Access Control
 infrastructure. This mechanism is also useful where the requester is under regulatory obligations
 to include evidence of Consent/Authorization on each transaction. When this Option is used in
 conjunction with the Basic Patient Privacy Consents (BPPC) Profile this would allow for the

requester to include in the XUA Assertion identification of a newly published BPPC Document. This option leverages the BPPC consent model.

See section ITI TF-2b:3.40.4.1.2.2 and ITI TF-2b:3.40.4.1.3.2 for transaction requirements.

3415 **13.5.3 PurposeOfUse**

As explained in the [IHE Access Control White Paper](#), there are Access Control decisions that are based on the ultimate use of the data. For example a Patient may have provided a BPPC Consent/Authorization for treatment purposes, but explicitly disallowed any use for research regardless of de-identification methods used. The purpose of use is also informative to the ATNA audit log to enable specific reporting of Accounting of Disclosures and Breach Notification. To enable this type of Audit Logging and Access Control decision there is a need to include in the XUA Assertion the intended purpose for which the data will be used. One specific purpose of use would be a Break-Glass / Emergency-Mode-Access.

3420 This option recommends that the Value-Set be derived from the codes found in ISO 14265, or XSPA. Implementations should expect that the Value-Set used may be using locally defined values. The use of the IHE Sharing of Value-Sets (SVS) Profile may assist with this.

3425 See section ITI-TF-2b:3.40.4.1.2.3 and ITI-TF-2b:3.40.4.1.3.3 for transaction requirements.

13.6 Grouping

13.6.1 Audit Trail and Node Authentication (ATNA)

3430 The X-Identity Assertion is valuable and must be protected against confidentiality risks. In some Profiles (e.g., XDS), there is already an inherited requirement to group with ATNA Secure Node or Secure Application Actor. This grouping forces the network transactions to utilize mutually authenticated and encrypted TLS or equivalent. This is leveraged by XUA to support the protection of the X-User Assertion to some risks to confidentiality and integrity. When ATNA

3435 Secure Node or Secure Application grouping is not required, there will need to be some other mechanism to protect the Provide X-User Assertion.

ITI TF-2b: 3.40.4.2 includes encoding rules for representing an X-User Assertion in an ATNA Audit Message.

13.6.2 Cross-Enterprise Document Sharing (XDS)

3440 When an XDS.b Document Consumer is grouped with X-Service User Actor, the XDS.b Document Consumer shall conform to all the requirements in the Provide X-User Assertion Transaction. The Document Consumer will obtain a properly scoped XUA Assertion targeted for the XDS.b Document Registry or XDS.b Document Repository. The method used may be through internal means, SAML 2.0 Core protocols, WS-Trust, or any other means.

3445 The XDS.b Document Registry and XDS.b Document Repository when grouped with the XUA X-Service Provider Actor shall conform to all the requirements in the Provide X-User Assertion Transaction. The XUA Profile does not constrain how the Assertion can be used (e.g., ignored, access control, etc.).

13.6.3 Enterprise User Authentication (EUA)

3450 An application that groups EUA and XUA Actors may use WS-Trust to get the X-User Assertion from the Security Token Service (STS). In this case the AuthnContextClassRef element of the SAML assertion shall be:

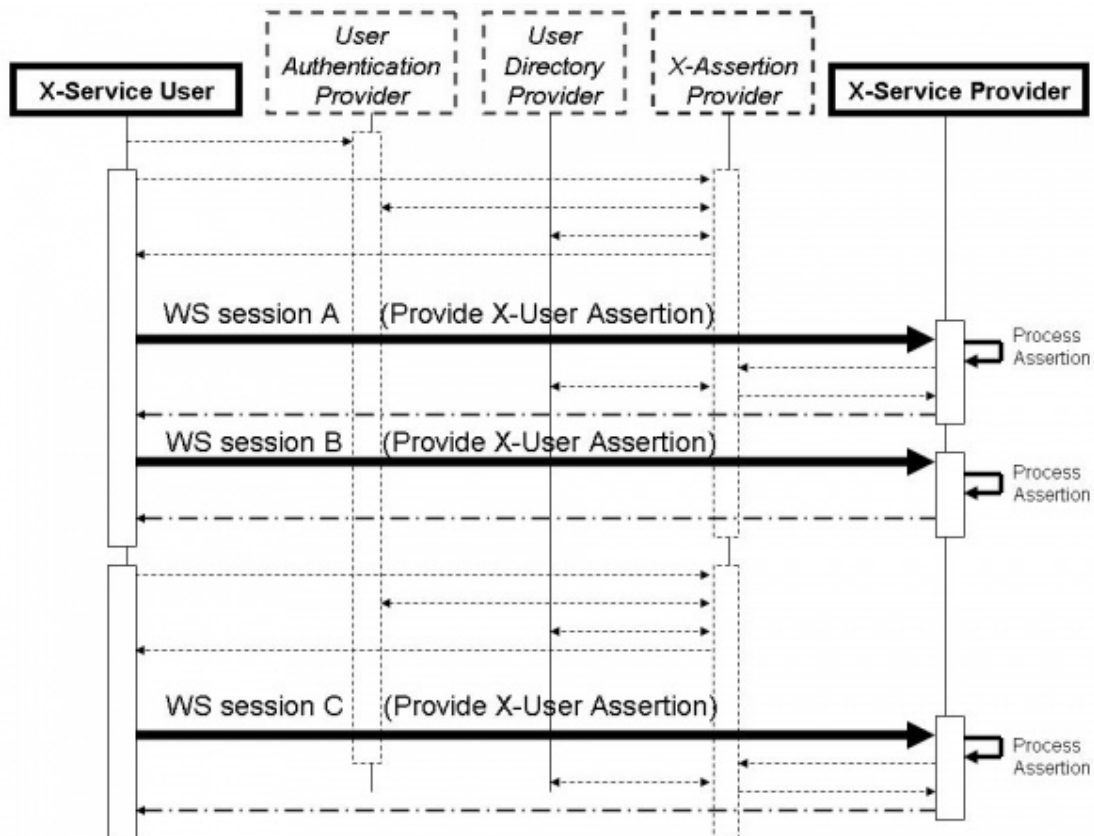
urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos

3455 This conversion from one security token format to another is documented in the WS-Trust standard, and not further profiled by IHE.

13.6.4 Any Web-Services Transaction that leverages ITI TF-2x: Appendix V

3460 Any Actor that uses Web-Services according to ITI TF-2x: Appendix V may be grouped with the appropriate XUA Actors. The Actor grouped with X-Service User Actor, the Requesting Actor, shall conform to all the requirements in the Provide X-User Assertion Transaction. The method used may be through internal means, SAML 2.0 Core protocols, WS-Trust, or any other means. The actor grouped with the X-Service Provider Actor shall conform to all the requirements in the Provide X-User Assertion Transaction. The XUA Profile does not constrain how the Assertion can be used (e.g., ignored, access control, etc.).

13.7 Process Flow



3465

Figure 13.6-1: Cross-Enterprise User Assertion Process Flow

3470 In the above flow we are showing more actors than are specified in this profile. This is a diagram showing a possible grouping with IHE-EUA (User Authentication Provider), IHE-PWP (User Directory Provider), and a SAML Identity Provider (X-Assertion Provider). The User Authentication Provider, User Directory Provider and X-Assertion Provider are not profiled here, but rather are shown to give a context to the XUA transactions.

3475 In this figure the dark lines represent the X-User Assertion transaction. The dashed lines represent other standards based transactions that may be used. Web-Services session A and B show an example where one X-User Assertion is used to cover two Web-Services transactions, where Web-Services Session C is using a different X-User Assertion. This may be due to a different user, timeout of the previous X-User Assertion, or some other reason.

13.8 Security Considerations

3480 The security risk assessment for XUA enumerates assets, threats, and mitigations. The security risk assessment for the Actors that are grouped (e.g., Registry Stored Query and Retrieve Document Set) with the XUA Actors are out of scope of the XUA Profile, please look at those transactions for the Security Considerations. The complete risk data are stored and available from IHE. The purpose of this risk assessment is to notify vendors and healthcare providers of some of the risks that they are advised to consider in implementing XUA Actors. For general IHE risks and threats, please see ITI TF-1: Appendix L. The vendor is also advised that many risks cannot be mitigated by the IHE profile and instead responsibility for mitigation is transferred to the vendor, and occasionally to the affinity domains, individual enterprises and implementers. In these instances, IHE fulfills its responsibility to notify affected parties through the use of the following sections.

3490 **14 Patient Administration Management (PAM) Integration Profile**

14.1 Patient Administration Management Use Cases

The Patient Administration Management Integration Profile defines transactions based on message exchanges to support patient identity and encounter information, as well as movements within an acute care encounter. These can be represented by the following use cases.

3495 **14.2 Patient Identity Management Use Case**

3500 A Patient Registration application decides to create a new patient John Smith, based on patient information input from Hospital Sun. At this time, however, there is a limited set of personal information traits of John Smith available. His date of birth, home address, and home phone number are unknown. The registration application creates the patient identity and sends a Patient Creation message to its downstream applications with the set of known personal information traits.

The next day, detailed personal information about John Smith becomes available. The registration application updates its patient identity record, and sends out a Patient Update message.

3505 After a week, the registration application creates a temporary patient identity John Doe based on input from Imaging Center Moon. After reconciliation of the temporary patient, it updates John Doe's demographics to (a new instance of) John Smith, and changes the temporary Patient Identifier originally assigned to a permanent identifier.

3510 After human inspection, it turns out that these two identities of John Smith represent the same person. The operator decides to merge the second identity to the previously established identity John Smith. A Patient Merge is communicated downstream.

14.2.1 Patient Encounter Management Use Case

3515 Patient Alan Alpha arrives for an annual exam at a clinic. The registration system sends the patient registration information to the local ancillary systems, and the affiliated hospital's ADT system.

3520 The exam of Alan Alpha reveals a serious condition, and an immediate hospital admission is recommended. Alan Alpha is referred to the affiliated hospital for admission. He is pre-admitted in the hospital for relevant diagnostic tests. The tests confirm the condition, and the patient is admitted in the hospital's ICU. During the stay in the ICU, the patient's insurance is verified, and the updated information is sent from the hospital's ADT system to the hospital's ancillary systems.

3525 After a day in the ICU, Alan Alpha's condition has improved, and he is transferred to a regular bed. The nurse recording the transfer makes a mistake, and enters the wrong room and bed. After discovering the error, the transfer is canceled, and the correct transfer is recorded. The patient is now recovered and about to leave the hospital. According to the hospital's procedures, he is

transferred to an outpatient unit for administering follow-up tests. The patient is registered in the Hospital Outpatient Registration System.

The outpatient encounter of Alan Alpha is completed; based on satisfactory test results, he is discharged from the hospital and the Outpatient Registration system.

3530 In this use case, two patient encounter management systems (the hospital ADT system and the hospital Outpatient Registration system) cooperate as peers.

14.2.2 Actors/ Transactions

3535 Figure 14.1-1 shows the actors directly involved in the Patient Administration Management Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved because of their participation in other IHE Integration Profiles, such as Radiology Scheduled Workflow, Patient Identity Cross-Referencing Integration Profiles, etc., are not shown.

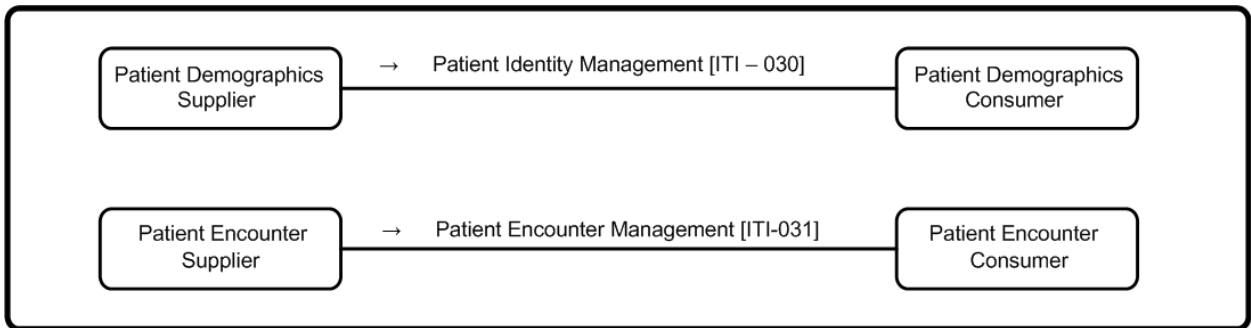


Figure 14.1-1: Patient Administration Management Actor Diagram

3540 Table 14.2-1 lists the transactions for each actor directly involved in the Patient Management Integration Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). A complete list of options defined by this Integration Profile that implementations may choose to support is listed in Table 14.2-1.

3545 **Table 14.2-1: Patient Administration Management - Actors and Transactions**

Actors	Transactions	Optionality	Section
Patient Demographics Supplier	Patient Identity Management [ITI-30]	R	ITI TF-2b: 3.30
Patient Demographics Consumer	Patient Identity Management [ITI-30]	R	ITI TF-2b: 3.30
Patient Encounter Supplier	Patient Encounter Management [ITI-31]	R	ITI TF-2b: 3.31
Patient Encounter Consumer	Patient Encounter Management [ITI-31]	R	ITI TF-2b: 3.31

14.3 Patient Administration Management Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 14.3-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

3550

Table 14.3-1: Patient Administration Management - Actors and Options

Actor	Options	Vol. & Section
Patient Demographics Supplier	<i>Merge (Note 1)</i>	ITI TF-2b: 3.304.1
	<i>Link / Unlink (Note 1)</i>	ITI TF-2b: 3.304.2
	<i>Acknowledgement Support</i>	ITI TF-2b: 3.304.4
	<i>Ambulatory Patient Data</i>	ITI TF-2b: 3.304.5
Patient Demographics Consumer	<i>Merge (Note 1)</i>	ITI TF-2b: 3.304.1
	<i>Link / Unlink (Note 1)</i>	ITI TF-2b: 3.304.2
	<i>Acknowledgement Support</i>	ITI TF-2b: 3.304.4
Patient Encounter Supplier	<i>Inpatient / Outpatient Encounter Management</i>	ITI TF-2b: 3.31.5.2
	<i>Pending Event Management (Note 2)</i>	ITI TF-2b: 3.31.5.3
	<i>Advanced Encounter Management</i>	ITI TF-2b: 3.31.5.4
	<i>Temporary Patient Transfer Tracking</i>	ITI TF-2b: 3.31.5.5
	<i>Historic Movement</i>	ITI TF-2b: 3.31.5.6
	<i>Acknowledgement Support</i>	ITI TF-2b: 3.31.5.7
	<i>Maintain Demographics</i>	ITI TF-1:14.3.9 ITI TF-2b:3.31.5.8
	<i>Ambulatory Patient Data</i>	ITI TF-2b: 3.31.5.8
Patient Encounter Consumer	<i>Inpatient / Outpatient Encounter Management</i>	ITI TF-2b: 3.31.5.2
	<i>Pending Event Management (Note 2)</i>	ITI TF-2b: 3.31.5.3
	<i>Advanced Encounter Management</i>	ITI TF-2b: 3.31.5.4
	<i>Temporary Patient Transfer Tracking</i>	ITI TF-2b: 3.31.5.5
	<i>Historic Movement</i>	ITI TF-2b: 3.31.5.6
	<i>Acknowledgement Support</i>	ITI TF-2b: 3.31.5.7
	<i>Maintain Demographics</i>	ITI TF-1:14.3.9 ITI TF-2b:3.31.5.8

Note 1: An IHE National Extension shall select at least one of the Merge and Link / Unlink Options, and shall mandate the same option for both the Patient Demographics Supplier and the Patient Demographics Consumer implementations in its realm to ensure interoperability.

3555

Note 2: The Pending Event Management Option depends on the Inpatient / Outpatient Encounter Management Option. An implementation supporting the Pending Event Management Option must also support the Inpatient / Outpatient Encounter Management Option.

3560 The Patient Encounter Consumer and Patient Encounter Supplier shall support the Basic Subset of messages defined in ITI TF-2b:3.31.5.1.

The Patient Demographics Consumer and Patient Demographics Supplier shall support at least one of the Merge Option or the Link/Unlink Option. These options define additional messages that shall be supported by these actors.

3565 The PAM Profile also offers a large number of options to support the exchange of patient demographic and encounter data in a wide variety of environments. Particularly, this profile addresses both acute care settings and ambulatory healthcare organizations. It is unlikely that one particular environment will need all the options.

3570 On one hand, an ambulatory care community might need only the pair of actors Patient Demographics Supplier/Patient Demographics Consumer, using transaction ITI-30. On the other hand, the exchange of patient demographic and encounter data between a hospital patient administration system and its ancillary systems (laboratory, radiology, cardiology, etc.) might be fully satisfied, for example, with the pair of actors Patient Encounter Supplier/Patient Encounter Consumer, using transaction ITI-31 with the Basic Subset of messages and the two options “Inpatient/Outpatient Encounter Management” and “Maintain Demographics.”

3575 Hence, the first decision that must be made by a healthcare organization for the deployment of this profile is to select the proper actors and the appropriate set of options to cover its needs, ensuring that each selected option will be supported by the actors on both ends of the transactions.

3580 Furthermore, as an IT Infrastructure profile, the PAM Profile may not be used standalone. Rather, its actors and transactions will be leveraged by other domain integration profiles (in radiology, cardiology, laboratory, or in cross enterprise document sharing). Here again, the first decision that will be taken by the IHE committee that wishes to leverage PAM for its domain, will be to select the proper set of options and to ascertain the consistent use of these options in its domain.

3585 Thus, during the building process of IHE domain technical frameworks, as well as in the deployment process, the PAM Profile will be constrained to reduce its original number of options.

3590 However, to accommodate situations in which a consumer application would not support an option implemented by a supplier application, the PAM Profile requires that the consumer application shall application-reject a message that it does not support (see ITI TF-2x: C.2.3).

14.3.1 Merge Option

The Merge Option defines the information exchange needed to manage the merging of patient identifiers.

14.3.2 Link / Unlink Option

3595 The Link / Unlink Option defines the information exchanges needed to manage the linking and unlinking of patient identifiers, respectively.

14.3.3 Inpatient / Outpatient Encounter Management Option

3600 The Inpatient / Outpatient Encounter Management Option extends the basic patient encounter management functions by defining the information exchanges needed for pre-admitting a patient and for transferring a patient from one location to another location in the enterprise, as well as for changing patient class.

14.3.4 Pending Event Management Option

3605 The Pending Event Management Option extends the basic patient encounter management functions by defining the information exchanges needed for supporting pending events, e.g., admission, transfer, and discharge.

14.3.5 Advanced Encounter Management Option

The Advanced Encounter Management Option extends the basic patient encounter management functions by defining a set of messages for handling patient temporary absence, changing attending doctor in an encounter, and moving accounts among different patient identities.

14.3.6 Temporary Patient Transfer Tracking Option

3610 The Temporary Patient Transfer Tracking Option defines the information exchange needed for tracking a temporary leave / return of a patient from / to a care facility.

14.3.7 Historic Movement Option

3615 The Historic Movement Option extends the basic patient encounter management functions, as well as the following Options:

- Inpatient / Outpatient Encounter Management
- Pending Event Management
- Advanced Encounter Management Options

3620 The Historic Movement Option provides a means to uniquely identify any movement event conveyed in the underlying information exchange. This enables updates of such events at any later time point after they were initially reported.

14.3.8 Acknowledgement Support

An actor that claims support for the Acknowledgement Support Option shall be capable of using the enhanced acknowledgement mode as defined in the HL7 v2.x standard.

3625 **14.3.9 Maintain Demographics Option**

The Maintain Demographics Option extends patient encounter management functions by defining a set of messages for merging patient identifiers and updating patient information in the context of a particular encounter, using transaction ITI-31.

14.3.10 Ambulatory Patient Data Option

3630 If the Patient Demographics Supplier supports the Ambulatory Patient Data Option, it SHALL be capable of providing the patient address.

If the Patient Encounter Supplier supports the Ambulatory Patient Data Option, it SHALL be capable of providing the patient address, the patient ambulatory status and the referring doctor.

See ITI TF-2b: 3.30.4.5 and ITI TF-2b:3.31.5.8.

3635

14.4 Patient Administration Management Integration Profile Actor Grouping

14.4.1 Actor Grouping of Patient Encounter Supplier

3640 In order to obtain patient identity and demographics information to serve its patient encounter message functions in transaction ITI-31, a Patient Encounter Supplier is required to be grouped with either a Patient Demographics Supplier or a Patient Demographics Consumer, as shown in Figure 14.4-1.

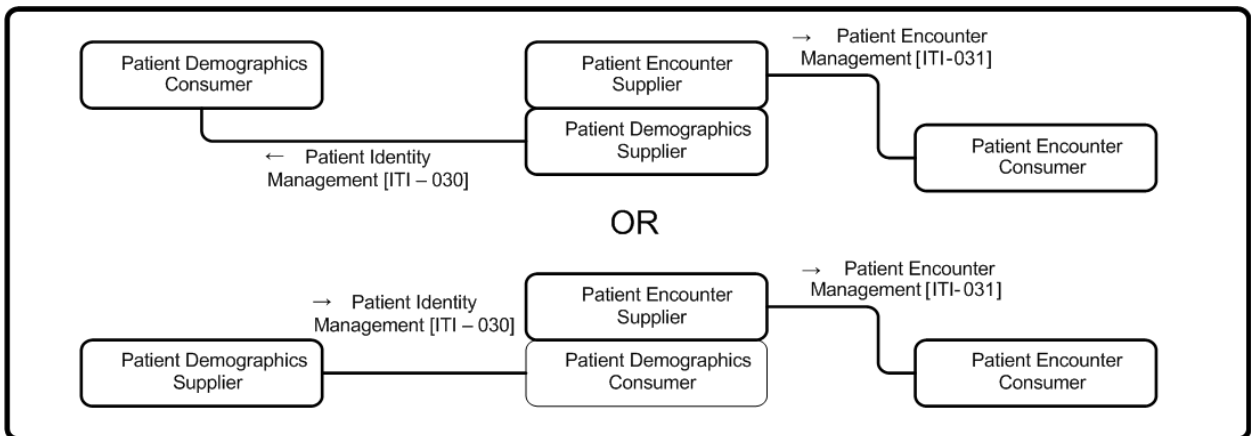


Figure 14.4-1: Patient Encounter Supplier Grouping Requirements

3645 On the other hand, transaction ITI-31 is self-contained in a sense that the Patient Encounter Supplier sends both patient encounter information and patient identity and demographics information (in the context of the encounter data) to the Patient Encounter Consumer. In addition, transaction ITI-31 also allows the Patient Encounter Supplier to send messages to the Patient Encounter Consumer for patient identity maintenance in the encounter context, including

3650 patient update and identity merge. There is no required grouping for the Patient Encounter Consumer.

14.4.2 Actor Grouping with other IHE Actors

3655 The PAM Profile provides an infrastructure in a healthcare enterprise or across a number of enterprises to distribute the patient identity, demographics, and encounter information, in order to enable various clinical functions in clinical settings. The PAM actors can be grouped with actors in other IHE Integration Profiles.

3660 One possible grouping is between the Patient Demographics Supplier Actor in the PDQ Profile and either the Patient Demographics Supplier Actor or the Patient Demographics Consumer Actor in this profile, to add query support defined in the Patient Demographics Query transaction to the same set of patient information managed in the PAM Profile.

3665 Furthermore, the Patient Demographics Supplier Actor in the PDQ Profile can be grouped with the Patient Encounter Supplier Actor of this profile. Due to the required grouping of the Patient Encounter Supplier Actor (see ITI TF-1: 14.4.1), such a grouping can provide query support defined in both the Patient Demographics Query and Patient Demographics and Visit Query transactions to the same set of patient and encounter information that is managed in the PAM Profile.

3670 These are some examples of possible grouping of the PAM actors with other IHE actors. Many other possibilities may be useful (either to provide additional values or to allow profile structure simplification). For example, in the radiology scheduled workflow (SWF) Profile, the Order Placer and Order Filler actors can be grouped with the Patient Encounter Consumer Actor.

14.5 Patient Administration Management Process Flow

14.5.1 Patient Identity Management

The Patient Identity Management incorporates the following process flows. This refines the use case shown in ITI TF-1: 14.1.1.

3675 14.5.1.1 Patient Identity Creation and Maintenance

- 3680 • Create Patient. The Patient Demographics Supplier decides to create a new patient John Smith, based on patient information input from Hospital Sun. At this time, however, there is a limited set of personal information traits of John Smith available. His date of birth, home address, and home phone number, e.g., are unknown. The Patient Demographics Supplier creates the patient identity and sends a Patient Creation message to the Patient Demographics Consumer with the set of known personal information traits.
- 3685 • Update Patient Demographics. The next day, detailed personal information about John Smith becomes available. The Patient Demographics Supplier updates its patient identity record, and sends out a Patient Update message, including date of birth, home address and home phone number.

- Create Temporary Patient. After a week, the Patient Demographics Supplier creates a temporary patient identity John Doe based on input from Imaging Center Moon.
- 3690 • Update Patient Demographics and Change Patient Identifiers After reconciliation of the temporary patient, the Patient Demographics Supplier updates John Doe’s demographics to (a new instance of) John Smith, and changes the temporary Patient Identifier originally assigned to a permanent identifier
- 3695 • Merge Patient Identifiers. After human inspection, it turns out that the two patients named John Smith in the Patient Demographics Supplier actually represent the same real-world patient. The operator decides to merge the two patient identities. The Patient Demographics Supplier sends a Patient Merge message to the Patient Demographics Consumer.

The following diagram shows the process flow:

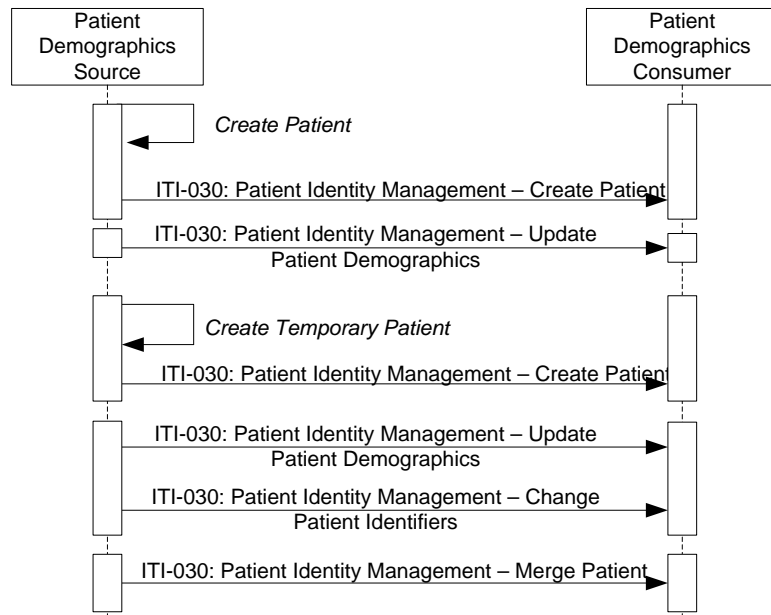


Figure 14.5-1: Patient Identity Management Process Flow in PAM Profile

3700 **14.5.1.2 Alternative Process Flow**

- Link Patient Identifiers. A similar situation as that mentioned above, except that the local procedures request the Patient Demographics Supplier to link these two duplicated patient records instead of merging them. The operator performs the link function. The Patient Demographics Supplier sends a Patient Identifiers Link message to the Patient Demographics Consumer.

3705

The following diagram shows the alternate portion of the process flow:




Figure 14.5-2: Patient Identity Management Alternate Process Flow in PAM Profile

3710

14.5.2 Patient Encounter Management

The Patient Encounter Management incorporates the following process flows:

14.5.2.1 Inpatient/Outpatient Encounter, Maintain Demographics, and Pending Event Management

3715 In this section, inpatient/outpatient encounter management process flow is described in an environment that involves a number of instances of Patient Encounter Supplier and Patient Encounter Consumer. This refines the use case shown in ITI TF-1: 14.1.2

3720 In some institutions, there may be one central Patient Encounter Supplier, while others may have multiple Patient Encounter Suppliers serving patient encounter management functions in different clinical settings (e.g., hospital inpatient, hospital outpatient, clinics). It is the responsibility of a healthcare institution to define the actor roles of its systems, as well as to configure the relationship of a Patient Encounter Supplier and its Patient Encounter Consumers, to satisfy their business process models.

3725 As shown in Figure 14.5-3, in the healthcare institution of this process flow, there are three Patient Encounter Suppliers, each of which serves a number of Patient Encounter Consumers in a specific clinical setting of the institution.

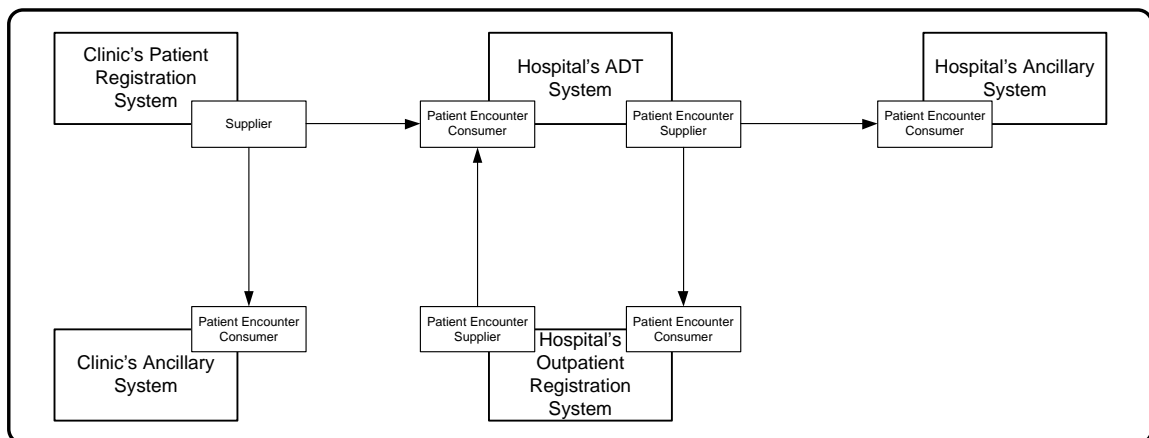


Figure 14.5-3: System and PAM Actor Role Configuration

The systems involved in this process flow implement the following PAM roles:

- 3730 • Clinic Registration System as Patient Encounter Supplier

- Clinic Ancillary System as Patient Encounter Consumer
- Hospital ADT system as both Patient Encounter Supplier and Patient Encounter Consumer
- Hospital Ancillary system as Patient Encounter Consumer
- 3735 • Hospital Outpatient Registration System as both Patient Encounter Supplier and Patient Encounter Consumer

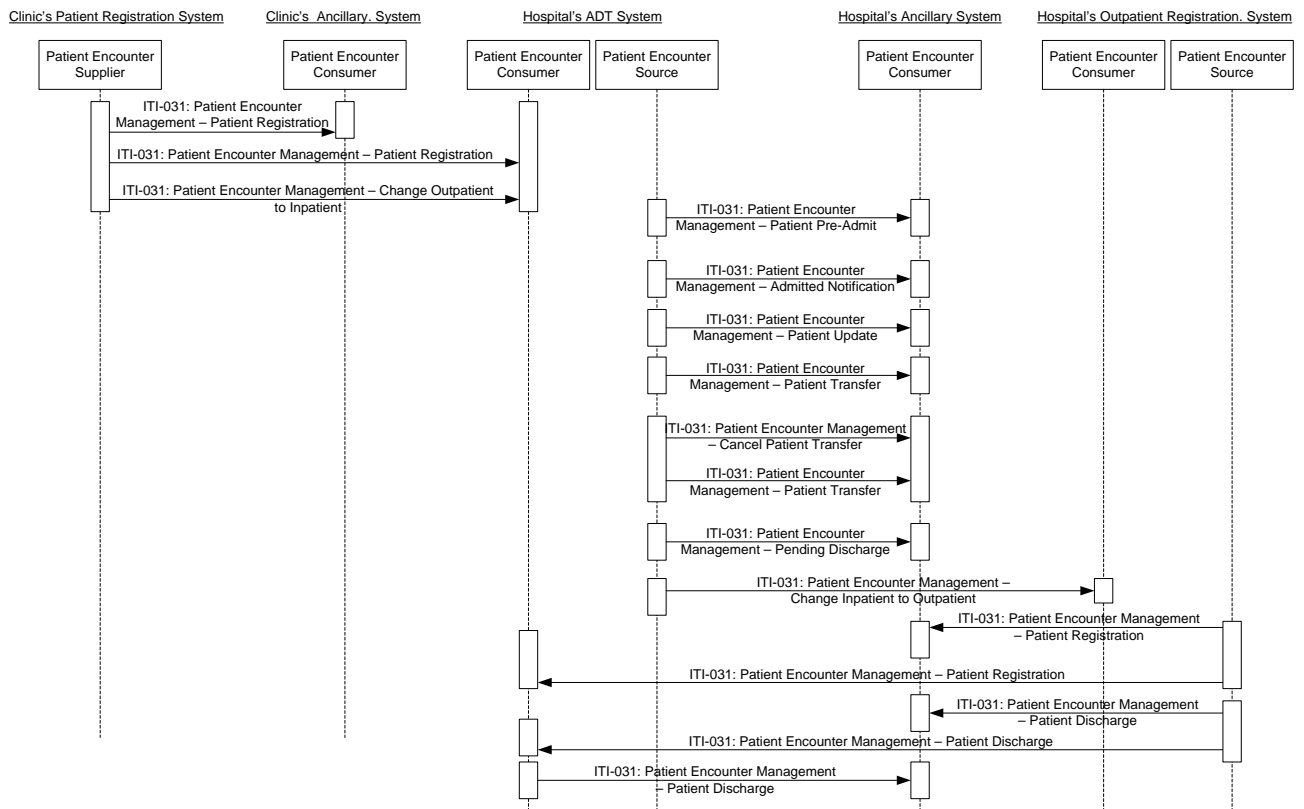
3740 Note that the Hospital ADT and Outpatient Registration Systems play both the roles of Patient Encounter Supplier and Patient Encounter Consumer, and cooperate as peers. The relationship between the Patient Encounter Supplier and Patient Encounter Consumer in the same system is dependent on the clinical application logic implemented in the institution, and the definition of this relationship is beyond the scope of the PAM Integration Profile.

The process flow in Figure 14.5-4 is described in the following:

- 3745 • **Patient Registration:** A patient arrives for an annual exam at a clinic. The patient record has been created previously by a Patient Demographics Supplier, and exists in the clinic's registration system through its grouping with the Patient Demographics Supplier Actor. The clinic's registration system sends the Patient Registration message to the local ancillary systems, and the affiliated hospital's ADT system.
- 3750 • **Change Outpatient to Inpatient:** The exam reveals a serious condition of the patient, and an immediate hospital admission is recommended. The patient is referred to the affiliated hospital for admission. A Change Outpatient to Inpatient message is sent to the hospital's ADT System.
- 3755 • **Pre-admit Patient for Hospitalization:** The patient is pre-admitted in the hospital for relevant diagnostic tests. The hospital ADT system sends Patient Pre-Admit message to the Hospital Ancillary System.
- 3760 • **Patient Admitted Notification:** The tests confirm the condition, and the patient is admitted to the hospital's ICU. The hospital ADT system sends an Admission Notification message to the Ancillary System.
- **Patient Insurance Information Update:** During the stay in the ICU, the patient's insurance is verified, and the updated information is sent from the hospital ADT to the Hospital Ancillary System.
- 3765 • **Patient Location Transfer:** After a day in the ICU, the patient's condition has improved, and the patient is transferred to a regular bed. The hospital ADT system sends a Patient Transfer message to the Hospital Ancillary System.
- **Patient Location Transfer Error Reconciliation:** The nurse recording the transfer makes a mistake, and enters the wrong room and bed. After discovering the error, the hospital ADT system sends a Cancel Patient Transfer message to the Hospital Ancillary System, followed by a new Patient Transfer message.

- 3770 • **Patient Pending Discharge:** The patient is now recovered and about to leave the hospital. The ADT system sends a Patient Pending Discharge message to the Hospital Ancillary System.
- **Change Inpatient to Outpatient:** According to the hospital’s procedures, the patient is transferred to an outpatient unit for administration of follow-up tests. The ADT system sends a Change Inpatient to Outpatient message to the Hospital Outpatient Registration System.
- 3775 • **Register Patient as Outpatient:** The patient is registered in the Hospital Outpatient Registration System, which sends a Patient Registration message to the Hospital ADT system and the Hospital Ancillary System.
- **Patient Discharged from Outpatient System:** The outpatient encounter is completed. A Patient Discharge message is sent to the Hospital ADT System and to the Hospital Ancillary System.
- 3780 • **Patient discharged from Hospital ADT System:** Based on satisfactory test results, the patient is discharged. The hospital ADT system sends a Patient Discharge message to the Hospital Ancillary System.

The following diagram shows the process flows of the discussed use cases:



3785

Figure 14.5-4: Inpatient / Outpatient Encounter Management Process Flow in PAM Profile

14.5.2.2 Advanced Encounter Management

- 3790 • **Attending Physician Change:** A patient's attending physician changes during an inpatient stay. The Patient Encounter Supplier sends a notification message that contains the name of the new attending doctor to the Patient Encounter Consumer.
- 3795 • **Cancellation of Attending Physician Change:** A notification of change of a patient's attending physician was sent in error. The Patient Encounter Supplier sends a cancellation message that contains the name of the old attending doctor to the Patient Encounter Consumer.
- 3800 • **Leave of Absence:** An inpatient is authorized a weekend leave of absence from the medical center. The Patient Encounter Supplier sends a notification message to the Patient Encounter Consumer that contains the date and time of the leave of absence and of the expected return.
- 3805 • **Cancellation of Leave of Absence:** A notification that an inpatient was authorized a weekend leave of absence was sent in error. The Patient Encounter Supplier sends a cancellation message to the Patient Encounter Consumer.
- 3810 • **Return from Leave of Absence:** An inpatient returns to the medical center from a weekend leave of absence. The Patient Encounter Supplier sends a notification message to the Patient Encounter Consumer that contains the date and time of the expected return and of the actual return.
- 3815 • **Cancellation of Return from Leave of Absence:** A notification that an inpatient returned from a weekend leave of absence was sent in error. The Patient Encounter Supplier sends a cancellation message to the Patient Encounter Consumer.
- 3815 • **Move Account:** The Patient Encounter Supplier sends a message that incorrectly associates Account 12345 with Patient A; in fact, Account 12345 should be associated with Patient B. To effect a correction, the Patient Encounter Supplier sends a message to the Patient Encounter Consumer that contains the account identifier and the identifiers of the patient records between which the account association is to be moved.

The following diagram shows these discussed use cases:

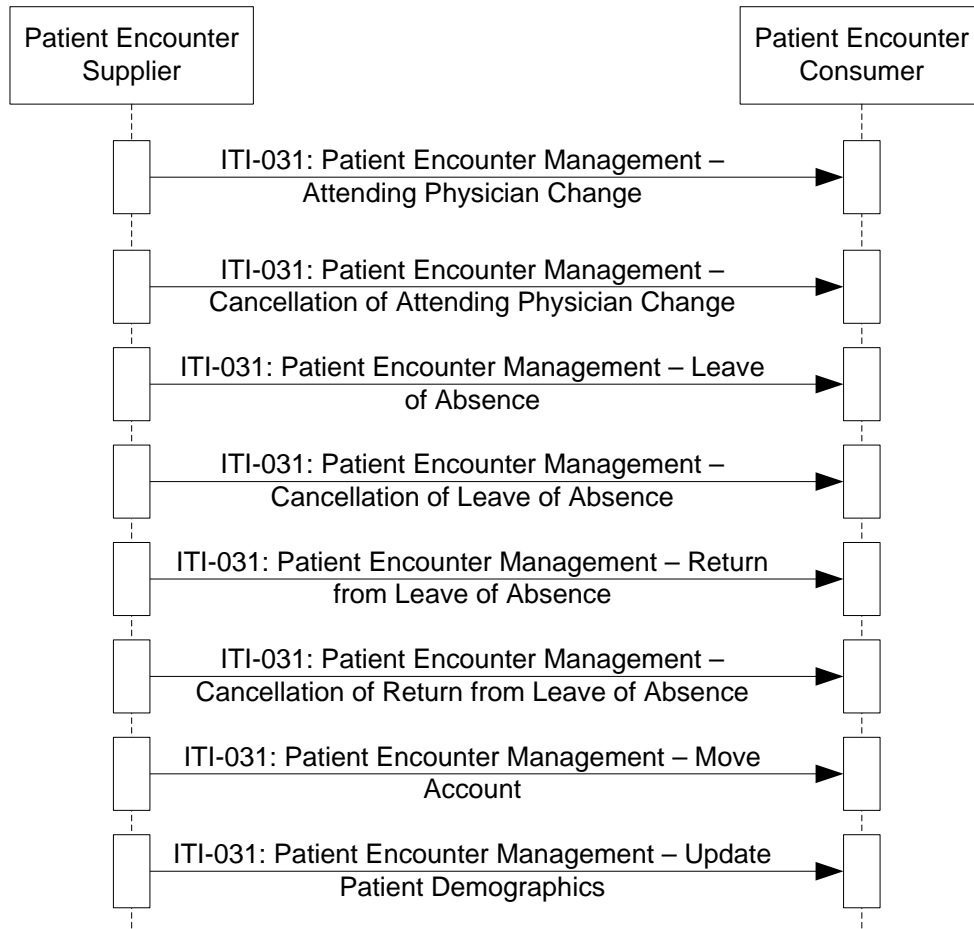


Figure 14.5-5: Advanced Encounter Management Process Flow in PAM Profile

3820 **14.5.2.3 Historic Movement Management**

Historic tracking of patient admissions, discharges, and other movements may be needed in some healthcare institutions. Such historic events may need to be tracked even beyond the boundary of an episode of care. In order to facilitate this tracking, the Patient Encounter Supplier may send the messages in ITI TF-1: 14.5.2.1 and 14.5.2.2 to the Patient Encounter Consumer, with the addition of an identifier for the particular encounter with which the patient admission, discharge, or movement is associated.

3825

- **Patient Location Transfer:** A patient is transferred to bed 23 of Room B after a few days of stay in ICU. The hospital ADT system sends a Patient Transfer message (including the elements provided in the Historic Movement Management Option) to the downstream applications.

3830

- **Update Previous Transfer Event.** After two days, the operator of the ADT system detects that the transfer destination and time in the previously sent Patient Transfer message were

wrong. He corrects the errors and an Update Historic Patient Transfer message is sent out, to communicate the true room / bed information and the true transfer time.

3835 The following diagram shows these use cases:

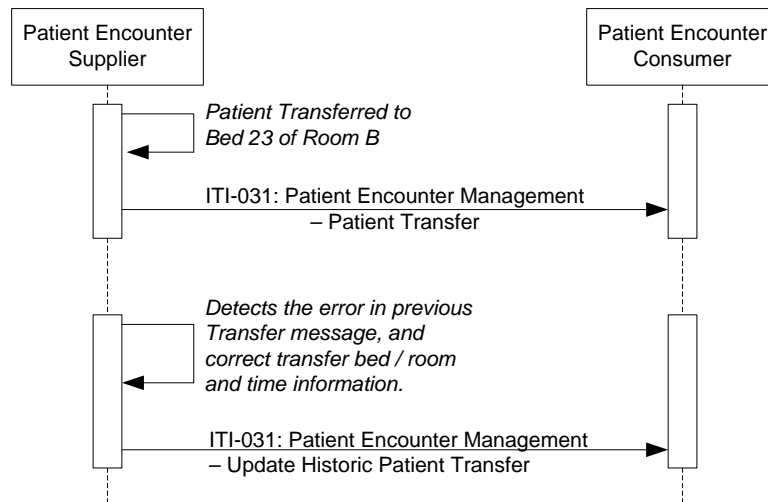
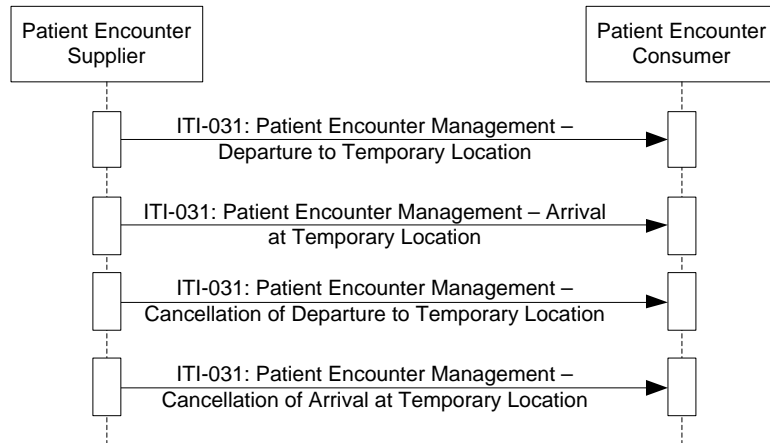


Figure 14.5-6: Historic Movement Management Process Flow in PAM Profile

14.5.2.4 Temporary Patient Transfer Tracking

- 3840
- **Departure to Temporary Location:** A chest X-ray is scheduled for an inpatient. To perform this service, the patient needs to be moved from her inpatient bed in the medical service to the Radiology department. When the patient departs from her inpatient bed, the Patient Encounter Supplier sends a notification message to the Patient Encounter Consumer that contains the temporary location to which the patient is being moved.
- 3845
- **Arrival at Temporary Location:** When the patient arrives at the Radiology department, the Patient Encounter Supplier sends a notification message to the Patient Encounter Consumer that contains the temporary location to which the patient has been moved.
 - **Cancellation of Departure to Temporary Location:** It is incorrectly communicated that a patient left her inpatient bed to move to the Cardiology department for treatment. The Patient Encounter Supplier sends a cancellation message to the Patient Encounter Consumer that contains the patient's location(s) (permanent and / or temporary) prior to the time of the erroneously communicated departure.
- 3850
- **Cancellation of Arrival at Temporary Location:** It is incorrectly communicated that a patient, having left her inpatient bed, arrived in the Surgery department for treatment. The Patient Encounter Supplier sends a cancellation message to the Patient Encounter Consumer that contains the patient's location(s) (permanent and / or temporary) prior to the time of the erroneously communicated arrival.
- 3855

The following diagram shows these discussed use cases:



3860

Figure 14.5-7: Temporary Patient Transfer Tracking Process Flow in PAM Profile

15 XDR Integration Profile

3865 Cross-Enterprise Document Reliable Interchange (XDR) provides document interchange using a reliable messaging system. This permits direct document interchange between EHRs, PHRs, and other healthcare IT systems in the absence of a document sharing infrastructure such as XDS Registry and Repositories.

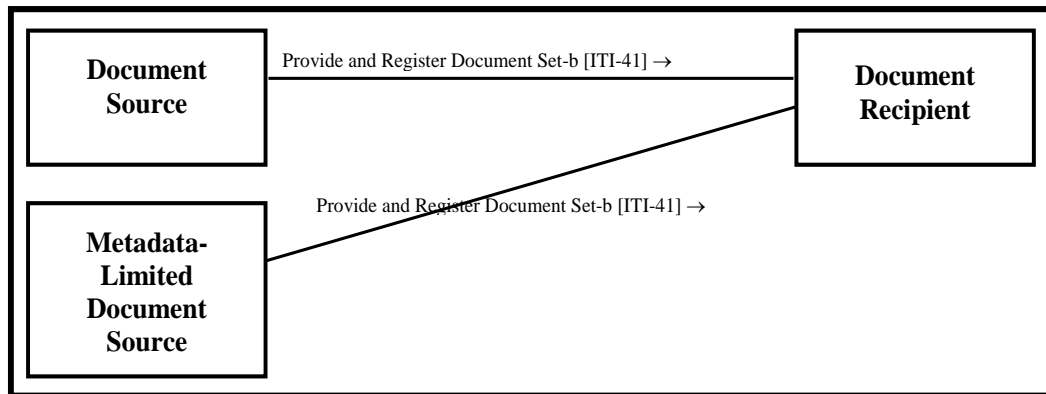
XDR provides a reliable and automatic transfer of documents and metadata for one patient between EHR systems even in the absence of an XDS infrastructure. XDR supports the reuse of the Provide and Register Set transaction-b with Web-Services as transport. Transfer is direct from source to consumer, no repository or registry actors are involved.

3870 XDR is document format agnostic, supporting the same document content as XDS and XDM. Document content is described in XDS Document Content Profiles. Examples are XDS-MS, XD-LAB, XPHR, and XDS-SD.

XDR defines no new metadata or message formats. It leverages XDS metadata with emphasis on patient identification, document identification, description, and relationships.

3875 15.1 Actors/ Transactions

Figure 15.1-1 shows the actors directly involved in the XDR Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in XDS, PIX or XUA are not shown.



3880

Figure 15.1-1: XDR Actor Diagram

3885 Table 15.1-1 lists the transactions for each actor directly involved in the XDR Profile. In order to claim support of this Integration Profile with one or more actors, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in Volume 1, Section 15.2.

Table 15.1-1: XDR Integration Profile - Actors and Transactions

Actors	Transactions	Optional ity	Section in Vol. 2
Document Source	Provide and Register Document Set-b [ITI-41]	R	ITI TF-2b:3.41
Metadata-Limited Document Source	Provide and Register Document Set-b [ITI-41]	R	ITI TF-2b:3.41
Document Recipient	Provide and Register Document Set –b [ITI-41]	R	ITI TF-2b:3.41

3890

15.2 XDR Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 15.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

3895

Table 15.2-1: XDR - Actors and Options

Actor	Options	Vol. & Section
Document Source	<i>Basic Patient Privacy Enforcement</i>	ITI-TF-2b: 3.41.4.1.3.1
Metadata-Limited Document Source	<i>Basic Patient Privacy Enforcement</i>	ITI TF-2b: 3.41.4.1.3.1
Document Recipient	<i>Basic Patient Privacy Enforcement</i>	ITI-TF-2b: 3.41.4.1.3.1
	<i>Accepts Limited Metadata</i>	ITI TF-1:15.2.3

15.2.1 Intentionally Left Blank

15.2.2 Basic Patient Privacy Enforcement Option

For this option, see ITI-TF-2b: 3.41.4.1.3.1

15.2.3 Accepts Limited Metadata

3900

When the Document Recipient declares this option it will accept metadata entries from a Metadata-Limited Document Source which use the less rigorous metadata attribute requirements as shown in ITI TF-2b:3.41.4.1.2 Table 3.41.4.1.2-2.

15.3 XDR Process Flow

3905

XDR describes the exchange of a set of a patient's documents between healthcare providers, such as: physicians, hospitals, special care networks, or other healthcare professionals.

Where XDS Registry/Repositories are not yet implemented or available for the exchange of information, XDR is the viable approach.

3910 In a situation where the information is going to an automated application or robust system capable of automated storage or processing of documents relative to one patient, XDR is the appropriate profile.

The XDR Integration Profile is intended only for exchange of patient related medical documents and not intended to address all cross-enterprise EHR communication needs.

Use Cases:

3915 1. Dr. Primary refers his aging patient Mr. Robinson to his first appointment with a gastroenterology specialist.

3920 Since there is no XDS repository available at the gastro clinic, Dr. Primary cannot use XDS to communicate the XDS-MS referral to Dr. Gastro. Also, since there is no affinity domain linking Dr. Primary and Dr. Gastro, XDR is preferable to XDS for the exchange of Mr. Robinson's referral information. XDR is also appropriate for Dr. Gastro's documents communication to Dr. Primary.


3925 2. Mabel is transferred from a hospital setting to her retirement home for long-term care.

XDR: Mabel's information can be transferred from the hospital to the long-term care facility's EHR application for future review by her attending physicians and nurses, through XDR.

3930 3. Stanley's recent MRI has generated unusual results that Stanley's primary physician would like to consult with another specialist in a specialized cancer facility located across the state. Since there is not likely to be an affinity domain between the remote health environments, XDR can be used instead.

3935 4. Mrs. Sweettooth has been diagnosed with adult diabetes and her specialized circle of care has not yet gotten organized to provide shared access to a common repository. Until they do, they will need to exchange her information peer-to-peer using XDR.

3935 This profile is only defining the digital transport mechanism used for such use cases, content transported will be detailed by Content Profiles such as the ones defined by the IHE PCC (Patient Care Coordination) domain.



3940

Figure 15.3-1: Process Flow in XDR Profile

15.4 Digital communication

It is a web service based HTTP message.

15.5 Security Considerations

3945 The Profile assumes that the health organizations that are using Document Source and Document
 Recipient have an agreement defining when they can interchange PHI. This may require an
 explicit patient consent (depending on the regulation) and an agreement on how to manage the
 potential inconsistency between the security policies. The main aspects that should be covered by
 this agreement are similar to XDS – See Appendix L. In the case of XDR, the EHR-to-EHR (or
 3950 PHR) communication is a transient XDS Affinity Domain. In addition, the following aspects
 should be covered:

- Management of Patient identification in order to perform patient reconciliation correctly upon importation of the documents.

Both Actors for this Profile require a grouping with Secure Node.

3955 **16 Cross-Enterprise Document Media Interchange (XDM) Integration Profile**

3960 Cross-Enterprise Document Media Interchange (XDM) provides document interchange using a common file and directory structure over several standard media types. This permits the patient to use physical media to carry medical documents. This also permits the use of person-to-person email to convey medical documents. XDM supports the transfer of data about multiple patients within one data exchange.

Physician to patient to physician - Bob has an MRI and cancer is diagnosed. He is given a CD-R with his MRI results and referral information on it to give to the specialist of his choice.

3965 **Patient visiting ED** - In addition, Bob, the informed patient, maintains a copy of his EHR record at home and can bring the CD-R with him when he visits the ED for an unrelated emergency.

Physician to physician - Dr. Primary refers his aging patient Mr. Robinson to his first appointment with a gastroenterology specialist. He transfers relevant documents in a zip file attached to an email to the specialist.

3970 The common thread of these use cases is that they are person-to-person communications. The XDM solution is intended to be easy to implement with pre-existing email clients, CD burners and USB ports. XDM does not include any additional reliability enhancements. XDM requires that the recipient be able to support human intervention in order to manually control the importing of the data (patient ID reconciliation, selection of patient of interest from possibly multiple patients' documents on the media).


3975 XDM is document format agnostic, supporting the same document content as XDS and XDR. Document content is described in XDS Document Content Profiles. Examples are XDS-MS, XPHR, XDS-SD, and XD-LAB.

XDM defines no new metadata. It leverages XDS metadata with emphasis on patient identification, document identification, description, and relationships.

3980 A directory and file structure is documented for populating the media. This structure maintains separate areas for each patient listed and is supported on all referenced media types. Media and the structure were selected based on experience with media interoperability in Radiology, i.e., PDI Profile. The media selected are the widespread CD-R, USB removable media, and email with ZIP attachment.

3985 **16.1 Actors/ Transactions**

Figure 16.1-1 shows the actors directly involved in the XDM Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in XDS, PIX or PDI are not shown.



3990

Figure 16.1-1: XDM Actor Diagram

Table 16.1-1 lists the transactions for each actor directly involved in the XDM Profile. In order to claim support of this Integration Profile with one or more actors, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 16.2.

3995

Table 16.1-1: XDM Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Portable Media Creator	Distribute Document Set on Media [ITI-32]	R	ITI TF-2b: 3.32
Portable Media Importer	Distribute Document Set on Media [ITI-32]	R	ITI TF-2b: 3.32

16.2 XDM Integration Profile Options

4000

Options that may be selected for this Integration Profile are listed in Table 16.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 16.2-1: XDM - Actors and Options

Actor	Options	Vol. & Section
Portable Media Creator	<i>USB (Note 1)</i>	ITI TF-1: 16.2.1
	<i>CD-R (Note 1)</i>	ITI TF-1: 16.2.2
	<i>ZIP over Email (Note 1)</i>	ITI TF-1: 16.2.3
	<i>Basic Patient Privacy Enforcement</i>	ITI TF-2b: 3.32.4.1.4.1
	<i>Zip over Email Response (Note2)</i>	ITI TF-1: 16.2.4
Portable Media Importer	<i>USB (Note 1)</i>	ITI TF-1: 16.2.1
	<i>CD-R (Note 1)</i>	ITI TF-1: 16.2.2
	<i>ZIP over Email (Note 1)</i>	ITI TF-1: 16.2.3
	<i>Basic Patient Privacy Enforcement</i>	ITI TF-2b: 3.32.4.1.4.1
	<i>Zip over Email Response (Note2)</i>	ITI TF-1: 16.2.4

4005

Note 1: At least one of these options is required for each Actor. In order to enable a better interoperability, is highly recommended that the actors support all the options.

Note 2: This option requires the ZIP over Email Option.

16.2.1 USB Option

4010 In this option the Portable Media Creator writes a set of documents on USB media. The media is physically transported to the Portable Media Importer which then imports the document set.

16.2.2 CD-R Option

In this option the Portable Media Creator writes a set of documents on CD-R media. The media is physically transported to the Portable Media Importer which then imports the document set.

16.2.3 ZIP over Email

4015 In this option the Portable Media Creator creates an ordinary ZIP file of the virtual media containing document set(s). The ZIP file is attached to an Email sent to the Portable Media Importer which then retrieves the Email and imports the ZIP file containing the document set.

16.2.4 ZIP over Email Response

4020 In this option the Portable Media Importer sends a response (MDN Based) to the Portable Media Importer to acknowledge that the Import operation of the Document Set(s) received was successful.

If this option is supported, the ZIP over Email option shall be supported.

16.3 XDM Process Flow

4025 XDM describes the exchange of a set of a patient's documents between healthcare providers, such as: physicians, hospitals, special care networks, or other healthcare professionals.

Where XDS is not desirable or available for one of the participants in the exchange of information, XDM is a viable option.

4030 XDM should be used in a situation where the information receiver is an individual who will manually interpret or examine the data and associated documents as though they were using physical media. XDM also allows for the exchange of documents relating to multiple patients, since the data will be interpreted manually by human intervention.

4035 The XDM Integration Profile is intended only for exchange of personal medical documents and not intended to address all cross-enterprise EHR communication needs. Some use cases may require the use of other IHE integration profiles such as XDS, DSG, PIX, and ATNA. Other use cases may only be partially supported, while still others may require future IHE integration profiles.

Use Cases:


1. Dr. Primary refers his aging patient Mr. Robinson to his first appointment with a gastroenterology specialist.

4040 In a case where either Dr. Primary's office or Dr. Gastro's clinic was not able to handle secure email, or other sustained online point-to-point communications (e.g., http over VPN), the XDM Profile would provide further solutions for the simpler

environment, such as the use of physical media, or email where the interchanged document set will be manually interpreted by a human intervention.

- 4045 2. In a hospital that does not have an XDS infrastructure; the XDS-MS Content Profile discharge use case can also be handled by XDM. For example:
- In a hospital, or in the case of a family physician not using robust EHR, the patient could be handed a CD or USB media with their discharge information on it to bring with them to their follow-up visit with their family physician.
- 4050 3. Mabel is transferred from a hospital setting to her retirement home for long-term care.
- If the hospital does not have an EHR application that automatically interprets her medical data and shares it with the necessary members of her health team, the information can be transferred manually directly to the file clerk, intake coordinator, records manager, or primary physician depending on the organization's resource model.
- 4055
4. Stanley's recent MRI has generated unusual results that Stanley's primary physician would like to consult with another specialist in a specialized cancer facility located across the state. Since there is not likely to be an affinity domain between the remote health environments, XDM can be used instead.
- 4060 5. Bob, the informed patient, maintains a copy of his Personal Health Record (PHR) at home. In this situation, Bob can be given a copy of his medical information on physical media such as a CD-ROM to take home with him. Bob now has an advantage that he can continue to have his complete medical record available with him on sudden emergency department visits, even when he is on an out-of-state trip where the new ED would have no access to the repository of his home affinity domain.
- 4065

This profile is only defining the digital transport mechanism used for such use cases. Content transported will be detailed by Content Profiles such as the ones defined by the IHE PCC (Patient Care Coordination) domain.



4070 **Figure 16.3-1: Process Flow in XDM Profile**

16.4 Digital communication

16.4.1 Actual Media Type

The media can be either CD-R or a USB media device, because these are the most common media types in other industries for the portable transport of electronic information. This

4075 supplement requires using one of these media types, depending on the use case. The benefit and risks of the reusability of the media deployed should be taken into account, especially when the media is under the control of the patient.

4080 Note: 1. Because the size of documents to be exchanged rarely requires more than the capacity of a CD, and the format for storing data on various different recordable DVD media is not totally stable yet, this profile is following the restriction defined in the IHE RAD PDI Profile, to not use recordable DVD media at this time.

2. CD-RW is excluded from this profile because field experiences with CD-RW in radiology with this media showed significant interoperability problems and significant accidental damage levels.

4085 3. The CD-R media is limited to the 74 minute blanks because the long playing CD-R format gains the larger capacity by eliminating one level of error correction and detection. The resulting much higher undetected error rate is considered unacceptable for medical data.

16.4.2 Virtual Media over a Network

The media can be a ZIP file containing the document set and sent via a secure email message.

16.4.3 Media Content

4090 The requirements for media content are intended to promote the simple transfer of medical documents, including patient summaries, lab results, discharge letters and reports, and to allow for the viewing of such documents on general purpose computers by care providers or patients.

Created media are required to contain documents and the relevant associated metadata.

The media contains one or more Submission Sets including the documents and the associated metadata, organized in a well-defined directory structure starting at the root level.

4095 The media content can be made web viewable by a web browser by providing optional files containing HTML content. This content must be based on the original documents in order to ensure consistency. Any ordinary web browser can be used to read these files. The Portable Media Importer ignores these files. They are just intended for the human recipient.

4100 Additional content may be present (files, directories), but can be ignored by the Portable Media Importer.

To summarize, the Portable Media Importer has two complementary ways to access the media and its content through a basic web browser:

- By inspecting in the directory dedicated to XDM all the subdirectories that contain a specifically named metadata file compatible with XDM
- 4105 • By presenting to the user the HTML index file that lists the submission sets and documents contained in the media.

Access to the content of an individual document is outside the scope of this Integration Profile and shall be addressed in specific IHE document content Integration Profiles.

16.5 Security considerations

4110 The Profile assumes that the Healthcare delivery organizations that are using Portable Media Creator and Importer have an agreement defining when they can interchange PHI. This may require an explicit patient consent (depending on existing regulations) and an agreement on how

4115 to manage the potential inconsistency between the security policies. The main aspects that should be covered by this agreement are similar to XDS – See ITI TF-1: Appendix L. In addition, the following aspects should be covered:

- Management of Patient identification in order to perform patient reconciliation correctly upon importation of the documents.
- Measures taken to avoid or limit loss of media or email, and detect that which occurs.

4120 In the case of physical media, security responsibilities for confidentiality and integrity are transferred to the patient by providing the media to the patient. In this case it is the patient's responsibility to protect the media, and the patient has the authority to disclose the contents of the media as they choose. They disclose the contents by providing the media.

4125 The Portable Media Creator in most cases does not know who the ultimate Importer will be, thus rendering encryption impractical.

In the case of transfer over email using a ZIP attachment, the transaction is secured by the use of S/MIME.

Both Actors for this Profile require a grouping with an ATNA Secure Node or Secure Application.

4130

17 Retrieve Form for Data Capture (RFD) Profile

4135 The Retrieve Form for Data Capture (RFD) Profile provides a method for gathering data within a user's current application to meet the requirements of an external system. RFD supports the retrieval of a form from a form source, the display and completion of the form, and the return of instance data from the display application to a receiving application. In addition, RFD provides a mechanism to amend data that was previously captured.

4140 Consider the case where a healthcare provider site uses an Electronic Health Record (EHR) to document patient care. In this case, the EHR acts as the local home application for the provider's personnel. Suppose an external agency, through some contractual arrangement, requires data from the provider, some of which reside in the EHR's database, the rest requiring data entry by the EHR's users. RFD enables the EHR user to retrieve a data capture form from the external agency, to fill out the form, and to return the data to the external agency without leaving the provider's local home application, the EHR. The profile also permits the external agency to indicate that there is a need to clarify points about the data so captured and provides the mechanisms to allow the data to be modified.

4150 Many potential uses of RFD want the form to dynamically pre-populate forms from the host application's database, that is have the form delivered with host application database values filled in to appropriate fields of a form. RFD permits automatic form population and provides a generic mechanism by which this can be accomplished. However, the profile does not speak to the issue of content, remaining silent on normative vocabularies and other enablers of semantic interoperability. Specific domain groups – clinical trials, drug safety, bio-surveillance – will build on RFD by contributing content specifications or by evaluating and recommending existing content standards that will operate within RFD. When RFD, as an infrastructure profile, integrates with domain-specific content standards, a much greater level of interoperability will result.

4155 The RFD Profile provides a generic polling mechanism to allow an external agency to indicate issues with data that have been captured and enable the healthcare provider to correct the data. The profile does not dictate the mechanism employed or content required to achieve such corrections.

4160 In this profile, the external agency provides data capture forms in a schema appropriate to its domain. The profile intends to minimize the work that the displaying application should do, and to bring over fully functional forms that carry with them the instruction necessary to complete the form. RFD also supports archiving a copy of the completed form.

4165 RFD offers the capability to leverage industry standards that address both the structure and content of forms used for data capture. HL7's Individual Case Safety Record (ICSR) and CDISC's Operational Data Model (ODM) provide examples.

The infrastructure provided by the RFD Profile can be utilized by many domain groups and the following domain-specific use cases illustrate the wide variety of uses to which RFD can be made.

4170 **17.1 Use Cases**

The following use cases indicate how this profile might be used by various disciplines. The RFD Profile enables all of these use cases. It does not implement any of them. Actual discipline specific profiles that specify both the use of RFD and the rules for data objects are expected in future domain-specific IHE profiles.

4175 **17.1.1 Investigational New Drug Clinical Trial Use Case**

4180 The setting for the clinical trial use case is a physicians' practice where patient care is delivered side-by-side with clinical research. The site, Holbin Medical Group, is a multi-site physician practice, employing over 100 physicians in a variety of specialties. Holbin's CEO encourages the physicians to participate as site investigators for pharmaceutical-sponsored clinical trials; Holbin provides support for clinical research activities in the form of a Research Department of twelve dedicated study coordinators, mostly RNs, along with clerical and data-entry support personnel. Holbin Medical Group uses an Electronic Health Record (EHR) and a number of sponsor-provided Electronic Data Capture (EDC) systems for documenting clinical trial activities. (For our purposes, an EHR is any application which is the primary site for documenting patient care, and retrieving patient care information. Thus we include in our span of interest many systems installed today that are not quite EHRs in the strictest sense, but which would still benefit from this approach.)

4185 Holbin's involvement in a clinical study begins when the Research Department receives a request for proposal from a study sponsor. A Study Coordinator, Patricia Zone, RN, evaluates the RFP for business viability and clinical appropriateness, and provides the requested documentation back to the sponsor. After being selected as a site for the trial, identified as #1234, and providing the required regulatory documentation to the sponsor, the physician identified as the Principal Investigator and other study personnel receive protocol-specific training from the sponsor. During the trial set-up period, Patricia ensures that the appropriate system security is in place for this protocol, recruits patients to participate as subjects according to inclusion and exclusion criteria described in the study protocol, schedules patient visits, manages data capture and data entry, and performs all the attendant financial tasks.

4195 Patricia contacts Corey Jones, a patient at Holbin, about participating in the trial, and Corey agrees to participate as a subject. Patricia registers Corey in the EHR as a subject in trial #1234, using the EHR's patient index. She schedules Corey's study visits using the EHR scheduling module, and flags the visits as pertaining to the trial #1234. After the set-up stage, the site initiates clinical trial care and trial-specific documentation.

4200 The use case continues with current state and desired state scenarios, which describe data capture utilizing EDC technology during a patient clinical trial visit before and after the RFD implementation.

4205 **17.1.1.1 Current State**

Corey Jones arrives at the clinic for a scheduled trial visit and meets with Patricia Zone for a face-to-face interview. Patricia logs into the EHR and documents the visit with a terse entry: 'Mrs. Jones comes in for a clinical trial visit associated with study #1234.' Patricia interviews

4210 Mrs. Jones, makes some observations, and records her observation on a source paper document. She looks up recent lab results in the EHR and records them in the Case Report Form (CRF). The EHR provides only a portion of the data required to complete the form, the rest comes from the interview and observations. (Estimates on the percentage of data required for a clinical trial that would be available in an EHR vary from 5% to 40%. Even in the best case, the EHR
4215 typically captures only a subset of the data required by a study protocol.)

The completed source document is forwarded to Bob, the data entry person. Bob identifies the CRF as belonging to trial #1234, and selects the trial #1234 EDC system, which may be housed on a dedicated laptop provided by the sponsor or may be accessible via a browser session connected to the Sponsor's EDC system via the Internet. He takes a three ring binder off the
4220 shelf and refers to his 'crib sheet' to get the instructions for how to use this particular system. He logs into the EDC application, using a user name and password unique to this system, and enters the data into the correct electronic case report form (eCRF) for that trial visit. Once the source document has been processed, Bob files it in a 'banker's box'⁴ as part of the permanent source record of the trial (in order to meet the requirements of the Federal Code of Regulations 21CFR
4225 312:62).

In addition to trial #1234, Bob performs data entry on eight additional EDC systems, five on dedicated laptops and three that are web-based. The web-based EDC systems save on table space, but still require entries in the three ring binders where Bob puts his 'crib sheets'. It is a chore to make sure that data from a particular trial gets entered into the corresponding laptop with its unique login ritual and data capture form, so Bob experiences much frustration in dealing
4230 with this unwieldy set of systems. Bob is a conscientious employee, and stays current in his work. But in many other sites the data entry person holds the CRF for a period of time before entering the data, perhaps entering data twice a month, or entering the data the week before the monitor visit occurs.

4235 **17.1.1.2 Desired State**

Mrs. Jones arrives for a visit and Patricia logs into the EHR, pulls up Mrs. Jones's record, and identifies the scheduled clinical trial visit. Because of the patient identification and scheduling steps that took place in the set-up stage, the EHR recognizes Mrs. Jones as a subject in Trial
4240 1234, and requests an electronic case report form from trial #1234's EDC system, using RFD. If the trial is sufficiently complex, the retrieved form may contain a list of relevant forms from the EDC system for Patricia to choose from. When the correct context is established between the EHR and the EDC, Patricia selects the clinical research tab within the EHR application to reveal the appropriate form. The EHR checks Patricia's credentials, confirms that she is empowered to view the form, and displays the form. The data capture form is essentially the same form that the
4245 EDC system would offer for this visit, and its presentation may take on some of the look and feel of the EHR's user interface. The use of a crib sheet may still be necessary, although sophisticated forms should carry with them information on how to fill out the form.

⁴ See the definition: http://www.archivists.org/glossary/term_details.asp?DefinitionKey=1193

Patricia interviews Mrs. Jones and enters data into the clinical trial form. Data from the EHR database may be pre-populated into the proper data fields (which have built-in edit checks).
4250 Upon completing the form, Patricia hits the submit button, and the EHR returns the complete form to the EDC system, using RFD. A copy of the document is archived in the site clinical trial document vault as part of the permanent source record of the trial.

17.1.2 Public Health Reporting Use Cases

17.1.2.1 Public Health Scenario 1

4255 17.1.2.1.1 Current State

Mrs. Smith presents to the Emergency Department of the Community Hospital with digestive complaints. The health care provider sends samples to the lab. The laboratory identifies cryptosporidium. The laboratory personnel query the laboratory database for weekly required public health reporting. Cases are identified, and information from the laboratory information
4260 system is copied to the public health form, printed, and sent to the public health authority. The public health officials review the reports submitted from the health care providers in the jurisdiction and identify that multiple cases of cryptosporidium have been presenting to area hospitals. Notification of the event is communicated to health care providers in the area to notify them to watch for additional cases. Water supplies servicing the affected areas are tested and
4265 treated accordingly. However, with the delay in the detection process caused by the paper-based process, numerous additional cases of cryptosporidium infection present for care.

17.1.2.1.2 Desired State

Mrs. Smith presents to the Emergency Department of the Community Hospital with digestive complaints. The health care provider sends samples to the lab. The laboratory identifies
4270 cryptosporidium. The laboratory system identifies this test result as a required public health report and sends it to the state DOH using PHIN standards as soon as the result is verified in the laboratory system. In addition or alternatively, a form is retrieved using the RFD Profile from the Biowatch public health system. The case reporting form is presented to the provider, pre-populated with EHR mapped data. The healthcare provider fills out the remaining supplemental
4275 information and submits this data electronically to the public health authority. The public health authority receives numerous electronic reports from laboratories and health care providers in the jurisdiction. Notification is sent to area health care providers and laboratories in the area to notify them to watch for additional cases. Water supplies servicing the area are tested and treated accordingly. With the early detection through process automation, further illness in the
4280 community is minimized.

17.1.2.1.3 Anthrax and Avian Influenza Scenarios: Disease Monitoring Based on Presumptive Diagnoses and/or Patient ‘Problems’

4285 Anthrax: Patient presents at ED with rapidly progressive respiratory symptoms. Gram stain of sputum reveals gram positive rods, chest X-ray reveals a widened mediastinum, and patient's condition rapidly deteriorates. Culture of sputum in laboratory is suspicious for *Bacillus anthracis*. State DOH contacted and specimens sent for confirmation. Once confirmed, the state

4290 DOH notifies appropriate local, regional, state, and federal officials (e.g., CDC, FBI, USAMRID), and notifies local hospitals, providers, and media. (This involves a bioterrorist scenario on the back end after ID confirmation – the influenza scenario below does not, but probably invokes the same pathways.)

4295 Once notified of the potential for additional cases, the ED performs STAT Gram stains on sputa and PA/Lateral Chest Xrays for all patients presenting with rapidly progressive respiratory symptoms. Presence of Gram positive rods in sputum is entered directly into the lab system OR by designated ER staff into a specific ADT field on the patient ADT screen in the CIS for internal / external surveillance reporting. Rapid reading of Chest X-ray with mediastinal widening is entered in a specific ADT field by designated staff (e.g., Radiology technician) on behalf of physician. Entry of information in these fields creates a transaction of the information to the local public health department biosurveillance system (BIS) as presumptively diagnosed inhalational anthrax. The BIS aggregates information received from multiple sites to present the location, origin and extent of presumptive and defined case presentation.

4300 Influenza: Physicians around hospital and hospital ED get rapidly increasing number of patients with respiratory symptoms suggestive of a viral infection, but no increased prevalence of similar symptoms in surrounding hospitals. Rapid test for influenza A/B is positive in many of the patients and epidemic influenza is circulating in the community. Respiratory culture is negative for bacterial pathogen at 24 hr., but viral culture is positive for influenza A. AH5N1 is suspected due to association of patients with each other and “dead chickens”. All specimens are sent to state DOH ASAP for ID. State lab identifies AH5N1. Follow-up similar to #1 above. The follow-up once notification is disseminated from health department(s) to local providers, is similar to the presumptive diagnosis information transmission to public health BIS. A more robust method for collection of presumptive diagnoses in either scenario (but not near-term) is to use standardized “problem” terms (using SNOMED) for selection of presumptive problems as part of routine operations of a CIS for physician order entry and for physician and nursing documentation.

4310 The difference in these two scenarios is that the Anthrax case involves syndromic surveillance (severe respiratory symptoms and a widened mediastinum on X-ray: need radiology surveillance and cross-correlation to ED and Lab – much more complex.)

17.1.3 Pharmaco-vigilance Scenario

4320 A community-based physician, Dr. Cramp, sees a patient in an outpatient clinic and accesses the patient’s electronic health record which reveals that the patient is on one of the new statin drugs. The physical examination turns up muscle weakness in the patient’s calves, which the physician recognizes as a possible adverse reaction to the statin. He orders a total creatinine kinase lab test to help in diagnosing the problem.

17.1.3.1 Current State

4325 Dr. Cramp exits the EHR and, using a web browser, goes to <http://www.fda.gov/medwatch/>. He brings up form FDA 3500, for ‘voluntary reporting of adverse events noted spontaneously in the course of clinical care’. He navigates through several screens of routing and instructions to arrive at the first screen of the actual form, which requests patient identifier, age at time of event or date

4330 of birth, sex, and weight; the second screen requests seven entries: a classification of the event, classification of outcome, event date, report date, description, relevant tests (he notes that a test has been ordered), and other relevant history (the last three fields are text entry); the third and fourth screens ask for details about the product ; and so forth. In actuality, the current state is that this form is seldom completed.

17.1.3.2 Desired State

4335 Dr. Cramp sees the patient and accesses the EHR as above. Upon finding the potential problem, he clicks on an ‘Adverse Event Reporting’ button which brings up FDA form 3500, using the EHR user interface. The form is presented with the demographics already completed. The product name is part of the working context of the EHR session, and is automatically loaded into the appropriate field. Dr. Cramp completes the empty fields of the form and submits directly to the FDA Medwatch site.

4340 RFD takes care of retrieving the form from MedWatch, displaying it, and returning the form to FDA. Note that the profile does not address whether or not the EHR stores a copy of the form or preloads it with EHR data. Simply using the EHR to display, complete, and submit the form is sufficient. The EHR and the site might decide to capture and store the form in the EHR database, which would be a permitted extension of the profile, but not necessary.

4345 17.1.4 Cardiology Research Use Cases

17.1.4.1 Cardiology Use Case 1 - Submission to National, State and Regional Data Registries

4350 Several jurisdictions have mandatory requirements for submission of data for particular cardiac procedures, (e.g., New York State for angioplasty and cardiac surgery, or the US for implantation of cardioverter defibrillators in Medicare patients). Additionally, many institutions participate in voluntary regional or national data registries, notably the NCDR™ National Cardiovascular Data Registry.

4355 A single cardiac patient’s data may be submitted to multiple registries. It is therefore useful for data collections for multiple submissions to be done simultaneously, so that the nurse preparing the data can review the patient medical record once and extract relevant data to each of the submission forms. Additionally, the patient’s “medical record” is in fact spread across several electronic and paper-based systems, so that repeated access in the preparation of multiple submissions must be minimized.

4360 Most of the cardiac registry submissions require data from several encounters. E.g., the NCDR gathers data on patients who undergo diagnostic cardiac catheterization followed by a percutaneous coronary intervention (PCI). If the patient had presented to the Emergency Department with an ST-elevation infarction, only a small portion of the NCDR-required data is gathered in association with the catheterization procedure. The following information is needed to complete the NCDR data set: Date of previous CABG, date of previous PCI, time of arrival in the ER, baseline laboratory data (BUN, creatinine), information from the patient’s history (family history of CAD, history of stroke, pulmonary and renal disease, etc.), measured cardiac

ejection fraction prior to PCI, QCA findings, inventory of the devices used (including bar codes), and medications administered.

4370 Thus, the preparation of the submission must be done incrementally at each encounter, and/or retrospectively at a time that all the information can be determined. Incremental preparation is problematic, since at the initial encounters it is not known what procedures the patient will undergo, and hence what registries' data forms need to be filled in. Purely retrospective data collection is similarly problematic, as it is better to obtain the data when it is produced, rather than needing to search through the record for it.

4375 Carl Cardiac, a patient, presents at the ED with chest pain, and based on ECG and history is whisked to the cath lab for a diagnostic and interventional procedure. During the PCI, while things are slow during the angioplasty balloon inflation, Ted Tech, the cath lab technologist, calls up the (empty) state and national angioplasty registry forms from the forms repository onto the cath lab logging system, and begins filling in relevant information from the case. During
4380 post-procedure clean-up, he completes as much information as he knows, and stores the partially filled-in forms back to the forms repository.

At the end of the month, Nancy Nurse is assigned the task of completing the registry data collection for that month's cath patients. She retrieves a list of cath patients, and for each one pulls up partially completed forms. When she gets to Carl's name, she pulls up the forms as
4385 partially completed by Ted, and accesses Carl's lab results, cath procedure report, nursing notes from the CCU, and discharge summary report. She fills in the remainder of the registry forms, and stores the completed forms back to the repository.

At the end of the quarter, Adele Admin uses a specialized application to retrieve all the completed forms for the national registry for the quarter from the repository, and to prepare the
4390 submission. She does a similar task with an application that processes the state registry forms.

17.1.4.2 Cardiology Use Case 2 – Performance Measures

A major issue in cardiology is improving the quality of care by monitoring select performance measures. There is a strong collaborative arrangement between the ACC, AHA, CMS, JCAHO, and AHRQ on the development and use of performance measures, such as the new ACC/AHA
4395 Clinical Performance Measures for Adults With ST-Elevation and Non-ST-Elevation Myocardial Infarction.

These performance measures require data collection, similar to the collection of data for submission to registries. However, after collection of data for a particular time period, further analysis on the total patient population must be applied to obtain an appropriate denominator for
4400 the reported measures (i.e., certain patients must be retrospectively excluded from the population data set).

17.1.5 Radiology Use Case – Clinical Impact Registry

As part of the effort to assess the impact of PET imaging on cancer patient management, the Centers for Medicare and Medicaid Services have predicated reimbursement, for a number of
4405 otherwise non-reimbursed procedures, on the submission of study data to a National Oncologic

PET Registry (NOPR) operated by the American College of Radiology at www.cancerpetregistry.org.

This use case involves a sequence of forms which must be submitted for a given patient study and includes overlaps with the billing process.

4410 PET Facilities are required to register their site with NOPR. Because access to NOPR is limited to registered facilities and because the facility depends on complete submission to get the reimbursement, the PET Facility has the primary responsibility and direct access for submitting all data. The referring physician does not have access to NOPR.

4415 Paul Positron, a patient, presents with indications of stomach cancer (or other indication covered only by participation in the NOPR). His physician, Dr. Jones, refers him to PET-Pros, a participating PET facility. PET-Pros obtains basic demographic information from Dr. Jones and submits this information to NOPR via a Web form, at which time a Registry case number is assigned by NOPR.

4420 Once a Registry case number is created, NOPR emails Dr. Jones the Pre-PET Form that must be completed with case specific clinical details and forwarded to PET-Pros for entry into the NOPR database by midnight of the day of the PET scan.

4425 At some time before the PET study, or when Paul arrives for the PET scan, PET-Pros provides Paul with the ACR IRB-approved standard NOPR Patient Information Sheet. Paul can contact the NOPR directly for more information, if necessary. Paul indicates his NOPR consent verbally to staff at the PET facility, either on the day of the PET study or within two working days after the PET study is completed. Written consent is not required. PET-Pros notes in the PET Report Form, if the patient gave or withheld consent for use of his data in future NOPR research.

Once the PET scan has been performed and reported, PET-Pros submits a study completion form and a report form (including the report provided to Dr. Jones) to NOPR.

4430 NOPR emails Dr. Jones the Post-PET Form for completion. This form collects information relating to the impact of the scan. It also includes an ACR IRB-approved Referring Physician Information Sheet and indication whether physician consent for use of the response data in future NOPR research has been given or withheld. The Post-PET form must be completed, forwarded to PET-Pros and entered into the NOPR database within 30 days of the PET scan.

4435 The NOPR database notifies PET-Pros when all case data have been entered so that the facility can bill CMS for the study. PET-Pros can check on the case status of their patients at any time using the PET Facility Reporting Tools available on the NOPR Web site.

17.1.6 Data Clarification

4440 There is a need for a clarification process that enables a sponsor organization to highlight data that needs to be examined and potentially corrected. These are detected by sponsor-initiated checks (edit checks) that result in sponsor data queries for clarification, correction, or verification relating to previously submitted data. These queries about previously submitted data are provided to the EHR system upon request. Note that there is no automated notification to the EHR that these queries for clarification / correction / verification exist. It is up to the EHR to

4445 periodically make requests when working with a sponsor that performs these edit checks.
Performing these longitudinal edit checks on submitted data does not apply to all use cases.

17.1.6.1 Current State - query process

4450 Edit checks built in to eCRFs can facilitate accurate and complete data capture; however, it is
probable that during the course of a trial, some data elements will need to be reviewed by the site
for clarification, correction, or verification. As data managers review the data (through manual
and/or system-supported validation processes), they identify missing, incomplete, or potentially
discrepant data (e.g., a site reports a patient was prescribed penicillin for a headache). Data
queries are generated through an EDC system and sent back to the site for clarification/
correction/ verification by the research coordinator. For each data query, the coordinator must
4455 reference the source record where the data element was originally documented and compare the
queried data element to the source. On occasions, the site may need to contact the patient if the
source is incomplete (e.g., a stop date on a medication). Clarifications to the data are documented
by the coordinator in the source and if it is determined that the source record is in error,
corrections are clearly documented in the source per GCP guidelines. The coordinator then
4460 responds to the query in the EDC system providing a reason for any updates to the original
record which the system captures in the audit trail. The data manager can then review the updates
and the response and close the query if no further information is required.

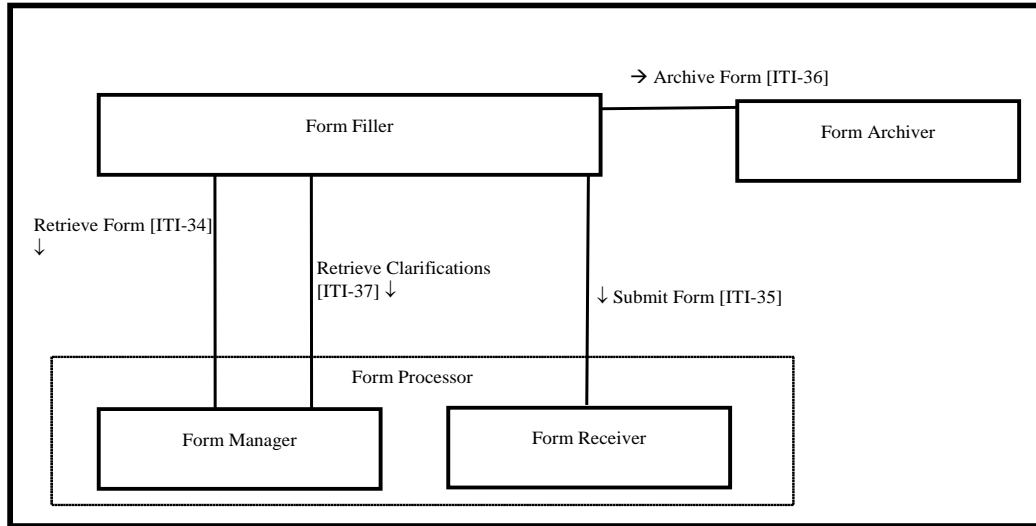
17.1.6.2 Future State - query process

4465 Edit checks built into trial-specific forms and eCRFs in the EHR system can facilitate accurate
and complete data capture; however, it is probable that during the course of a trial, some data
elements will need to be reviewed by the site for clarification, correction, or verification.

4470 As data managers review the data (through manual and/or system-supported validation
processes), they identify missing, incomplete, or potentially discrepant data (e.g., a site reports a
patient was prescribed penicillin for a headache). Data queries are generated through the sponsor
system and prepared to the site for clarification/ correction/ verification by the research
coordinator. The EHR study coordinator accesses and reviews each data query through the EHR
system referencing the EHR data in order to respond to the query. On occasions, the site may
need to contact the patient if the EHR data is incomplete (e.g., a stop date on a medication). The
coordinator documents clarifications to the data in the EHR system if needed and submits a
4475 query response as well as any data updates to the sponsor system and to the investigator site
archive. The query response includes a reason for any changes made which is included as part of
the audit trail in the EHR system, sponsor system, and the investigator's site archive. The data
manager of the sponsor can then review the response and the updates in the sponsor system and
close the query if no further information is required.

4480 17.2 Actors/ Transactions

Figure 17.2-1 shows the actors directly involved in the Retrieve Form for Data Capture
Integration Profile and the relevant transactions between them. Actors that may be indirectly
involved due to their participation in other profiles are not shown.



4485

Figure 17.2-1: Retrieve Form for Data Capture Actor Diagram

Table 17.2-1 lists the transactions for each actor directly involved in the Retrieve Form for Data Capture Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile that the implementations may choose to support is listed in Section 17.3.

4490

Table 17.2-1: Retrieve Form for Data Capture Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section in Vol. 2
Form Filler	Retrieve Form [ITI-34]	R	ITI TF-2b: 3.34
	Submit Form [ITI-35]	R	ITI TF-2b: 3.35
	Archive Form [ITI-36]	O	ITI TF-2b: 3.36
	Retrieve Clarifications [ITI-37]	O	ITI TF-2b: 3.37
Form Manager	Retrieve Form [ITI-34]	R	ITI TF-2b: 3.34
	Retrieve Clarifications [ITI-37]	R	ITI TF-2b: 3.37
Form Receiver	Submit Form [ITI-35]	R	ITI TF-2b: 3.35
Form Archiver	Archive Form [ITI-36]	R	ITI TF-2b: 3.36
Form Processor	Retrieve Form [ITI-34]	R	ITI TF-2b: 3.34
	Submit Form [ITI-35]	R	ITI TF-2b: 3.35
	Retrieve Clarifications [ITI-37]	R	ITI TF-2b: 3.37

4495

17.2.1 Actors

17.2.1.1 Form Manager

4500 The Form Manager supplies forms to Form Fillers based upon form retrieval requests. In some cases, the Form Manager may simply return a form from a store of forms, whereas in other cases the returned form may be selected or even constructed based upon context information supplied in the form retrieval request. Additionally, forms from a store may be modified based upon whether or not the Form Filler supplies additional information about a Form Archiver. A Form Manager may return a form instance id along with a form in response to a request to retrieve a form. The Form Manager constructs forms such that form data is submitted to a Form Receiver.

17.2.1.2 Form Filler

4505 The Form Filler Actor retrieves forms from a Form Manager as and when required. When requesting a form, the Form Filler Actor can optionally provide EHR context information by providing pre-population xml data in the request for use by the Form Manager, as well as workflow data that may be used to facilitate form selection. A form instance id may be provided to identify use of previously submitted data.

4510 The Form Filler may also specify a Form Archiver Actor. The Form Archiver Actor specified by the Form Filler is in addition to any Form Archiver actors specified by the Form Manager.

17.2.1.3 Form Receiver

4515 The Form Receiver Actor receives and processes completed or partially completed forms instance data from a Form Filler. Form Receiver processing is out of the scope of the profile.

17.2.1.4 Form Archiver

The Form Archiver Actor receives completed or partially completed forms instance data and stores these for archival purposes.

17.2.1.5 Form Processor

4520 The Form Processor is an integrated Form Manager and Form Receiver, supporting all of the transactions and options of those actors.

The Form Processor constructs forms such that form data is submitted back to itself.

17.2.2 Transactions

17.2.2.1 Retrieve Form

4525 The Retrieve Form transaction carries the form identifier from a Form Filler to a Form Manager, or Form Processor. The transaction allows a Form Filler to optionally specify a Form Archiver Actor. Additional data containing context information as well as workflow information may be supplied with the request to facilitate the selection and pre-population of the requested form. The value of the assigned form identifier determines the format of the form. Assignment of form

4530 identifiers is not profiled and is assumed to take place as a part of the setup configuration process necessary between Form Fillers and Form Managers, or Form Processors.

17.2.2.2 Submit Form

4535 The Submit Form transaction allows a Form Filler to submit form instance data to a Form Receiver Actor, or Form Processor Actor. For instance data submits to a Form Receiver when the form was retrieved from a Form Manager, or back to the Form Processor that created the form.

17.2.2.3 Archive Form

The Archive Form transaction allows a Form Filler to submit form instance data to a Form Archiver Actor.

4540 **17.2.2.4 Retrieve Clarifications**

The Retrieve Clarifications transaction allows a Form Filler to request the set of clarifications for a given organization from a Form Manager, or Form Processor. The value of the assigned organization identifier determines the named option format of the clarifications form. Assignment of organization identifiers is not profiled and is assumed to take place as a part of the setup configuration process between Form Fillers and Form Managers.

17.3 Retrieve Form for Data Capture Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 17.3-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

4550

Table 17.3-1: Actors and Options

Actor	Options	Vol. & Section
Form Filler	Archive Form	ITI TF-2b: 3.36
	Data Clarifications	ITI TF-2b: 3.37
	XForms	ITI TF-1: 17.3.2
Form Manager	XForms	ITI TF-1: 17.3.2
Form Processor	XForms	ITI TF-1: 17.3.2

17.3.1 Archive Form Option

The Archive Form Option allows a Form Filler to submit, for archival purposes, the form instance data to a Form Archiver.

4555 **17.3.2 Data Clarifications Option**

The Data Clarifications Option allows a Form Filler to retrieve clarifications from a Form Manager, or Form Processor, and submit updates to a Form Receiver, or Form Processor, for data that have been previously submitted.

17.3.3 XForms Option

4560 The XForms Option allows Form Fillers, Form Managers, and Form Processors to exchange forms in XForms format. See ITI TF-2b: 3.34.4.1 for constraints that apply to this option.

17.4 Retrieve Forms for Data Capture Process Flow

4565 This section describes the process and information flow when a form is retrieved for data capture and subsequently submitted upon partial or full completion. The criteria for determining whether or not the form is “complete” is outside the scope of this profile.

Five cases are distinguished.

Case 1: This case illustrates a simple, Retrieve Form using a known formID.

4570 The identifier of a form, the formID, is known to the Form Filler, such as may happen during the registration process for participation in a Clinical Trial. formID values could also be communicated by publication of form directories or by personal communications. The method of acquisition of the formID is outside the scope of this profile and is a precondition for the Retrieve Form request.

Two actor configurations are possible:

4575 A Form Manager and Form Receiver are grouped and functioning as the form source.
Figure 17.4-1

A Form Processor exists. Figure 17.4-1b

4580 The Form Filler makes a Retrieve Form request to a Form Manager or a Form Processor. The Form Manager or Form Processor either returns the requested form, or an error indicating no form is available. When a form is returned, the Form Filler will subsequently submit the form instance data to a Form Receiver or back to the Form Processor using the Submit Form transaction. When a Form Manager and Form Receiver are grouped, there may be communications between the Form Receiver and the Form Manager, as would be necessary to support partially completed forms, but these communications are internal and are not IHE transactions.

4585




Figure 17.4-1: Case 1: Retrieve Form and Submit Form; Form Manager grouped with Form Receiver

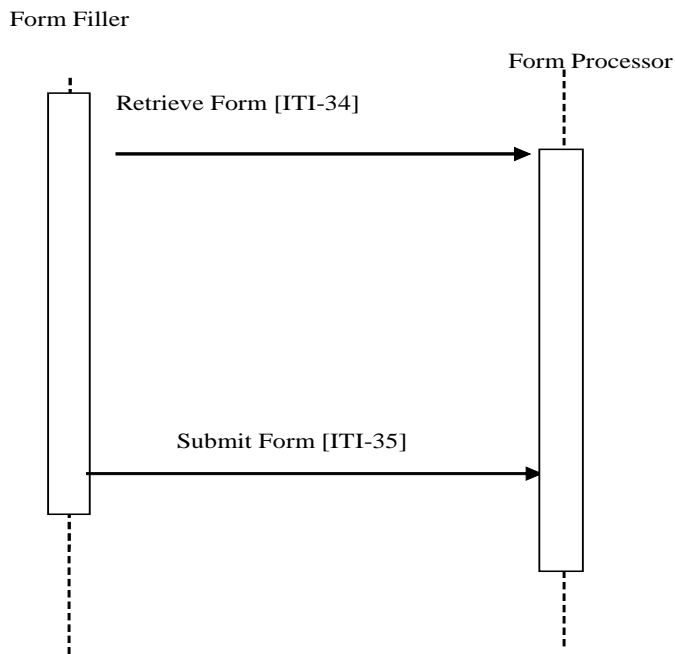


Figure 17.4-1b: Case 1: Retrieve Form and Submit Form; Form Processor

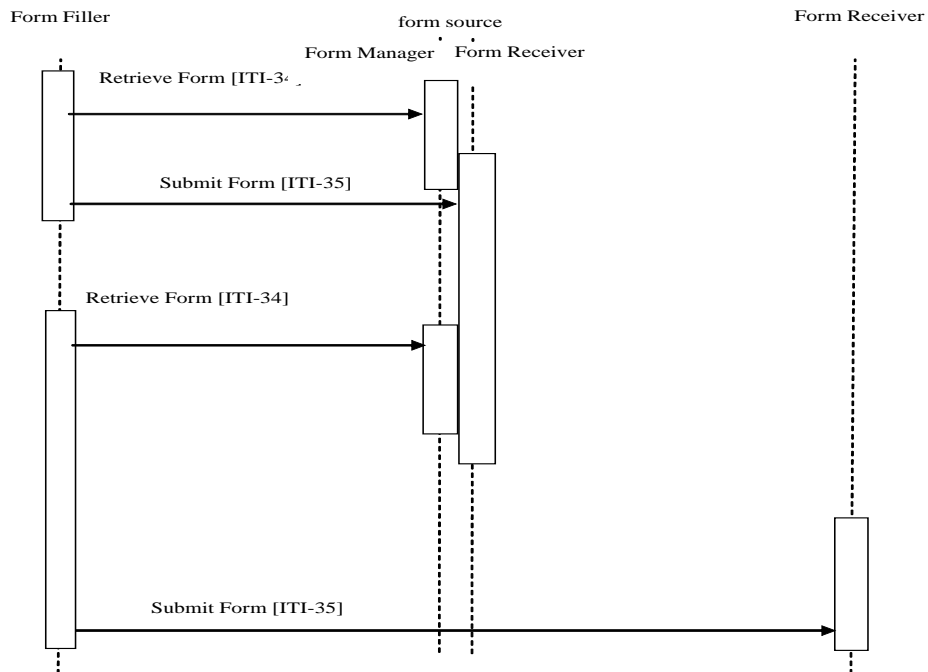
4590

4595 **Case 2:** This case illustrates that a Form Receiver may be standalone (i.e., not grouped with a Form Manager).

In this illustration there are two Form Receivers: 1) the intermediate Form Receiver, is grouped with the Form Filler; 2) the final, ungrouped Form Receiver.

4600 The identifier of a form, the formID, is known to the Form Filler; there is a grouped Form Manager and Form Receiver on one system supporting intermediate form storage, and a separate Form Receiver on a different system for final storage of form data.

4605 The Form Filler makes a Retrieve Form request to a Form Manager. The Form Manager either returns the requested form or an error indicating no form is available. When a form is returned, the Form Filler submits partially complete forms to the intermediate Form Receiver. This partially completed form can be retrieved with another Retrieve Form request to the Form Manager, and final completed form data can be submitted to the final storage, standalone, Form Receiver, such as a national data registry. The action upon submit is controlled by the form, hence the Form Manager is responsible for defining the post-submit action by selection of, or generation of, the desired action during the Retrieve Form transaction processing.



4610

Figure 17.4-2: Case 2: Retrieve Form, Submit Form; Form Manager separate from Form Receiver

Case 3: In this case the Form Filler uses the Archive Option.

4615 The Form Filler makes a Retrieve Form request to a Form Manager, specifying that archival is necessary to a specific Form Archiver. The Form Manager either returns the requested form or an error indicating no form is available. The Form Manager constructs the form to perform an

4620 archive transaction to the Form Archiver specified in the Form Filler’s Retrieve Form request. When the form is returned and subsequently submitted, form instance data is submitted to the Form Receiver and also to the Form Archiver.

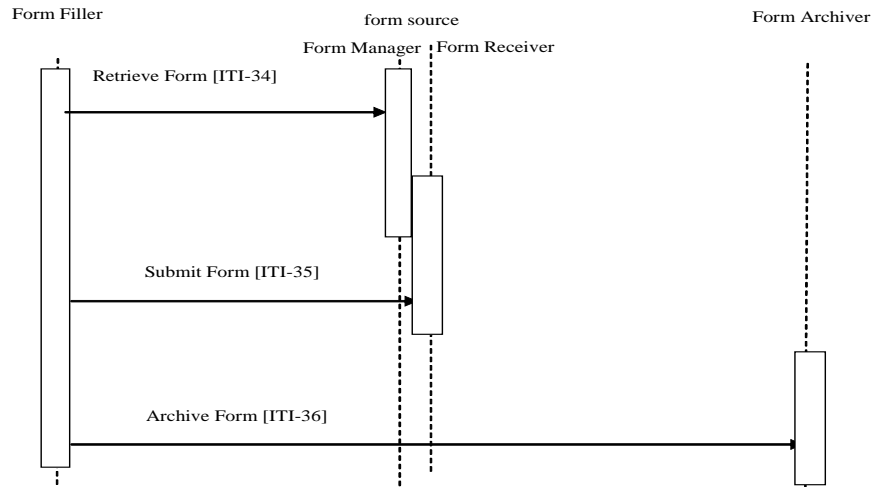
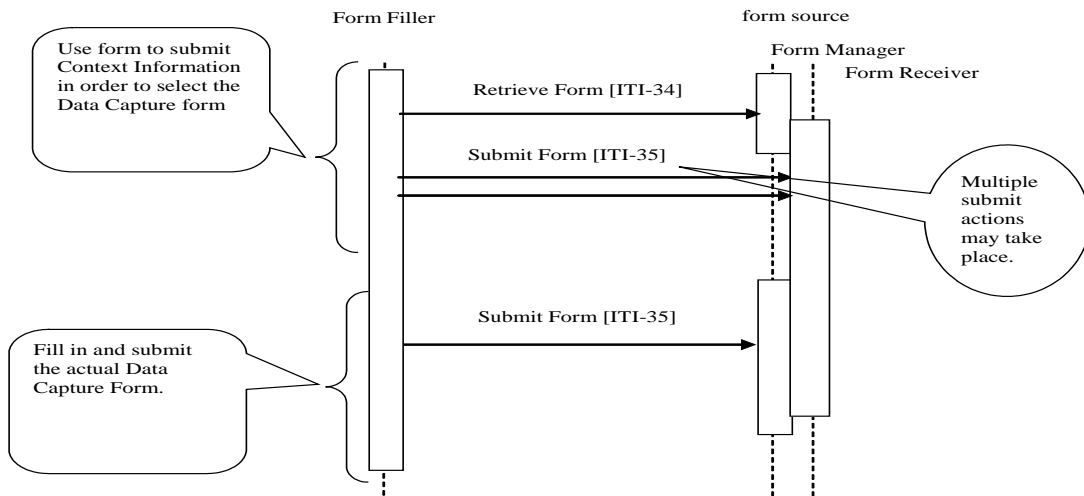


Figure 17.4-3: Case 3: Retrieve Form, Submit Form, Archive Form

4625 **Case 4:** This case illustrates one way to use Form design to solve the issue where a formID is not known in advance.

The identifier of a form, the formID, is not known to the Form Filler, but a set of context value (name, value) pairs is known. A context form where these values could be entered would have a formID. Information collected by the instance of a context form would be used by the Form Manager to determine the appropriate data capture form to return to the Form Filler.

4630 The Form Filler has enough information to request a context form that collects information that can help the Form Manager determine the actual data capture form. The Form Filler completes the context form, submits this to the Form Receiver which returns either new instance data, or a new form.



4635

Figure 17.4-4: Case 4: Retrieve Form; Submit Form

Case 5: In this case the Form Filler supports the Data Clarifications Option.

The Form Filler makes a Retrieve Clarifications request to a Form Manager. The interactions of Form Receiver and Form Manager are outside of the scope of this profile. An example of a solution for providing clarification information to a Form Manager is to group the Form Manager with the Form Receiver, as shown in figures 17.4-5 and 17.4-6. The request made by the Form Filler contains an organization identifier allowing the Form Manager to return only the set of clarifications relevant to the organization making the request. The Form Manager returns a form containing the necessary information to allow the site or organization making the request to amend the data as required. These Retrieve Clarifications requests must be periodically executed by the Form Filler. The frequency of request is likely based upon some duration as defined or agreed upon by the Form Manager / Form Receiver.

The Form Manager can return either a form containing the data to be modified or a form containing a list of references to other forms. In the second case, the references are used to obtain the individual forms using the Retrieve Form transaction. In both cases the data are then modified and submitted to the Form Receiver using the Submit Form transaction. Submitted data may then be evaluated by the data manager of the sponsor for proper handling.

The profile does not distinguish between the two responses, the content returned within the form allows the user of the Form Filler to process the form returned in the appropriate manner.

4655

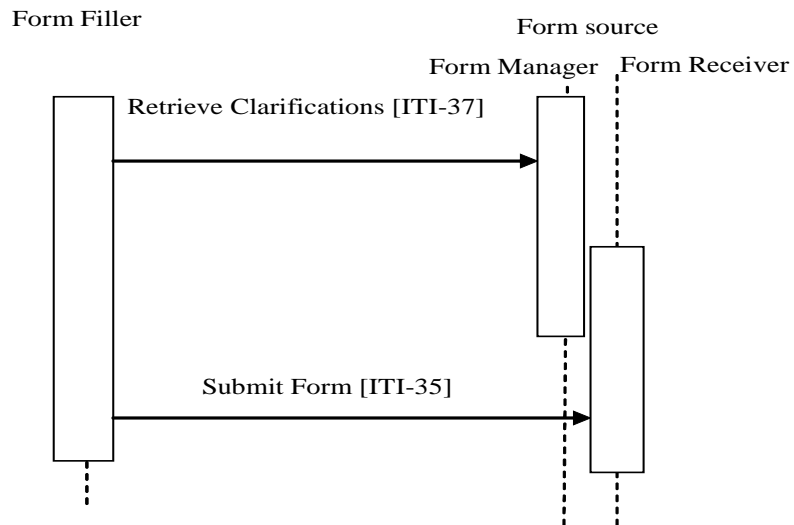


Figure 17.4-5: Case 5: Form Filler supporting Data Clarifications Option

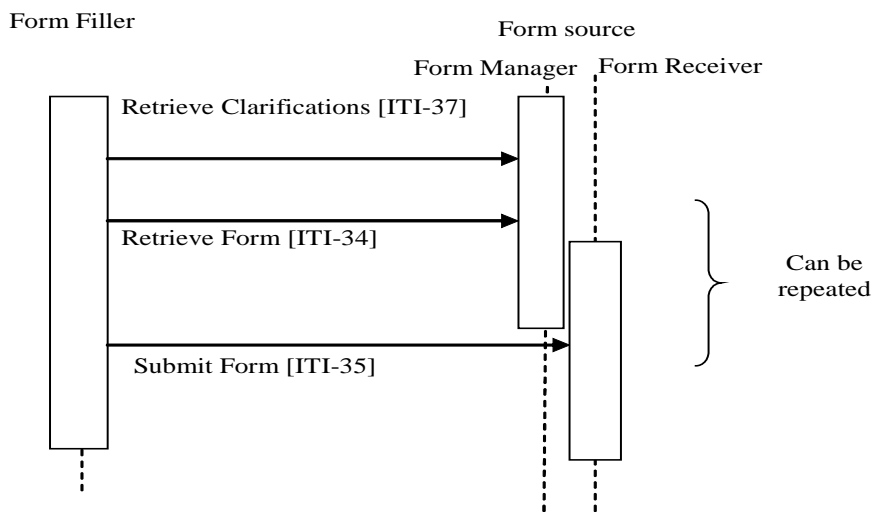


Figure 17.4-6: Case 5: Form Filler supporting Data Clarifications Option

4660 **17.5 Security Considerations**

17.5.1 RFD Risk analysis Risk Assessment

The risk analysis for RFD enumerates assets, threats, and mitigations. The complete risk data is stored and available from IHE⁵.

⁵ The risk analysis data may be found at: ftp://ftp.ihe.net/IT_Infrastructure/iheityr5-2007-2008/Technical_Cmte/Profile_Work/RFD/RFD%20Risk%20Analysis%202007-05-15.xls

4665 The purpose of this risk assessment is to notify vendors of some of the risks that they are advised to consider when implementing RFD actors. For general IHE risks and threats, please see ITI TF-1: Appendix L. The vendor is also advised that many risks cannot be mitigated by the IHE profile and instead responsibility for mitigation is transferred to the vendor, and occasionally to the affinity domains, individual enterprises and implementers. In these instances, IHE fulfills its responsibility to notify affected parties through the use of the following sections.

4670 **17.5.2 Recommendations**

The high impact risks include: accuracy errors, mismatch between data and schema, disclosure of trade secrets. This profile includes the mitigations:

M1. If the user notices that the wrong form has been retrieved they will discard the form. Since Form Retrieval is stateless, a discard of the form shall cause no problems.

4675 **M2.** When using the XForm Option, the XForms model provides for schema validation of the data model. The XForms plugins responsible for processing and displaying XForms, which are outside of this profile, are required to validate forms.

4680 **M3.** TLS may be implemented, so that those affinity domains and enterprises that need privacy protection and site authentication can use it. (Implementations must provide the TLS, but the decision to activate it is up to the affinity domain and enterprises.)

M4. Form validations will prevent submission of forms with missing data.

M5. The RFD Archive Form Transaction for saving source data to a trusted third party is an option that it is available to enterprises.

These mitigations are transferred to Vendors and Clients.

4685 **T1.** IHE recommends that providers evaluate and review forms as presented before entering data and submitting. Provider review is an essential part of the forms retrieval and submission process to ensure data is entered into the correct form and for the correct patient. Vendors are cautioned not to use RFD for unmediated treatment or diagnosis. A doctor must always intervene prior to treatment or diagnosis to ensure that errors that may occur in transit are checked by a human
4690 prior to engaging in any treatment or diagnosis of a patient.

T2. The supported format options allow for basic data validity checks within the form. It is the responsibility of the forms designers/implementers to take advantage of this to protect against entry errors, etc.

4695 **T3.** The need for partially filled forms identifies this as a workflow issue within the organization(s) supplying the data.

T4. Forms and workflow designers should break forms into sequential step forms if possible.

T5. Forms Design should facilitate evaluation of workflow and gaps.

T6. Access control and security at the client site are important mitigating factors to potential disclosures.

4700 **T7.** Policy controls are recommended to determine which systems may be used to perform the Form Filler Actor.

T8. Policy controls are recommended to determine which users may fill out forms.

T9. This profile does not require audit logging. An enterprise audit logging process is recommended to reduce errors and track malicious behavior.

4705 **T10.** An application feature to support roll back of forms data may be needed.

T11. Notification of the need to clarify data.

T13. Form Managers, Receivers, Archivers must be on well protected systems.

4710 **T14.** Network and Infrastructure and Systems robustness must be considered, especially for forms applications that are to be used during disasters, epidemics, and other situations where the local infrastructure may be significantly disrupted.

T15. Forms should be designed for high latency, low bandwidth links if they are for applications that are to be used during disasters, epidemics, and other situations where the local infrastructure may be significantly disrupted.

4715 **T16.** Form Fillers should be robust in the face of user error, network failure, and underlying hardware failures.

T17. Workflow must be addressed in the requirements gathering phase. Vendors are advised to discuss investigator workflow with clients.

T18. Vendors are advised to consider the implications of their logging and audit repository implementation.

4720

18 Cross-Community Access (XCA) Integration Profile

4725 The Cross-Community Access Profile supports the means to query and retrieve patient relevant
 4730 medical data held by other communities. A community is defined as a coupling of
 facilities/enterprises that have agreed to work together using a common set of policies for the
 purpose of sharing clinical information via an established mechanism. Facilities/enterprises may
 host any type of healthcare application such as EHR, PHR, etc. A community is identifiable by a
 globally unique id called the homeCommunityId. Membership of a facility/enterprise in one
 community does not preclude it from being a member in another community. Such communities
 may be XDS Affinity Domains which define document sharing using the XDS Profile or any
 other communities, no matter what their internal sharing structure.

18.1 Actors/ Transactions

Figure 18.1-1 shows the actors directly involved in the XCA Integration Profile and the relevant transactions between them.

4735 **Note:** The Document Consumer Actor is shown in Figure 18.1-1 to clarify the responsibility of
 the XDS Affinity Domain Option discussed in Section 18.2.

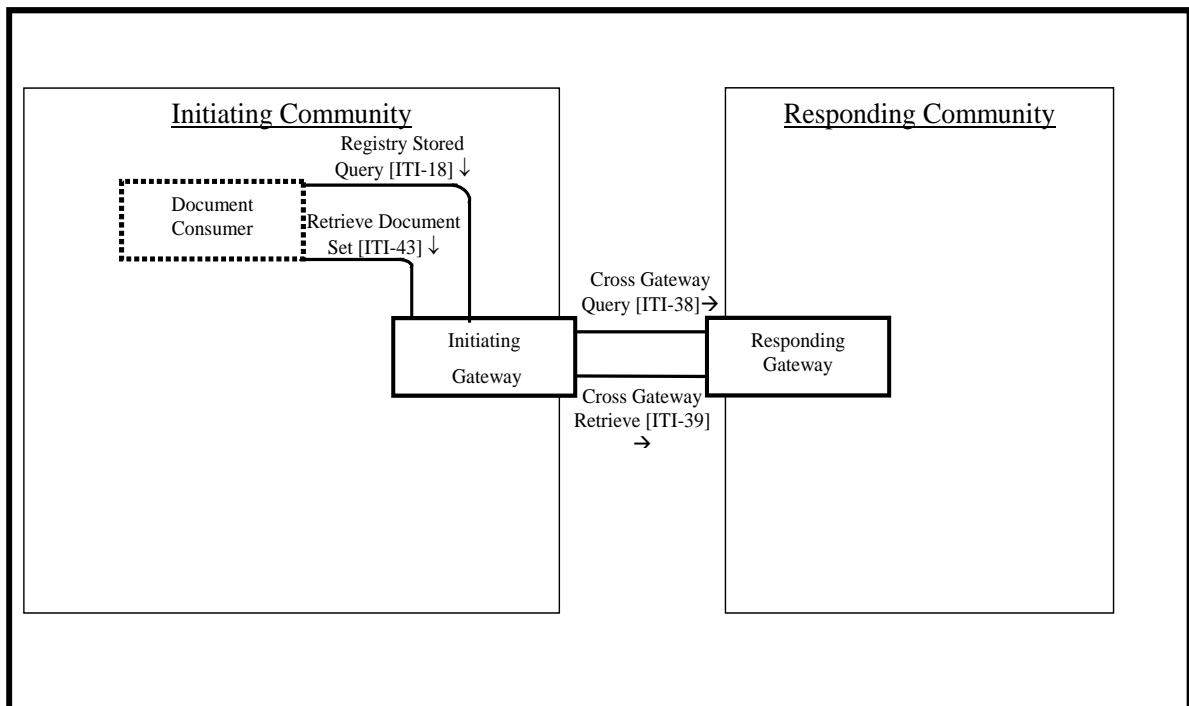


Figure 18.1-1: XCA Actor Diagram

4740 Table 18.1-1 lists the transactions for each actor directly involved in the XCA Profile. In order to
 claim support of this Integration Profile, an implementation must perform the required

transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in Section 18.2.

4745

Table 18.1-1: XCA Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Initiating Gateway	Cross Gateway Query [ITI-38]	R	ITI TF-2b: 3.38
	Cross Gateway Retrieve [ITI-39]	R	ITI TF-2b: 3.39
	Registry Stored Query [ITI-18]	O	ITI TF-2a: 3.18
	Retrieve Document Set [ITI-43]	O	ITI TF-2b: 3.43
Responding Gateway	Cross Gateway Query [ITI-38]	R	ITI TF-2b: 3.38
	Cross Gateway Retrieve [ITI-39]	R	ITI TF-2b: 3.39

Note: When an Initiating or Responding Gateway is grouped with a Document Consumer, there are additional requirements. See Section 18.2.3 for a description of grouping.

4750 **18.2 XCA Integration Profile Options**

Options that may be selected for this Integration Profile are listed in the Table 18.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 18.2-1: XCA Integration Profile - Actors and Options

Actor	Options	Vol. & Section
Initiating Gateway	<i>XDS Affinity Domain Option</i>	ITI TF-1: 18.2.1
	<i>Asynchronous Web Services Exchange</i>	ITI TF-1: 18.2.2
	<i>On-Demand Documents</i>	ITI TF-1: 18.2.4
Responding Gateway	<i>On-Demand Documents</i>	ITI TF-1: 18.2.4
	<i>Persistence of Retrieved Documents</i>	ITI TF-1: 18.2.5

4755 **18.2.1 XDS Affinity Domain Option**

Initiating Gateways which support the XDS Affinity Domain Option interact with Document Consumers within the XDS Affinity Domain served by the Initiating Gateway.

Initiating Gateway actors which support this option:

4760 shall **receive** Registry Stored Query [ITI-18] transactions from a local Document Consumer Actor and act on those requests on behalf of the Document Consumer. When receiving a Registry Stored Query from a local Document Consumer, shall require the homeCommunityId as an input parameter on relevant queries, and shall specify the homeCommunityId attribute within its Registry Stored Query responses. See Section 18.3.2 for description of homeCommunityId. Initiating

4765 Gateways which support this option shall adjust the patient identifier found in the

Registry Stored Query to an appropriate patient identifier known to the Responding Gateway receiving the Cross Gateway Query. See TF-2a:3.18.4.1.3 for details of the processing of the patient identifier.

4770 shall **receive** Retrieve Document Set [ITI-43] transactions from a local Document Consumer Actor and act on those requests on behalf of the Document Consumer. When receiving a Retrieve Document Set from a local Document Consumer, shall require the homeCommunityId as an input parameter.

4775 When an Initiating Gateway does not support the XDS Affinity Domain Option it is expected to be using non-IHE specified interactions to communicate remote community data to systems within its local community. These proprietary interactions are not further described within any IHE profile.

See the relevant transactions for further details regarding the homeCommunityId attribute.

18.2.2 Asynchronous Web Services Exchange Option

4780 Initiating Gateways which support Asynchronous Web Services Exchange shall support Asynchronous Web Services Exchange on the Cross Gateway Query [ITI-38] and Cross Gateway Retrieve [ITI-39] transactions. If the Initiating Gateway supports both the XDS Affinity Domain Option and the Asynchronous Web Services Option it shall support Asynchronous Web Services Exchange on the Registry Stored Query [ITI-18] and Retrieve Document Set [ITI-43] transactions.

4785 18.2.3 Grouping Rules

Grouping with a Document Consumer Actor is used in situations where an Initiating Gateway and/or Responding Gateway are supporting an XDS Affinity Domain

4790 When an Initiating Gateway is supporting an XDS Affinity Domain, it can choose to query and retrieve from local actors in addition to remote communities. This is accomplished by grouping the Initiating Gateway Actor with a Document Consumer Actor. This grouping allows Document Consumers such as EHR/PHR/etc. systems to query the Initiating Gateway to retrieve document information and content from both the local XDS Affinity Domain as well as remote communities. For details see Section 18.2.3.1. An Initiating Gateway Actor that is not grouped with a Document Consumer Actor is only able to return results from remote communities, so
4795 local EHR/PHR/etc. systems (Document Consumer Actors) must direct separate query and document retrieve transactions internally and externally.

4800 When a Responding Gateway is supporting an XDS Affinity Domain, it may resolve Cross Gateway Query and Cross Gateway Retrieve Transactions by grouping with a Document Consumer Actor and using the Registry Stored Query and Retrieve Document Set transactions. For details see 18.2.3.2

18.2.3.1 Initiating Gateway grouped with Document Consumer

Initiating Gateways that are grouped with a Document Consumer:

- shall support the XDS Affinity Domain Option

- 4805
- shall **initiate** Registry Stored Query [ITI-18] transactions to a local Document Registry to query local information in response to a received Registry Stored Query [ITI-18] from a local Document Consumer.
 - shall **initiate** Retrieve Document Set [ITI-43] transactions to a local Document Repository in response to a received Retrieve Document Set from a local Document Consumer which contains a homeCommunityID indicating the local community.

4810

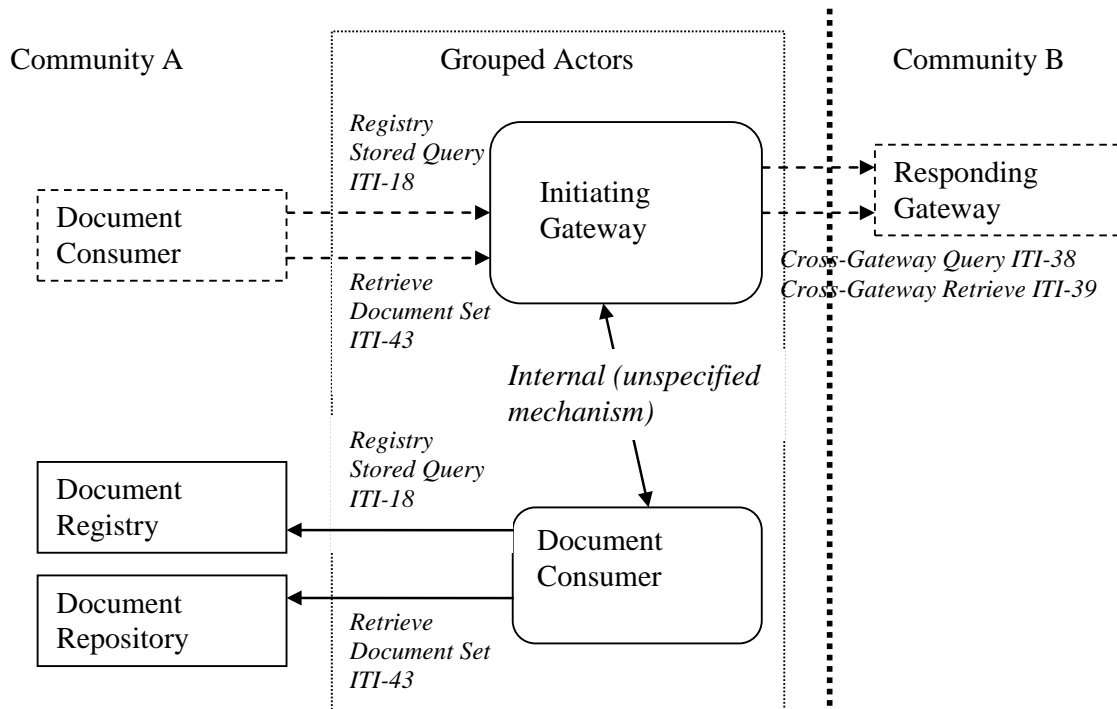


Figure 18.2.3.1-1: Initiating Gateway grouped with Document Consumer

18.2.3.2 Responding Gateway grouped with Document Consumer

Responding Gateways that are grouped with a Document Consumer:

- 4815
- shall **initiate** a Registry Stored Query [ITI-18] transaction to a local Document Registry to query local information in response to a received Cross Gateway Query [ITI-38]. The Document Registry response must be augmented with the homeCommunityId of the Responding Gateway's community prior to returning in the response to the Cross Gateway Query.
- 4820
- shall **initiate** a Retrieve Document Set [ITI-43] transaction to a local Document Repository to retrieve local information in response to a Cross Gateway Retrieve [ITI-39].

When a Responding Gateway is not grouped with a Document Consumer Actor it is expected to be using non-IHE specified interactions to collect local information in response to a Cross

4825 Gateway Query or Cross Gateway Retrieve. These proprietary interactions are not further described within any IHE profile.

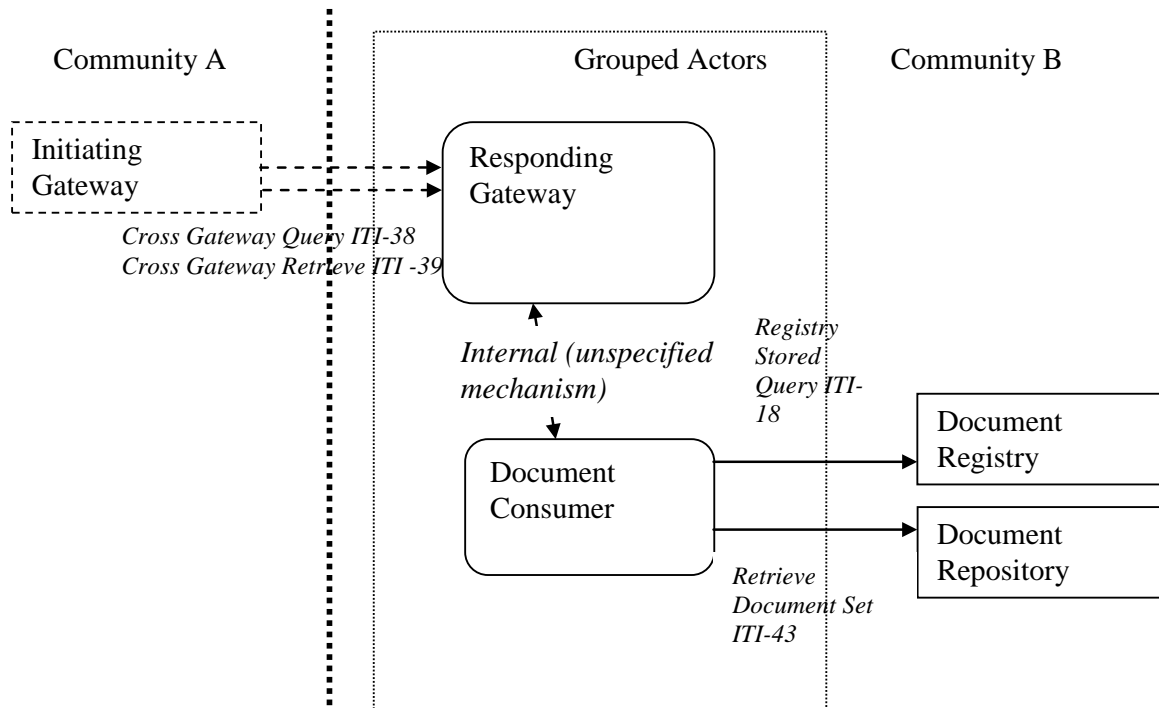


Figure 18.2.3.2-1: Responding Gateway grouped with Document Consumer

4830 18.2.4 On-Demand Documents

Initiating and Responding Gateways may declare support for On-Demand Document Entries. Refer to section ITI TF-1:10.4.13 for details about On-Demand Document Entries.

Requirements for an Initiating Gateway supporting this option differ depending on whether or not the Initiating Gateway also supports the XDS Affinity Domain Option:

- 4835
- An Initiating Gateway that supports the XDS Affinity Domain Option relies on an XDS Document Consumer Actor in its local community to generate query and retrieve requests which support On-Demand Document Entries. The Initiating Gateway does not modify content related to On-Demand Entries, but passes it in the Cross-Gateway Query and Retrieve messages to the Responding Gateway and returns full results to the Document Consumer. For this reason this option imposes no additional requirements on Initiating Gateways which support the XDS Affinity Domain Option.
- 4840
- An Initiating Gateway that supports the On-Demand Documents Option, but not the XDS Affinity Domain Option, shall:
 - be able to specify, in a Cross Gateway Query, a request for On-Demand Document Entries
- 4845

- be able to retrieve On-Demand Document Entries from one or more Responding Gateways

Requirements for a Responding Gateway supporting this option differ depending on whether or not the Responding Gateway is grouped with an XDS Document Consumer:

- 4850
- A Responding Gateway which groups with an XDS Document Consumer to interact with XDS Document Registry and Repository actors is dependent on that grouped actor to support the semantics needed for On-Demand Documents. The Responding Gateway does not modify On-Demand Document Entries obtained from the XDS Document Consumer and passes them within the Cross-Gateway Query and Retrieve response messages returned to the Initiating Gateway. For this reason this option imposes no additional requirements on Initiating Gateways which group with an XDS Document Consumer.
- 4855
- A Responding Gateways that is not grouped with an XDS Document Consumer shall:
- be able to respond with On-Demand Document Entries to a Cross Gateway Query requesting On-Demand Document Entries. Without this option, a Responding Gateway shall ignore the request for On-Demand Document Entries (i.e., return an empty list if no other type of entry matches the query parameters).
- 4860
- be able to respond to a Cross Gateway Retrieve requesting On-Demand Document Entries.

4865 **18.2.5 Persistence of Retrieved Documents**

Responding Gateways which support the Persistence of Retrieved Documents Option shall:

- also support the On-Demand Documents Option
 - make available, as a Stable Document Entry in response to a Cross Gateway Query, every document created as a result of receipt of a Cross Gateway Retrieve which specified the uniqueID of an On-Demand Document Entry
- 4870

See Section 18.3.3.2 for an overview of this.

18.3 XCA Process Flow

18.3.1 Use Cases

4875 Assume within a given domain, such as the State of California, we have several healthcare communities (or XDS Affinity Domains or RHIOs). One in Los Angeles is based on IHE-XDS. One in Sacramento is based on another form of healthcare sharing infrastructure. One in San Francisco is also based on IHE-XDS. A patient X, who travels frequently, has received healthcare in each of these communities. Patient X is admitted to a hospital in LA. The attending physician uses his hospital information system to query across multiple domains for healthcare information about this patient. Once found, references to patient data outside the local domain

4880 are cached locally for easy future reference.

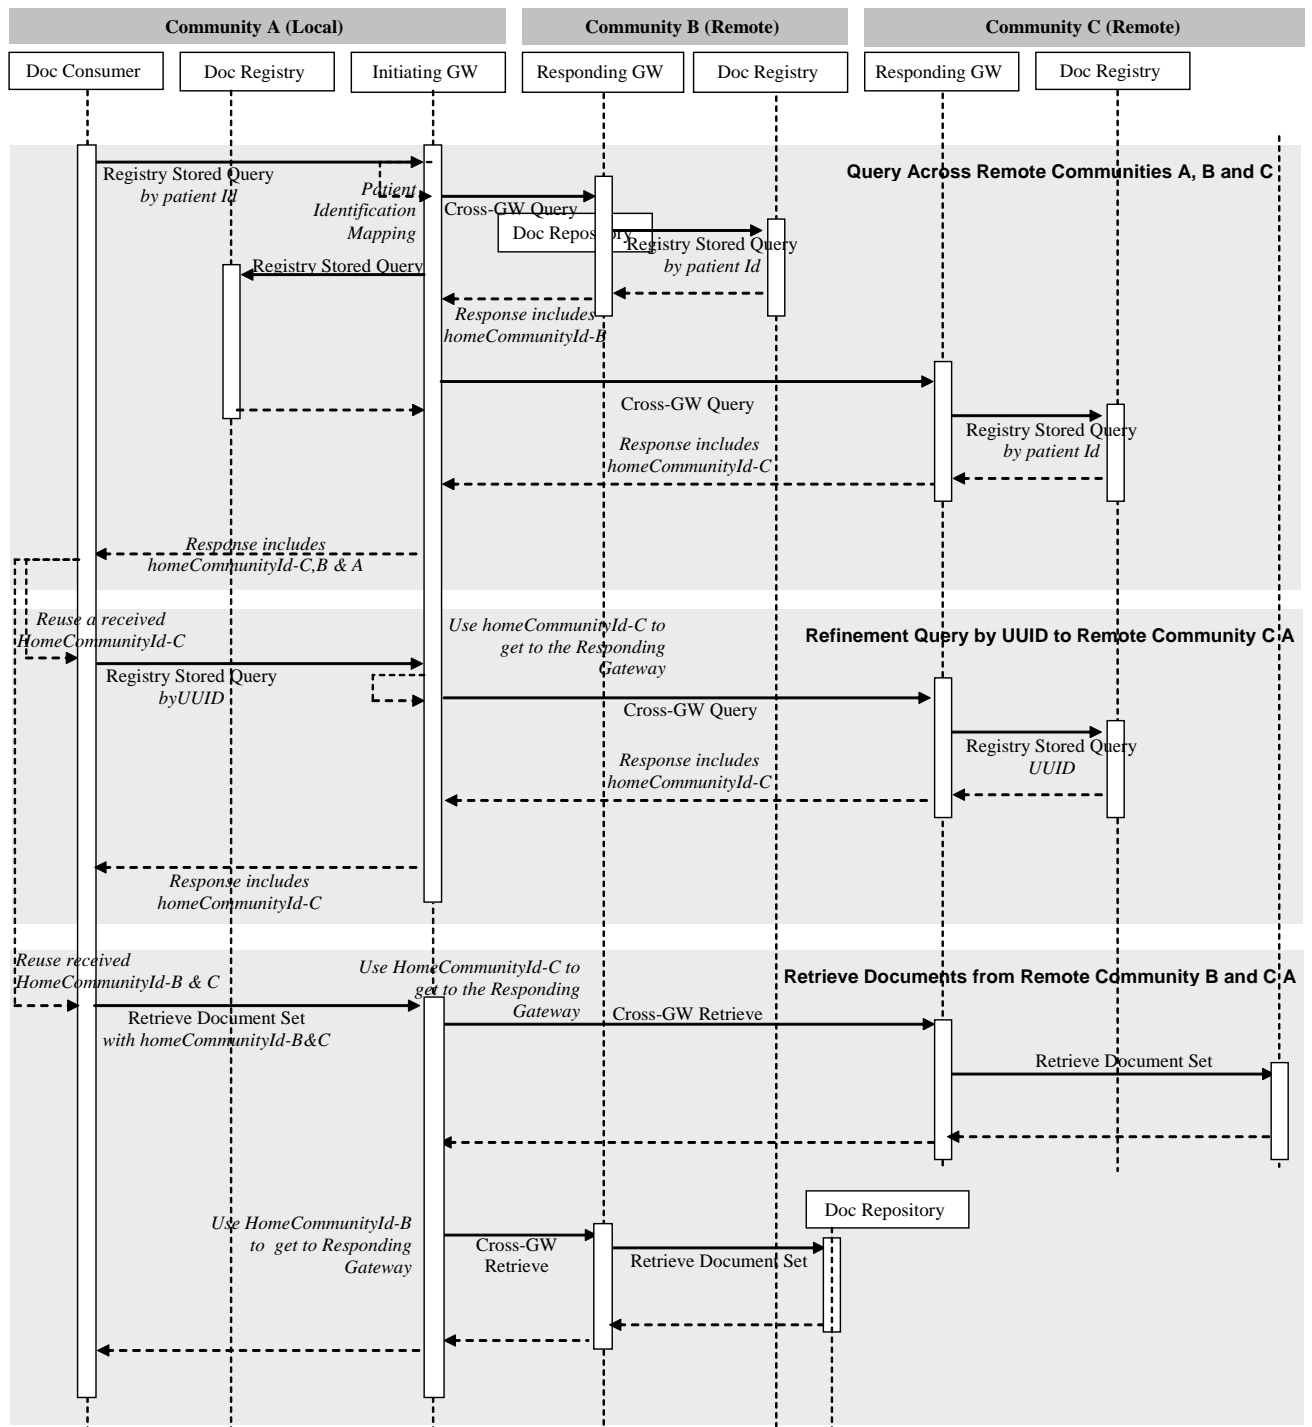
18.3.2 homeCommunityId defined

4885 This profile makes use of a homeCommunityId value which is a globally unique identifier for a community and is used to obtain the Web Services endpoint of services that provide access to data in that community. Specifically:

- It is returned within the response to Cross Gateway Query and Registry Stored Query transactions to indicate the association of a response element with a community. Document Consumers process the value in the response as an opaque unique identifier.
- 4890 • It is an optional parameter to Registry Stored Query requests, not requiring a patient id parameter, and Retrieve Document Set requests to indicate which community to direct the request.
- It is used by Initiating Gateways to direct requests to the community where the initial data originated.

18.3.3 Detailed Interactions

4895 The following diagram presents a high level view of the interactions between actors when both initiating and responding communities are XDS Affinity Domains i.e., use of the XDS Affinity Domain Option and the Initiating Gateway and Responding Gateway are each grouped with a Document Consumer. Details on each interaction follow the diagram.



4900

Figure 18.3.3-1: XCA Detailed Interactions

- **Document Consumer** initiates a *Registry Stored Query* request by patient id – the Document Consumer initiates the initial transaction by formatting a Registry Stored Query request by patient identifier. The consumer uses PDQ, PIX or some other means to

- 4905 identify the XDS Affinity Domain patient id, formats that information plus any other query parameters into a Registry Stored Query request and sends this request to an Initiating Gateway.
- 4910 • **Initiating Gateway processes Registry Stored Query by patient id request** – The Initiating Gateway receives a Registry Stored Query by patient id and must determine a) which Responding Gateways this request should be sent to b) what patient id to use in the Cross Gateway Queries. Detailed specification of these steps is not in the intended scope of this profile. Combination of this profile with other existing profiles (e.g., PIX/PDQ), future profiles or configuration mechanisms is possible. Please refer to ITI TF-1: E.10 XCA and Patient Identification Management for possible use of existing profiles PIX and PDQ. For each Responding Gateway identified, the Initiating Gateway shall update the query with the correct patient identifier corresponding to the Responding Gateway’s community and initiates a Cross Gateway Query transaction to the Responding Gateway. If the Initiating Gateway is grouped with a Document Consumer it shall also initiate a Registry Stored Query to the local Document Registry.
 - 4920 • **Responding Gateway processes Cross Gateway Query by patient id** – The Responding Gateway within an XDS Affinity Domain processes the Cross Gateway Query by using grouping as a Document Consumer and initiates a Registry Stored Query to the local Document Registry. The Responding Gateway shall update the response from the Document Registry to ensure that the homeCommunityId is specified on every applicable element. This updated response is sent as the response to the Cross Gateway Query.
 - 4925 • **Initiating Gateway processes Cross Gateway Query by patient id responses** – The Initiating Gateway collects the responses from all Responding Gateways it contacted. For each response it shall verify homeCommunityId is present in each appropriate element. If the Initiating Gateway initiated a Registry Stored Query to the local Document Registry it shall update the response to that transaction to contain the homeCommunityId value associated with the local community. Once all responses are received the Initiating Gateway consolidates all updated response data into one response to the Document Consumer. The Initiating Gateway shall return to the Document Consumer the same homeCommunityId attribute values that it received from Responding Gateways.
 - 4930 • **Initiating Gateway processes Cross Gateway Query by patient id responses** – The Initiating Gateway collects the responses from all Responding Gateways it contacted. For each response it shall verify homeCommunityId is present in each appropriate element. If the Initiating Gateway initiated a Registry Stored Query to the local Document Registry it shall update the response to that transaction to contain the homeCommunityId value associated with the local community. Once all responses are received the Initiating Gateway consolidates all updated response data into one response to the Document Consumer. The Initiating Gateway shall return to the Document Consumer the same homeCommunityId attribute values that it received from Responding Gateways.
 - 4935 • **Document Consumer receives Registry Stored Query by patient id response** – The Document Consumer receives the results of the query from the Initiating Gateway and must account for two unique aspects of the response; namely that a) the homeCommunityId attribute will be specified and b) the Document Consumer may not be able to map the repository id value directly to the Document Repository. There shall be a common coding/vocabulary scheme used across all communities. For example, all communities shall have common privacy consent vocabularies. The Document Consumer shall retain the values of the homeCommunityId attribute for future interaction with the Initiating Gateway.
 - 4940 • **Document Consumer receives Registry Stored Query by patient id response** – The Document Consumer receives the results of the query from the Initiating Gateway and must account for two unique aspects of the response; namely that a) the homeCommunityId attribute will be specified and b) the Document Consumer may not be able to map the repository id value directly to the Document Repository. There shall be a common coding/vocabulary scheme used across all communities. For example, all communities shall have common privacy consent vocabularies. The Document Consumer shall retain the values of the homeCommunityId attribute for future interaction with the Initiating Gateway.
 - 4945 • **Document Consumer initiates a Registry Stored Query by UUID** – Many Registry Stored Queries do not include patient id as a parameter, but instead require one of the entryUUID or uniqueID parameters, generically referred to as UUID. Both of these

- 4950 values are returned as part of the metadata from a query by patient id. The Document Consumer may do a patient id query to the Initiating Gateway prior to a query by UUID or shall have access to the correct homeCommunityId through some other means. In either case the consumer has the homeCommunityId attribute and shall specify it as a parameter of the query. The Document Consumer puts the homeCommunityId and UUID values plus any other query parameters into a Registry Stored Query request and sends this request to an Initiating Gateway.
- 4955 • **Initiating Gateway processes Registry Stored Query by UUID request** – The Initiating Gateway receives a Registry Stored Query by UUID and determines which Responding Gateway to contact by using the homeCommunityId to obtain the Web Services endpoint of the Responding Gateway. The process of obtaining the Web Services endpoint is not further specified in this profile. If the homeCommunityId represents the local community the Initiating Gateway will initiate a Registry Stored Query to the local Document Registry. The Initiating Gateway shall specify the homeCommunityId in the Cross Gateway Query by UUID which is associated with the Responding Gateway.
 - 4960 • **Responding Gateway processes Cross Gateway Query by UUID** – The Responding Gateway within an XDS Affinity Domain processes the Cross Gateway Query by grouping as a Document Consumer and initiating a Registry Stored Query to the local Document Registry. The response to the Cross Gateway query shall contain the homeCommunityId of the responding community. This processing is identical to processing of the Cross Gateway Query by patient id.
 - 4965 • **Initiating Gateway receives Cross Gateway Query by UUID response** – The processing of a Cross Gateway Query by UUID response is identical to the processing of a Cross Gateway Query by patient id response, except there is only one response, so consolidation of responses is not needed.
 - 4970 • **Document Consumer receives Registry Stored Query by UUID response** – The processing of a Registry Stored Query by UUID response is identical to the processing of a Registry Stored Query by patient id response.
 - 4975 • **Document Consumer initiates a Retrieve Document Set** – Prior to issuing a Retrieve Document Set the Document Consumer may issue a Registry Stored Query by patient id to the Initiating Gateway. The response to the Registry Stored Query by patient id or subsequent Registry Stored Query by UUID includes a) the document unique ID b) the repository unique ID c) the homeCommunityId attribute. If the Document Consumer did not issue a Registry Stored Query which returned this information then it shall have acquired the information through some other means. The Document Consumer shall specify these three parameters in its Retrieve Document Set transaction to the Initiating Gateway.
 - 4980 • **Initiating Gateway processes Retrieve Document Set** – The Initiating Gateway determines which Responding Gateways to contact by using the homeCommunityId to obtain the Web Services endpoint of the Responding Gateway. If the homeCommunityId represents the local community the Initiating Gateway will initiate a Retrieve Document Set to a local Document Repository. The Retrieve Document Set may contain more than
 - 4985

4990 one unique homeCommunityId so the Initiating Gateway shall be capable of initiating requests to more than one Responding Gateway and consolidating the results. The Initiating Gateway shall specify the homeCommunityId in the Cross Gateway Retrieve which identifies the community associated with the Responding Gateway.

- 4995 • **Responding Gateway processes Cross Gateway Retrieve** – The Responding Gateway within an XDS Affinity Domain processes the Cross Gateway Retrieve by grouping as a Document Consumer and initiating a Retrieve Document Set transaction to the Document Repository identified by the repository unique ID within the request. If the Cross Gateway Retrieve requests multiple documents with different repository unique IDs, the Responding Gateway shall contact multiple Document Repositories and consolidate the responses.

5000 **18.3.3.1 Sharing using On-Demand Documents**

5005 Figure 18.3.3-2 shows the use of an On-Demand Document Entry to allow access to the most recently available set of data from a responding community. This example does not show the Gateways interacting with XDS Actors and does not assume that XDS is used in any of the communities. This workflow includes persisting each document returned as a result of a retrieve and forming a stable document entry describing the new document which will be returned to future Cross Gateway Query requests. The semantics of the response to a retrieve of an On-Demand Document Entry are different than the semantics of the response to retrieval of stable documents. The response to a retrieve of an On-Demand Document Entry contains a new uniqueId, different than the uniqueId specified in the request, which is the same as the value
5010 within the ClinicalDocument id of the document returned. Use of an appropriate query against that new uniqueId will allow the requestor access to updated metadata about the document, including size, hash, etc. This workflow assumes that some change to the underlying data occurs between the time Community A retrieves on-demand document entry #5, and Community C
5015 retrieves on-demand document entry #5. That is why #6 is not re-used in response to Community C's request.

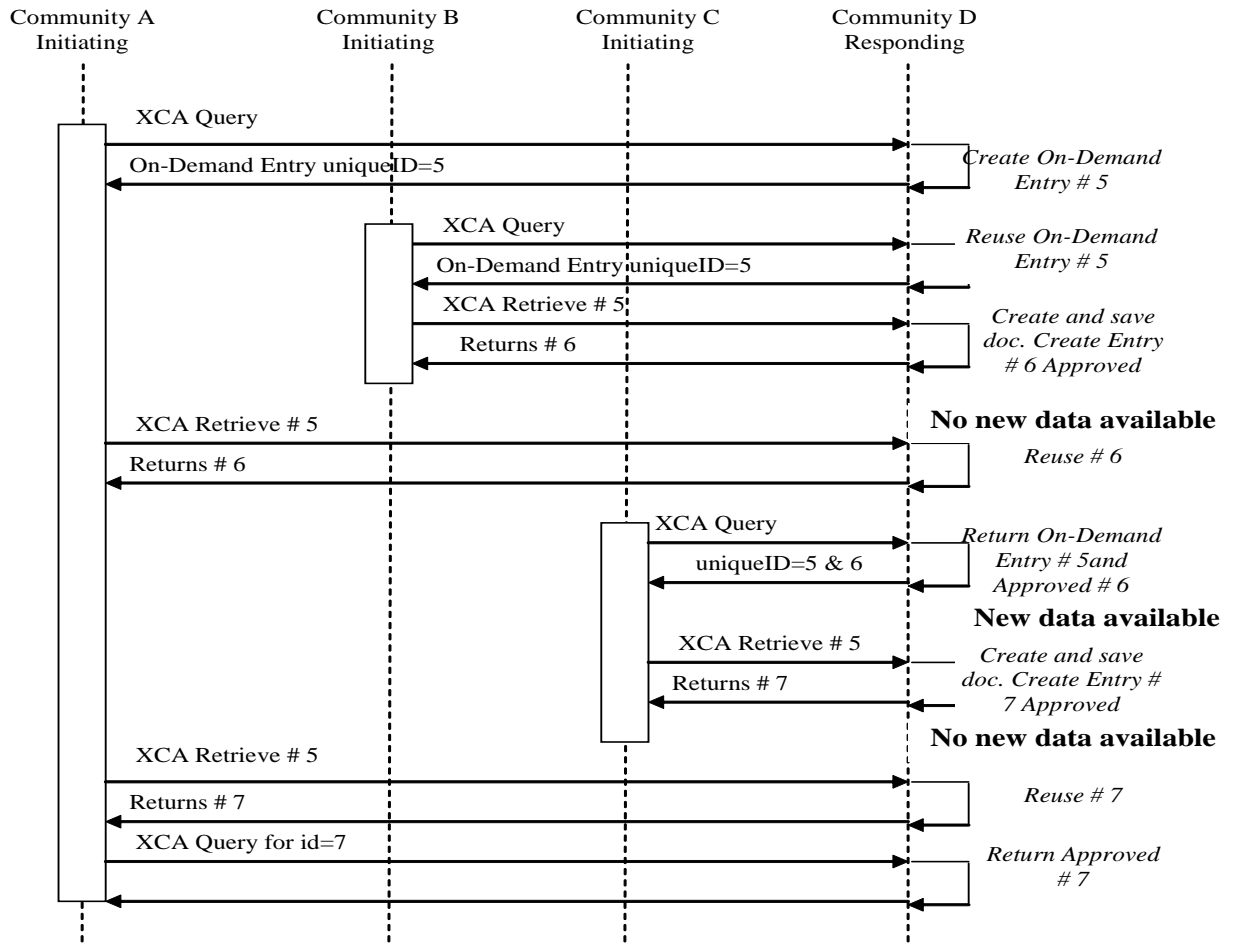


Figure 18.3.3-2: Dynamically created content with persistence

5020

Note: Figure 18.3.3-2 is a diagram of a possible interaction, not the required interaction. In particular, the ability of the Responding Gateway to determine whether there is new data available or not is an implementation detail. In the case where a Responding Gateway is able to make that determination it should work as presented. If the Responding Gateway is not able to make that determination then it is free to create a new document at every retrieve request.

18.4 XCA Security Considerations

5025

18.4.1 XCA Risk Assessment

The risk analysis for XCA enumerates assets, threats, and mitigations. The complete risk data is stored and maintained in a central location. The complete risk data is stored and available from IHE⁶.

⁶ The risk analysis data may be found at: ftp://ftp.ihe.net/IT_Infrastructure/iheityr5-2007-2008/Technical_Cmte/Profile_Work/XC/XCARiskAnalysis.xls

5030 The purpose of this risk assessment is to notify vendors of some of the risks that they are advised to consider in implementing XCA actors. For general IHE risks and threats please see ITI TF-1: Appendix L. The vendor is also advised that many risks cannot be mitigated by the IHE profile and instead the responsibility for mitigation is transferred to the vendor, and occasionally to the XDS Affinity Domain and enterprises. In these instances, IHE fulfills its responsibility to notify affected parties through the following section.

5035 **18.4.2 Requirements/Recommendations**

The following mitigations shall be implemented by all XCA actors. These mitigations moderate all high impact risks.

- **M1:** All actors in XCA shall be grouped with an ATNA Secure Node Actor (or Secure Application) and a CT Time Client Actor.
- 5040 • **M2:** Document metadata shall include a SHA1 hash of the document content. Applications shall have the ability to verify the SHA1 hash of the document with the SHA1 hash in the metadata, if corruption detection is requested.
- **M3:** Document Consumer implementations shall handle overloading through excessive volume of response data by discontinuing the read on the socket and closing it. The Initiating and Responding Gateways shall respond to disconnection by discontinuing processing of responses.
- 5045 • **M4:** Document Consumer implementations shall not issue a Registry Stored Query that is not patient specific, i.e., it shall either supply a patient identifier or a unique document entry identifier.
- 5050 • **M6:** Queries of unknown patient identifiers shall return either zero documents with no further information or XDSUnknownPatientId, depending on local policy. This applies to patient identifiers that are properly formatted or improperly formatted. By not using an error code indicating that the identifier is ill-formatted, you are able to reduce the ability of applications to fish for data. This applies only to Responding Gateways, if appropriate.

5055 The following mitigations address the risk of a document being maliciously changed. This mitigation is optional.

- **M5:** Documents may be digitally signed using the DSG Profile

The following mitigations are transferred to the vendors, XDS Affinity Domains, and enterprises.

- 5060 • **T1:** Backup systems for registry metadata, repository documents, and gateway configuration are recommended.
- **T2:** All implementations are recommended to ensure that all received data is propagated appropriately (i.e., without corruption and complete results) or an error is presented.
- 5065 • **T3:** Network protection services are recommended to be sufficient to guard against denial of service attacks on all service interfaces.

- **T4:** A process that reviews audit records and acts on inappropriate actions is recommended.
- **T5:** It is recommended that service interfaces be implemented with a good design to guard against corruption and denial of service attacks

5070 **18.4.3 Policy Choices**

Policy choices will not be addressed by this profile. Each community may have different policies. The profile has been designed with this fact in mind and an understanding of enough variety of policies so that any reasonable policy can be implemented without violating the profile.

5075

19 Basic Patient Privacy Consents Integration Profile

5080 The document sharing infrastructure provided by XD* allow for the publication and use of clinical documents associated with a patient. This profile allows for a Patient Privacy Policy Domain (e.g., an XDS Affinity Domain to have a number of Patient Privacy Policies that can be acknowledged (aka consent). This allows for more flexibility to support some patient concerns, while providing an important and useful dataset to the healthcare provider. Without BPPC, the XDS Profile requires that the administrators of an XDS Affinity Domain creates and agrees to a single document publication and use policy (See ITI TF-1: Appendix L). Such a single XDS Affinity Domain Policy is enforced in a distributed way through the inherent access controls of the systems involved in the XDS Affinity Domain.

5085 This profile will use terms consistent with ISO 22600 - Privilege Management and Access Control (PMAC), but is not restricted to systems that implement PMAC. This profile uses the term “Patient” to refer to the human-subject of health related data. In this context “Patient” is not to imply only those subjects under current treatment, this is sometimes referred to as “consumer”. This profile uses the term “Consent” to mean acknowledgement of a privacy policy, also known as an information access policy. In this context the privacy policy may include constraints and obligations. The systems involved in XDS are expected to support sufficient Access Controls to carry out the Policy of the XDS Affinity Domain⁷.

5095 Healthcare providers utilize many different sets of data to carry out treatment, billing, and normal operations. This information may include patient demographics, contacts, insurance information, dietary requirements, general clinical information and sensitive clinical information. This information may be published (e.g., to XDS, XDR, XDM, PACS) as independent documents with different sensitivity labels (i.e., confidentialityCode). This mechanism is not unique to BPPC, but is leveraged by privacy and security policies.

5100 Healthcare providers in different functional roles will have different needs to access these documents. For example, administrators may need to be able to access the patient demographics, billing and contact documents. Dietary staff will need access to the dietary documents but would not need access to insurance documents. General care providers will want access to most clinical documents, and direct care providers should have access to all clinical documents. This is an example of a Patient Privacy Policy that would be given a Patient Privacy Policy Identifier within a Patient Privacy Policy Domain. When a patient acknowledges this policy, the Patient Privacy Policy document would refer to the policy by the Patient Privacy Policy Identifier.

5110 This profile provides a mechanism by which an XDS Affinity Domain can create a basic vocabulary of codes that identify Patient Privacy Domain managed Privacy Policy Identifiers with respect to document sharing. Each Privacy Policy Identifier uniquely identifies a Privacy Policy which should identify in legal text what the acceptable use, re-disclosure uses, which functional roles may access which document and under which conditions, etc. The administration of the XDS Affinity Domain will assign each Privacy Policy Identifiers for use within the XDS

⁷ See the IHE white paper “HIE Security and Privacy through IHE” published on the IHE web site http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Whitepaper_Security_and_Privacy_2007_07_18.pdf

5115 Affinity Domain. Future profiles may include in addition to the legal text, a structured and coded expression of the consent policy that can be used to support even more dynamic understanding of the patient's directives (see HL7 and OASIS).

19.1 Basic Patient Privacy Consent Use-Cases

5120 This section gives examples of some possible patient privacy consent policies and how the systems publishing documents and using documents might act. This is an informative section and should not be interpreted as the only way to implement the BPPC Profile. Its purpose is to allow implementers of BPPC to more easily understand the principle of operation of BPPC.

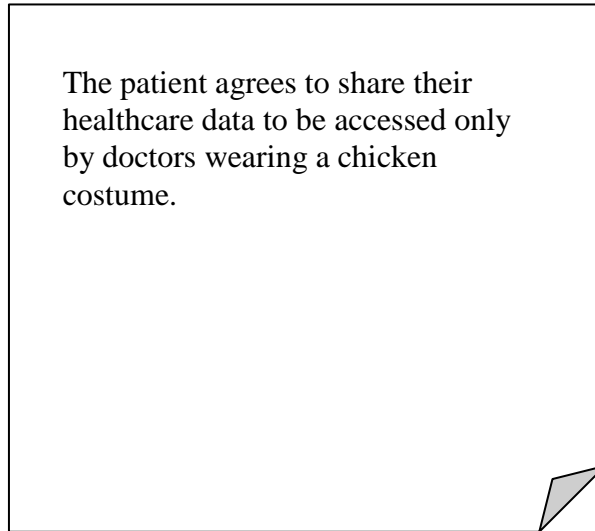
19.1.1 Implied Consent vs. Explicit Consent

5125 This profile supports both Implied Consent as well as Explicit Consent environments. In order to provide a profile with global appeal we have supported both environments. In an implied consent environment it would be normal for a Document Consumer to find no instance of a patient specific acknowledgement of a privacy consent policy in the XDS Affinity Domain, as capturing the act of acknowledging a privacy consent policy would not be required. Note: this may also be true in an Explicit Consent environment, where obtaining the acknowledgement is delayed due to medical reasons (e.g., emergency).

5130 An XDS Affinity Domain might have a paper document that describes their Privacy Consent Policy. In our example this Privacy Consent Policy will be given a local XDS Affinity Domain managed Privacy Policy Identifier (e.g., an OID such as: 9.8.7.6.5.4.3.2.1). The example in Figure 19.1-1 is ridiculous (i.e., chicken costume) but is provided to emphasize that IHE doesn't write these policies, and to make clear that the BPPC Profile could be used to enforce any policy that could be written in human readable form, provided that all actors can be configured to enforce that policy. This example also points out that the content of the policy is human readable

5135

text, and that we provide no structured or coded way to interpret. This example policy might look



like:

Figure 19.1-1: Policy Example

5140 **19.1.1.1 Opt-In**

A common structure for sharing clinical documents requires that the patient first acknowledge that they want this sharing to happen before any documents are actually shared. In this case the XDS Affinity Domain administrators would write a policy that indicates what should be shared, when it should be shared, when it can be used, etc. There would also be an overriding XDS Affinity Domain policy that indicates that no document will be shared until the patient has explicitly chosen to participate.

5145

19.1.1.2 Opt-Out

Equally as common is a structure for sharing documents that presumes that when the patient chooses to get care within a care setting, that they are implicitly agreeing to the normal sharing of their documents for treatment purposes. In this environment, there is usually a control that allows a patient to choose to NOT participate in this sharing. This is commonly referred to as “opt-out”.

5150

In this case the existence an acknowledgement to an opt-out policy would mean that documents should no longer be shared, and any documents that might appear should not be used. Clearly the XDS Affinity Domain administrators need to make the actual behavior clear in their policies.

5155

19.1.2 Wet Signature

An XDS Affinity Domain might have the patient acknowledge the consent through ink on paper. For Example:


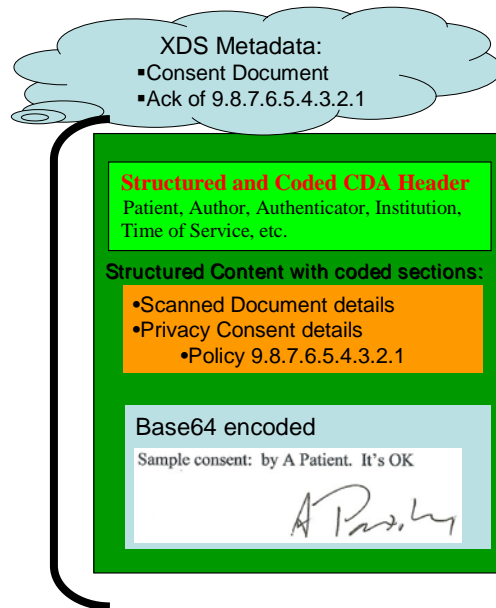


Figure 19-2: Simplistic Consent Example

5160 This acknowledgement is captured according to the XDS Scanned Document Content Profile (XDS-SD), with the additional parameters specified in the BPPC Content Profile also applied. This is submitted into the XDS Affinity Domain as proof that the patient has acknowledged policy 9.8.7.6.5.4.3.2.1.

The following shows this graphically:



5165

Figure 19.1.2-1: Graphical representation of consent with wet signature

If an XDS Affinity Domain wants to further provide non-repudiation protections it may choose to apply a digital signature using the IHE-DSG Content Profile to the whole package with the appropriate purpose and signed by an appropriate signing system/person.

5170 **19.1.3 Advanced Patient Privacy Consents**

An XDS Affinity Domain may have jurisdictional or organizational policies that require support for more complex patient privacy consent policies. These privacy policies may require that a patient explicitly consent to disclosure of protected or sensitive health information to specific

5175 entities. The BPPC Profile provides a starting point for implementing these types of privacy consent policies, but does not explicitly specify how additional information needed to enforce the policy would be conveyed. In these cases, the capability of BPPC may not be enough to support all types of needs. An example of an Advanced Patient Privacy Consent would be when a patient wants to name individuals that can access their documents.

19.2 Creating Patient Privacy Policies

5180 The administrators of the Patient Privacy Policy Domain (e.g., XDS Affinity Domain) will need to develop and publish an overall Policy for the Patient Privacy Policy Domain that clearly defines the overall appropriate use of the protected resources. This is the subject of ITI TF-1: Appendix L and is not further defined here.

5185 Within this Patient Privacy Policy Domain (e.g., XDS Affinity Domain) overall Policy is a defined set of acceptable use Patient Privacy Policies. A Patient Privacy Policy further explains appropriate use of the protected resources in a way that provides choices to the patient. The BPPC Profile places no requirements on the content of these policies nor the method used to develop these policies (See ITI TF-1: Appendix P for some guidance on developing these policies). BPPC only assumes that the overall Patient Privacy Policy Domain can be structured as a set of specific policies (A, B, C, D in the example below), where each one may be used independently or combined in relationship to publication and access of a specific type(s) of document.

5190

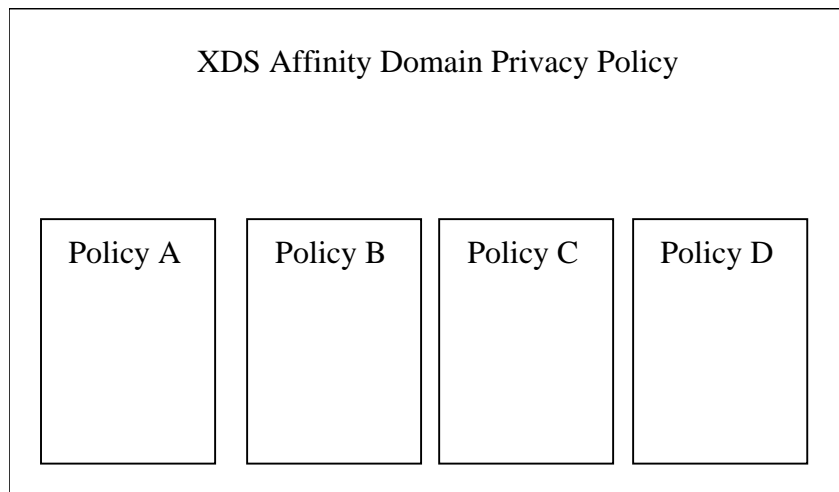


Figure 19.2-1: Example Patient Privacy Policy Hierarchy

5195 A Patient Privacy Policy will identify who has access to information, and what information is governed by the policy (e.g., under what conditions will a document marked as containing sensitive information be used by a specific type of individual for a specific use). The mechanism for publishing these policies is not described by this profile. The set of Patient Privacy Policies written by the Patient Privacy Policy Domain must be able to be implemented by the technologies in all of the systems that have access to the domain. This means that the Patient Privacy Policies must be created with great care to ensure they are enforceable.

5200

5205 Each Patient Privacy Policy will be given a unique identifier (OID) known as a Patient Privacy Policy Identifier. This is additionally used when capturing a patient's acknowledgement of a specific Patient Privacy Policy resulting in a Patient Privacy Policy Acknowledgement Document (i.e., an instance of a BPPC document).

5210 Finally, Privacy Consent Policies used within an XDS Affinity Domain will very likely be different than those used with the XDM or XDR Profiles as these profiles often are used to transfer documents in ad-hoc ways. The patient may provide a consent given to share information on media to the provider creating the media for specific use, rather than for more general sharing within an XDS Affinity Domain. When transferring information that originated in an XDS Affinity Domain to media (XDM), the Privacy Consent Policies found in the XDS Affinity Domain might be changed during the publication process. There are also differences in the sensitivity that should be considered for consents shared on media or transmitted through XDR and those shared in an XDS Affinity Domain. See the section Security Considerations later in this volume for more details.

19.2.1 Summary of the creation and publication of the policies

1. The Patient Privacy Policy Domain will write and agree to overall privacy policies (lots of lawyers involved).
- 5220 2. The Patient Privacy Policy Domain will include a small set of Patient Privacy Policies (more lawyers). These are text documents very similar to the privacy consent documents used today.
3. Each Patient Privacy Policy will be given an unique identifier (OID) called the Patient Privacy Policy Identifier
- 5225 4. The Policy of the Patient Privacy Policy Domain and all of the Patient Privacy Policies will be published in some way. It is expected that this will be sufficiently public to support local regulation.
5. When a patient acknowledges a Patient Privacy Policy, a Patient Privacy Policy Acknowledgement Document will be published with the Patient Privacy Policy Identifier of the policy that the patient acknowledged.

5230 19.3 Actors/Transactions

There are two actors in the BPPC Profile, the Content Creator and the Content Consumer. Content is created by a Content Creator and is to be consumed by a Content Consumer. The sharing or transmission of content or updates from one actor to the other is addressed by the use of appropriate IHE profiles described in the section on Content Bindings with XDS, XDM and XDR. in PCC TF-2: 4.1, and is out of scope of this profile. A Document Source or a Portable Media Creator may embody the Content Creator Actor. A Document Consumer, a Document Recipient or a Portable Media Importer may embody the Content Consumer Actor.



Figure 19.3-1: BPPC Actor Diagram

5240 Table 19.3-1 lists the transactions for each actor directly involved in the BPPC Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 19.4.

5245

Table 19.3-1: BPPC Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Content Creator	<i>Share Content</i>	R (note 2)	ITI TF-1: 19.4.3 ITI TF-1: 19.4.4
Content Consumer	<i>Share Content</i>	R (note 3)	ITI TF-1: 19.4.5

19.3.1 Grouping

19.3.1.1 Basic Patient Privacy Documents Bindings to XDS, XDR, XDM

5250 A BPPC Content Creator or Content Consumer can be grouped with appropriate actors from the XDS, XDM or XDR profiles to exchange Basic Privacy Consent documents. The metadata sent in the document sharing or interchange messages has specific relationships or dependencies (which we call bindings) to the content of the clinical document – a Basic Patient Privacy Consent document - described in ITI TF-3: 5.1.2 and 5.1.3.

- 5255
- A BPPC Content Creator shall be grouped with and XDS/XDR Document Source or and XDM Portable Media Creator.
 - A BPPC Content Consumer shall be grouped with an XDS Document Consumer, and XDR Document Recipient, or an XDM Portable Media Importer.

19.3.1.2 Basic Patient Privacy Grouping with XDS-SD

5260 The BPPC Content Consumer shall be grouped with a XDS-SD Content Consumer. This means that a Content Consumer for BPPC Content must also be able to display XDS-SD content. This is required due to the common practice of capturing Wet Signatures.

19.4 Basic Patient Privacy Consent Profile Options

Options that may be selected for this Integration Profile are listed in Table 19.4-1 along with the IHE actors to which they apply.

Table 19.4-1: Basic Patient Privacy Consents - Actors and Options

Actors	Option	Section
Content Creator	<i>Basic Patient Privacy Acknowledgement (note 1)</i>	ITI TF-3: 5.1.2
	<i>Basic Patient Privacy Acknowledgement with Scanned Document</i>	ITI TF-3: 5.1.3
Content Consumer	<i>Basic Patient Privacy Acknowledgement View (note 2)</i>	ITI TF-3: 5.1.2 ITI TF-3: 5.1.3

Note 1: Content Creator shall implement the Basic Patient Privacy Acknowledgement Option, and may choose to implement the Basic Patient Privacy Acknowledgement with Scanned Document Option

Note 2: Content Consumer shall implement the Basic Patient Privacy Acknowledgement View Option.

19.4.1 Intentionally Left Blank

19.4.2 Intentionally Left Blank

19.4.3 Basic Patient Privacy Acknowledgement Option

The Content Creator shall be able to create Patient Privacy Policy Acknowledgement Document Content as specified in ITI TF-3: 5.2.

A Patient Privacy Policy Acknowledgement Document is a kind of medical document. The content of a Patient Privacy Policy Acknowledgement Document shall include the effective time of the acknowledgement and Patient Privacy Policy Domain (e.g., XDS Affinity Domain) defined coded vocabulary identifying the Patient Privacy Consent Policy Identifier (OID) acknowledged by the patient. The content of the Patient Privacy Acknowledgement Document may include a text description of what the patient has acknowledged.

The Patient Privacy Policy Acknowledgement Document may be signed. There are cases, as seen in the use-cases, where the Content Creator would need to be grouped with a DSG Content Creator. The BPPC Profile does not require this grouping. This grouping can be fully specified in an IHE Integration Statement.

19.4.4 Basic Patient Privacy Acknowledgement with Scanned Document Option

A Basic Patient Privacy Policy Acknowledgement Document may include a scanned document. An example of the scanned document could be a wet signature by the patient on the text. The Content Creator that claims to support Basic Patient Privacy Policy Acknowledgement with Scanned Document Option shall be able to create a Patient Privacy Policy Acknowledgement with Scanned Document Content as specified in ITI TF-3: 5.1.3.

19.4.5 Patient Privacy Acknowledgement View Option

The Content Consumer shall be able to display the Patient Privacy Policy Acknowledgement Document Content as specified in ITI TF-3: 5.1.2 and ITI TF-3: 5.1.3.

5295 **19.5 Intentionally Left Blank**

19.6 BPPC Process Flow in an XDS Affinity Domain

This flow shows how an XDS Affinity Domain would use the BPPC Profile. Only a basic flow is shown, the profile supports many alternative flows.

19.6.1 Checking for a patient's acknowledgement of a privacy policy

5300 An XDS Document Consumer Actor that is enforcing policies registered by BPPC can query an XDS Affinity Domain for instances of Patient Privacy Policy Acknowledgement Documents that have been acknowledged by a specific patient. Through the XDS Metadata the Document Consumer can determine which Patient Privacy Policies have been acknowledged.

5305 Note if the local regulations allow, some XDS Affinity Domains may not publish the consent documents, so systems should be able to handle the configurations where no Patient Privacy Acknowledgment Document is in the XDS Affinity Domain for a specific patient (e.g., implied consent).

5310 Note if the local regulations allow, some patients may have documents shared before informed consent can be captured. In this case the XDS Affinity Domain policy needs to explain the default behavior, that behavior for the absence of a consent document.

19.6.2 Recording a patient's acknowledgement of a privacy policy

The Content Consumer Actor creates Patient Privacy Policy Acknowledgement Documents with or without a scanned document part. This document records the patient's acknowledgement of a specified policy.

5315 **19.6.3 Publishing documents against a consent policy**

All documents managed in an XDS Affinity Domain, or transferred using XDM/XDR, are labeled with a confidentialityCode. The administrators of an XDS Affinity Domain may need to define a vocabulary and meaning to that vocabulary.

5320 The XDS or XDR Document Source Actor that supports the Basic Patient Privacy Enforcement Option determines which of the XDS Affinity Domain – Privacy Consent Policies would allow the documents to be published. In some XDS Affinity domains this may require that the system check that a patient has indeed acknowledged a specific policy.

The Document Source Actor will set the XDS Metadata – confidentialityCode - to indicate the appropriate sensitivity for use/constraint (determined by the XDS Affinity Domain Policy)

5325 The XDS Document Registry validates that each of the confidentialityCode(s) are from the approved list of confidentialityCode for use within the XDS Affinity Domain.

19.6.4 Using published documents

5330 When an XDS Document Consumer that supports the Basic Patient Privacy Enforcement Option queries the XDS Affinity Domain it may utilize the confidentialityCode filter in the Registry Stored Query to restrict the documents returned to those that the Document Consumer can utilize.

The Document Consumer will enforce access controls based on the returned XDS metadata-confidentialityCode, current state of consent acknowledgements, system type, user, context, and any number of other factors that the system is capable of enforcing.

5335 The Document Consumer may be capable of querying for ‘Approved’ consent acknowledgement documents and using the resulting XDS Metadata as the list of currently Approved Patient Privacy Policy Acknowledgement Documents. There is no requirement for the Document Consumer system to retrieve the Patient Privacy Policy Acknowledgement Document content.

19.7 Security Considerations

5340 Consents stored in an XDS Affinity Domain are also governed by privacy policies. The content of a Patient Privacy Policy Acknowledgement Document may itself contain sensitive information. For example, a terminally ill patient may decide that his prognosis should not be shared with his family members, but that other information may be. Sharing the Patient Privacy Policy Acknowledgement Document with family members would potentially inform them of a
5345 negative prognosis. Thus the confidentialityCode placed on Patient Privacy Policy Acknowledgement Documents must be appropriately assigned (e.g., most will be assigned the broadest use confidentialityCode).

5350 However, Patient Privacy Policy Acknowledgement Documents stored in the clear on media (XDM), or transmitted through XDR, should not contain sensitive information. The rationale is that the receiver of the information must be able to read the consent that was used to share this information in order to understand how they must treat the information with respect to their own Patient Privacy Policies.

5355 Implementation of Patient Privacy Policies within a healthcare environment has different considerations and risks than implementing similar access control policies within other non-treatment environments. This is for the simple reason that failing to provide access to critical healthcare information has the risk of causing serious injury or death to a patient. This risk must be balanced against the risk of prosecution or lawsuit due to accidental or malicious disclosure of private information. The XDS Affinity Domain should take care in writing their Patient Privacy Policies to avoid this.

5360 One mitigation strategy that is often adopted in healthcare provides accountability through audit controls. That is to say that the healthcare providers are trusted not to abuse their access to private information, but that this is followed up by a policy of monitoring healthcare provider accesses to private information to ensure that abuse does not occur. This strategy reduces the risk of serious death or injury due to lack of access to critical healthcare information, at the increased
5365 risk of disclosure of private information. This is why the ITI Technical Committee created the Audit Trail and Node Authentication (ATNA) Integration Profile, and furthermore, why that profile is a requirement of XDS and related profiles.

5370 Another risk that must be resolved by an affinity domain is how to address the issues of sharing truly sensitive information in a registry (e.g., psychology documents). One strategy that might be recommended is that truly sensitive data not be shared within the XDS Affinity Domain; directed communications using XDR or XDM may be more appropriate.

20 Cross-Enterprise Sharing of Scanned Documents Content Integration Profile

5375 A variety of legacy paper, film, electronic and scanner outputted formats are used to store and exchange clinical documents. These formats are not designed for healthcare documentation, and furthermore, do not have a uniform mechanism to store healthcare metadata associated with the documents, including patient identifiers, demographics, encounter, order or service information. The association of structured, healthcare metadata with this kind of document is important to maintain the integrity of the patient health record as managed by the source system. It is
5380 necessary to provide a mechanism that allows such source metadata to be stored with the document.

This profile defines how to couple such information, represented within a structured HL7 CDA R2 header, with a PDF or plaintext formatted document containing clinical information. Furthermore, this profile defines elements of the CDA R2 header necessary to minimally
5385 annotate these documents. Such header elements include information regarding patient identity, patient demographics, scanner operator identity, scanning technology, scan time as well as best available authoring information. Portions of CDA R2 header, along with supplemental document registration information, are then used to populate XDS Document Entry metadata.

The content of this profile is intended for use in XDS, XDR and XDM. Content is created by a
5390 Content Creator and is to be consumed by a Content Consumer. The Content Creator can be embodied by a Document Source Actor or a Portable Media Creator, and the Content Consumer by a Document Consumer, a Document Recipient or a Portable Media Importer. Obligations imposed on the Content Creator and the Content Consumer by this profile are understood to be fulfilled by the software that creates the final document for submission and/or consumes profile
5395 conformant documents rather than any particular scanning technology.

20.1 Use Cases

20.1.1 Content Use Cases

Text Chart Notes

5400 Examples of this content include handwritten, typed or word processed clinical documents and/or chart notes. These documents are typically multi-page, narrative text. They include preprinted forms with handwritten responses, printed documents, and typed and/or word processed documents, and documents saved in various word processing formats. Appropriate formats are PDF, derived from the word processing format, or plaintext, if the text structure is all that needs to be conveyed. PDF is desirable because it most faithfully renders word processed
5405 document content and it preserves meaning embodied in non-textual annotations.

Graphs, Charts and/or Line Drawings

5410 Examples of this content include Growth Charts, Fetal Monitoring Graphs. Line drawings such as those described above are best rendered using PDF versus an image based compression, such as JPEG. However, when computer generated PDFs include lines or lossy compression is not acceptable for diagnostic purposes, PDF should be used.

Object Character Recognition (OCR) Scanned Documents

5415 Clinical documents can contain text and annotations that cannot be fully processed by optical character recognition (OCR). We call attention to the fact that the OCR text content may only partially represent the document content. These are best supported by converting to PDF format, which can mix the use of OCR'd text, compressed scanned text, and scanned image areas.

Electronic Documents

5420 Existing clinical documents that are electronically transmitted or software created (e.g., PDF, or plaintext) can be considered as actually scanned, previously scanned or virtually scanned before they are shared. In this context, “actually scanned” refers to electronic documents, newly created via some scanning technology from legacy paper or film for the purposes of sharing. “Previously scanned” refers to electronic documents that were previously produced via some scanning technology from legacy paper or film, but have existed in their own right for a period of time. “Virtually Scanned” electronic documents are existing electronic documents not derived from legacy paper or film that either are PDF/A or plaintext format or have been converted to one of these formats for the purposes of sharing. This content is covered by this profile.

5425

20.1.2 Content Creator Use Cases

5430 Content is created by a Content Creator. Impact on application function and workflow is implementation specific and out of scope of this content profile, though we note that they will be compliant with this content profile if they can produce CDA wrapped PDF, CDA wrapped plaintext or both. The following example use case is included to aid in the scoping of this content profile.

5435 Legacy Clinic is a small two-physician clinic. They presently store their patient's medical records on paper. The Clinic is trying to figure out what to do with its paper and word processing documents as it converts over to an electronic system. They would like to be able to view the files over their local intranet.

5440 Presently, most records are handwritten on preprinted paper forms that are inserted into specific sections of the patient's chart. More detailed encounter reports are dictated and sent to a transcription company that returns them in a word processing format. The medical records clerk at Legacy Clinic receives these files via e-mail, decrypts them, prints them out, and adds them to the patient's chart in the correct section.

5445 Over the years, Legacy Clinic has used a number of different transcription companies, and the documents are stored in a variety of word processing formats. Several years ago, they began to require that returned documents be in RTF format in an attempt to reduce frustrations induced by dealing with discrepant word processing formats. Only in some cases was patient and encounter metadata stored within the word processing document in a regular format, depending upon the transcription company used at the time. A third party presently handles labs for the clinic. These are usually returned to the Clinic as printed documents. The clerk inserts these into the labs section in the patient's chart.

5450 In the case of Legacy Clinic, the link between the word processing documents and the patient has been maintained for many of its documents, since the existing manual process maintains

5455 that association, and some of the files also contain the encounter metadata. However, the link to the specific encounter will need to be reestablished by interpreting the document content, which will require a great deal of manual effort for some of their documents which do not have it, and will still require custom handling depending upon the format used to store this metadata.

Legacy Clinic uses a transcription provider that can generate PDF documents, wrapped in a CDA Release 2.0 header. These are sent to Legacy Clinic via e-mail. While the same manual process is used, these documents are now in a format that is ready to be used by their new EHR system.

5460 20.1.3 Content Consumer Use Cases

Content is consumed by a Content Consumer. Impact on application function and workflow is implementation specific and out of scope of this content profile. However, we note that adoption of this profile will necessitate the Content Consumer, upon document receipt, support the processing of both CDA wrapped PDF and CDA wrapped plaintext.

5465 20.2 Actors/ Transactions

5470 There are two actors in the XDS-SD Profile, the Content Creator and the Content Consumer. Content is created by a Content Creator and is to be consumed by a Content Consumer. The sharing or transmission of content from one actor to the other is addressed by the appropriate use of IHE profiles described below, and is out of scope of this profile. A Document Source or a Portable Media Creator may embody the Content Creator Actor. A Document Consumer, a Document Recipient or a Portable Media Importer may embody the Content Consumer Actor. The sharing or transmission of content or updates from one actor to the other is addressed by the use of appropriate IHE profiles described in the section on Content Bindings with XDS, XDM and XDR.

5475 Figure 20.2-1 shows the actors directly involved in the Scanned Documents Content Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other profiles are not necessarily shown.



5480 **Figure 20.2-1: Scanned Documents Actor Diagram**

20.3 Scanned Documents Content Integration Profile Options

Options for Scanned Documents leverage those in the Patient Care Coordination (PCC) Technical Framework (TF). Options that may be selected for this Integration Profile are listed in the Table 20.3-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 20.3-1: XDS-SD - Actors and Options

Actor	Options	Vol. & Section
Content Creator	<i>No options defined</i>	--
Content Consumer	<i>View Option¹</i>	PCC TF- 2: 3.1.1
	<i>Document Import Option¹</i>	PCC TF-2: 3.1.2

Note 1: The Actor shall support at least one of these options.

20.4 Scanned Documents Bindings to XDS, XDR, XDM

Actors from the ITI XDS, XDM and XDR profiles embody the Content Creator and Content Consumer sharing function of this profile. A Content Creator or Content Consumer may be grouped with appropriate actors from the XDS, XDM or XDR profiles to exchange the content described therein. The metadata sent in the document sharing or interchange messages has specific relationships or dependencies (which we call bindings) to the content of the clinical document described in the content profile. The Patient Care Coordination Technical Framework (PCC-TF) defines the bindings to use when grouping the Content Creator of this Profile with actors from the IHE ITI XDS, XDM or XDR Integration Profiles. See PCC TF-2:4.

5500 **20.5 Scanned Documents Content Process Flow**

This profile assumes the following sequence of events in creation of an XDS-SD document.

1. A legacy paper document is scanned and a PDF/A is rendered. Alternatively, an electronic document is converted, if necessary, to PDF/A or plaintext format (see ITI TF-3: 5.2.1 and 5.2.1.1).
- 5505 2. Software, conformant to this profile and most likely with the aid of user input (e.g., to provide document title, confidentiality code, original author), renders the CDA R2 header pertaining to the PDF or plaintext produced. The document is wrapped and the XDS-SD document is completed (see ITI TF-3: 5.2.3).
- 5510 3. XDS metadata is produced from data contained in the CDA header and supplemental information (see ITI TF-3: 5.2.2).
4. The completed XDS-SD document and corresponding metadata is sent via the Provide a Register Document Set Transaction [ITI-15] or [ITI-41] of XDS/XDR, or the Distribute Document Set on Media Transaction [ITI-32] of XDM.

21 Sharing Value Sets Integration Profile (SVS)

- 5515 The Sharing Value Sets (SVS) Profile provides a means through which healthcare systems producing clinical or administrative data, such as diagnostic imaging equipment, laboratory reporting systems, primary care physician office EMR systems, or national healthcare record systems, can receive a common, uniform nomenclature managed centrally. Shared nomenclatures are essential to achieving semantic interoperability.
- 5520 A single Value Set Repository can be accessed by many Value Set Consumers, establishing a domain of consistent and uniform set of nomenclatures. It supports automated loading of Value Sets by Value Set Consumers, reducing the burden of manual configuration. This profile describes two Transactions for retrieving Value Sets from a Value Set Repository by a Value Set Consumer.
- 5525
- A single value set can be retrieved based on an OID value. This is aimed at meeting the needs of systems that are pre-configured to use specific value sets. These systems are often medical devices with strictly controlled functions that should not be modified without careful review. This transaction does not include metadata content, and provides just the value set concept list as uniquely identified in the request.
- 5530
- Multiple value sets can be retrieved based on metadata about the value sets. This is aimed at meeting the needs of systems and users that will be dynamically selecting value sets, deciding which value sets should be used, and creating new value sets based on the contents of existing value sets. This transaction supports a much richer selection criteria and provides metadata descriptions as well as the contents (expanded lists of coded values) of all the value sets that meet those criteria.
- 5535

Both transactions provide access to centrally managed value sets that have been assigned metadata, including group identification. The ability to identify groups of value sets is essential to achieving semantic interoperability and development of modular structures of electronic health records (EHR). Group identification can be used to identify, for example, all the value sets needed for a given purpose like filling in a particular kind of report.

5540

21.1 Actors/ Transactions

- 5545 Figure 21.1-1 shows the actors directly involved in the SVS Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other related profiles are not necessarily shown. As well, the method for creating a Value Set is not covered by this profile (this subject will be addressed once the basic infrastructure is in place).

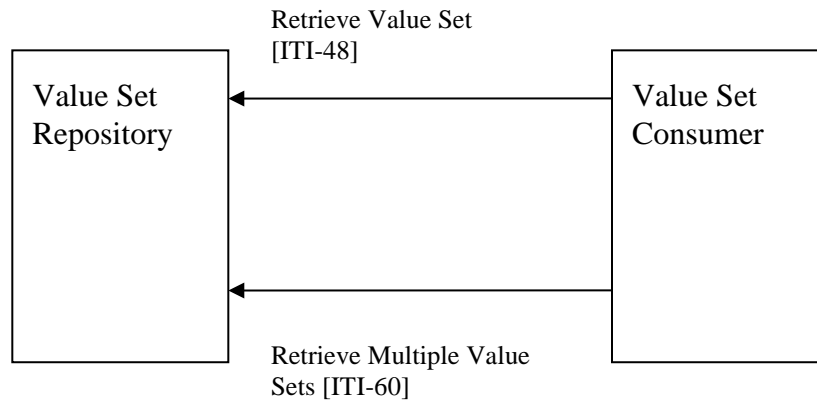


Figure 21.1-1: Actors and Transactions

5550 Table 21.1-1 SVS Integration Profile - Actors and Transactions lists the transactions for each actor directly involved in the SVS Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile is listed in Table 21.2-1.

5555

Table 21.1-1: SVS Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Value Set Repository	Retrieve Value Set [ITI-48]	R	ITI TF-2b: 3.48
	Retrieve Multiple Value Sets [ITI-60]	R	ITI TF-2b:3.60
Value Set Consumer	Retrieve Value Set [ITI-48]	R	ITI TF-2b: 3.48
	Retrieve Multiple Value Sets [ITI-60]	O	ITI TF-2b:3.60

21.1.1 Assumptions and background information

5560 A Value Set is a uniquely identifiable set of valid concept representations. A Value Set may be a simple flat list of concept codes drawn from a single code system, or it might be constituted by expressions drawn from multiple code systems (a code system is a system consisting of

designations and meanings, for example LOINC, SNOMED-CT, ICD-10, or ISO 639 Language Codes).

5565 This profile will address a flat list of concept codes, one of the simplest examples of a Value Set being shown in Table 21.1.1-1: The provinces of Canada.

Table 21.1.1-1: The provinces of Canada

Provinces of Canada ISO Code	Print Name
NL	Newfoundland
AB	Alberta
BC	British Columbia
SK	Saskatchewan
MB	Manitoba
ON	Ontario
QC	Quebec
NB	New Brunswick
NS	Nova Scotia
PE	Prince Edward Island

21.1.2 Value Set Unique ID and Value Set Version

5570 A Value Set must be uniquely identified to allow various applications and users to recognize it. When a Value Set is retrieved, the application or the user is retrieving a particular instance of it, or an *Expanded Value Set* (an Expanded Value Set is a set of concept representations that were in effect at a specific time for a particular version of a Value Set definition. The *Value Set* (definition) and the *Expanded Value Set* concepts are similar to the programming concepts of Class and Instance of Class.)

5575 This profile uses the *Value Set Unique ID* (using an ISO OID), and the *Value Set Version* attributes to allow flexible handling of the identification of a Value Set.

The actual set of codes derived from this definition of a Value set is an *Expanded Value Set*. SVS supports the sharing of Expanded Value Set with two different approaches to their identification:

- 5580 1. By unique identification of the Expanded Value Set itself, and no reference to the definition that produced it. Such an Expanded Value Set carries its own unique identifier (i.e., an OID and Version).
2. By reference to the Value Set definition (OID and Version) from which the Expanded Value Set was derived. In this case specific Expanded Value sets (derived from the same Value Set definition) are only distinguished by their expansion date and time.

5585

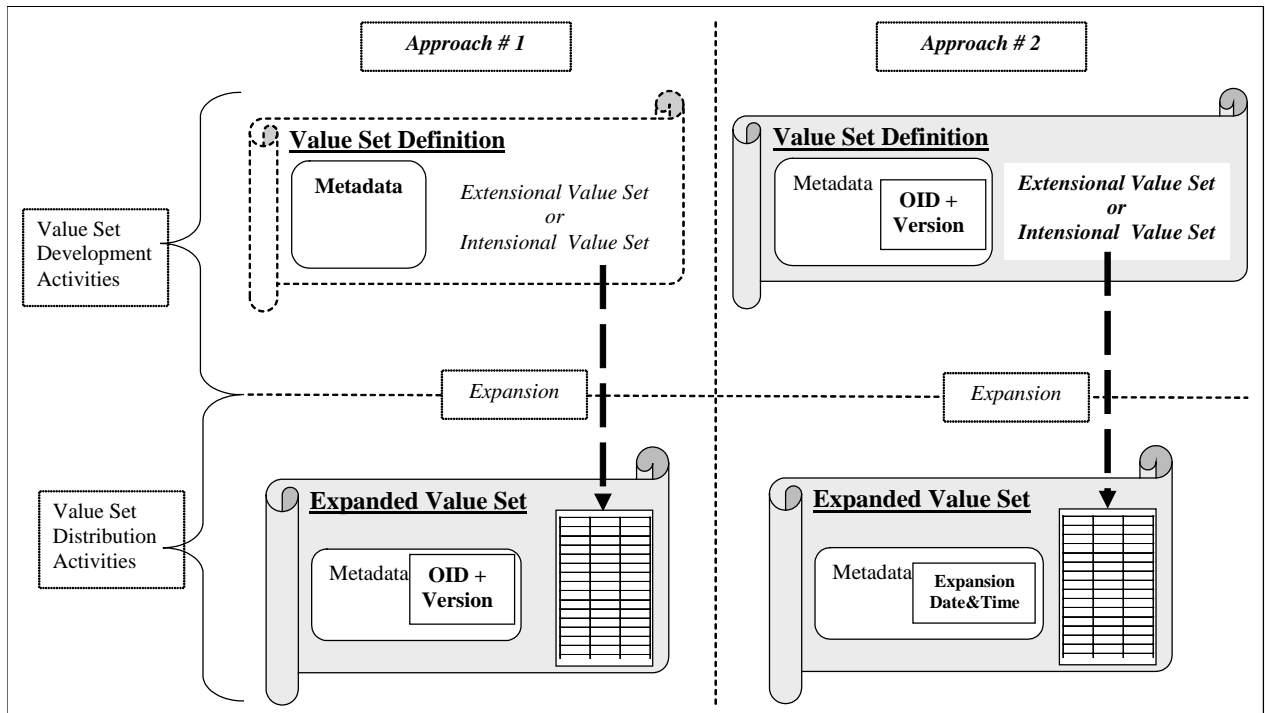


Figure 21.1.2-1: The two approaches for identifying Value Sets

21.1.3 The relationship between ITI SVS and CTS

5590 The Value Set Repository Actor can be supported by a system that implements a
 Terminology Server using the current HL7 CTS or the upcoming HL7 CTS2 specifications. It is
 important to note the complementary role of the HL7
 specification for CTS and CTS2, and that of the SVS Integration Profile. CTS defines an API
 (Application Programming Interface) supported by a terminology management service, and
 5595 CTS2 defines the functionality supported by a terminology management service leaving the
 specification of the API to the Object Management Group. SVS defines the transmission
 protocols for a network access to a terminology server focused specifically on the distribution of
 Value Sets.

5600 However there is functional consistency between SVS and CTS/CTS2. More
 specifically, all the properties of the Value Sets and concepts described in the
 Shared Value Sets Retrieve transaction are a subset of the properties defined in
 CTS and the CTS2 functional specification for the same entities. Note that SVS
 supports the distribution of Value Sets containing concepts from multiple
 code systems (e.g., DICOM and SNOMED issued) which is consistent with the CTS
 5605 capabilities, but which was not supported in the CTS specifications (but is supported
 in the CTS2 specification).

Informative references:

1. LexGrid Common Terminology Services.
<http://informatics.mayo.edu/LexGrid/downloads/CTS/specification/ctsspec/cts.htm>.
- 5610 2. Common Terminology Services 2 (CTS 2). Service Functional Model Specification. (See HL7 site for latest information.)

21.1.3.1 Value Set Distribution Flow

There are three types of value sets supported by the SVS Transactions:

- 5615 1. Intensional Value Sets are defined in terms of algorithmic and other methods. These value sets can be supported by the Value Set Repository, but this profile does not provide a means to convey the intensional form. Instead, these value sets are described using the metadata, and the appropriate resulting expanded value set contents are returned along with the Intensional Value Set definition and expansion metadata. This profile specifies how these can be retrieved using the Retrieve Multiple Value Sets Transaction (ITI-60).
- 5620 2. Extensional value sets are defined in terms of a list of concepts. As with intensional value sets, the definition and expansion metadata for these can be retrieved along with the appropriate expanded value set contents. This profile specifies how these can be retrieved using the Retrieve Multiple Value Sets Transaction (ITI-60).
- 5625 3. Expanded Value Sets result from the expansion of any Value Set definition (e.g., Intensional or Extensional), but their definition metadata is not important to the Value Set Consumer, only an identified instantiation defined in terms of a list of specific codes from specific vocabularies is shared. This profile describes how these can be retrieved using either the retrieve multiple value sets transaction (ITI-60), or the Value Set Retrieve Transaction (ITI-48).
- 5630 The developers of value sets may choose to work with one or more of these types, but the final consumers of value sets need to work with expanded value sets. There are efforts underway to develop standard methods for exchanging explicit definition of intensional and extensional value sets, but these are outside the scope of the SVS Profile. SVS provides only a way to distribute value sets that have been expanded.
- 5635 The SVS Profile also restricts the complexity of the expanded value sets. At present, it only supports unstructured value sets that are a list of codes from coded terminologies. Other internal structures such as hierarchy are not defined. This meets the needs of most, but not all, value sets.
- 5640 The process and rules associated with a value set expansion is not specified nor constrained by this profile. It is the responsibility of the value set developer or of the system supporting the SVS Repository Actor to perform the appropriate expansions. If the value set developer defines their standard distribution format as the expanded form of the value set, they have the appropriate procedures for this expansion. Value set developers that do not have a procedure defined for distributing the expanded
- 5645 form will need to establish one in order to use the SVS Profile.

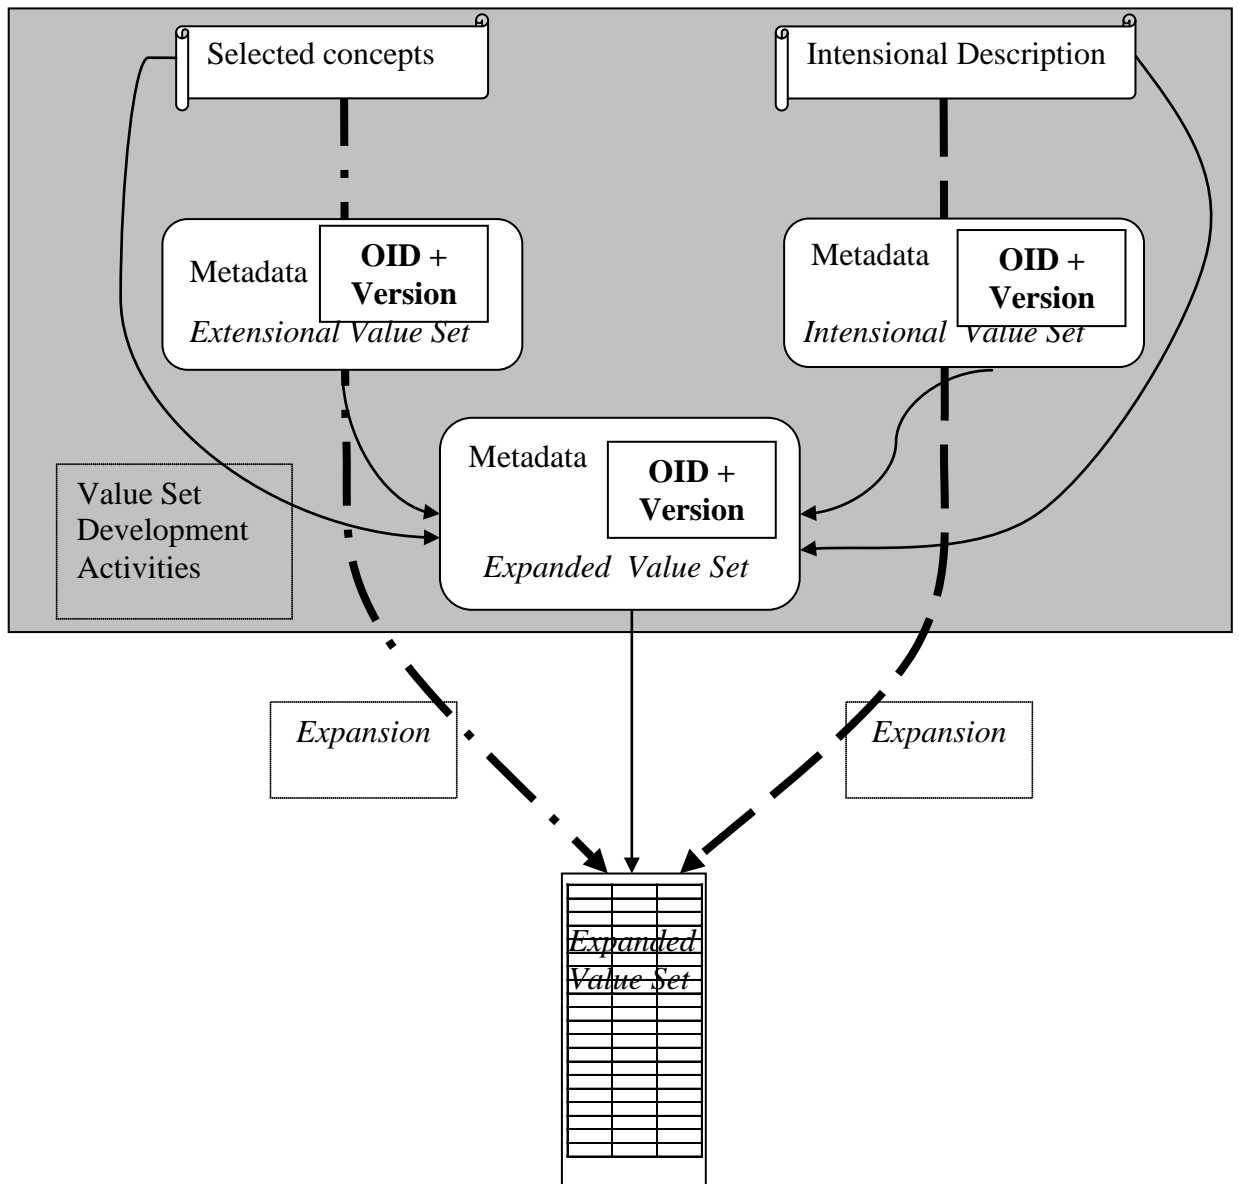


Figure 21.1.3.1-1: Development Flow for Value Sets

5650

A value set developer that defines and publishes expanded value sets should also establish the proper identification that identifies either this expanded value set or the definition that resulted in this expanded value set. They also define metadata that describes the value set. (Value set group descriptions will be discussed later.) The metadata is listed below, and includes descriptive information, links to further explanatory material, effective dates, etc. The SVS Profile provides two transactions for retrieving an expanded value set:

5655

1. Retrieve Value Set – This is appropriate for rapid retrieval of expanded value sets. It retrieves the expanded value set based on having the OID for the value set pre-configured into the system requesting the value set. This transaction does not retrieve the expanded

- 5660 value set metadata nor the value set definition metadata. It only retrieves the list of codes for that expanded value set.
2. Retrieve Multiple Value Sets – This is appropriate for retrieval of value sets based on metadata contents. It can still retrieve value set expansion based on the value set OID, but can also retrieve value set expansions based on contents of descriptions, OIDs and versions, group labels, dates, etc. This form of retrieval provides both the expanded value set contents for the retrieved value sets and the metadata for the value set.
- 5665 Value set developers that publish intensional and extensional value sets also defined OIDs for their value sets definitions. Note that a developer may publish multiple forms of related value sets, but will assign each form the proper OID. When publishing with SVS, the value set developer should provide an expanded form that should be provided along with the metadata.
- 5670 The SVS Profile provides one transaction for retrieving intensional and extensional value sets:
1. Retrieve Multiple Value Sets – This is appropriate for retrieval of value sets based on metadata contents, including value set OID, but can also be based on contents of descriptions, group labels, dates, etc. This form of retrieval provides both the expanded value set contents for the retrieved value sets and the metadata for the value set. Note that there are other standards efforts defining forms for intensional and extensional value sets. These other forms are intended for use by value set developers. SVS provides the expanded form primarily for value set consumers.
- 5675 A value set user that receives an intensional or extensional value set must be aware that the expansion is only for representational uses. The other metadata, such as effective dates and the descriptive material, must be consulted to determine the proper use of the expanded form. In practice, value sets change slowly and there is usually time for human review and decision making about the use of the expanded form.
- 5680 The SVS Profile does not specify how or when this expansion should take place. That is the responsibility of the value set developers and server maintainers. In many cases, the value set developer will provide an expanded form together with effective dates so that the organizations involved can manage change easily.
- 5690

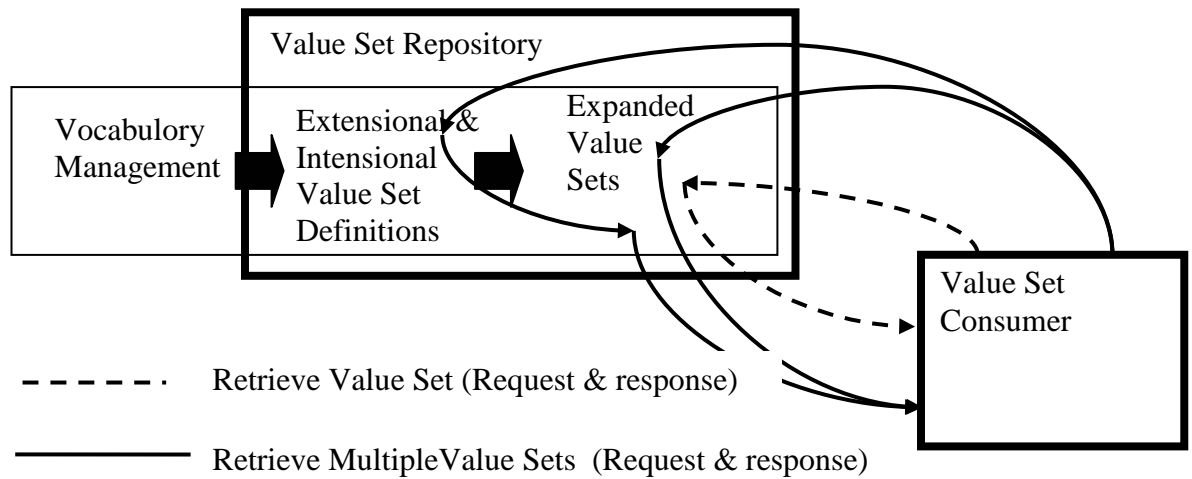


Figure 21.1.3.1-2: SVS Retrieve Transactions

21.1.3.2 Value Set Groups

5695 Value sets are also described by various grouping and tagging mechanisms. These groupings
 may be defined in parallel by many different organizations. It is expected that each organization
 is creating groups for their own purpose. One organization may assign groups like “value sets
 associated with H1N1”, while another group may assign groups like “value sets associated with
 5700 clinical trial xyz reports”, and a third may assign groups like “formulary for treatment of H1N1
 influenza”. Each of these organizations may assign key words so that retrieval requests can find
 the relevant groups, and they may assign OIDs for these groups.

To simplify maintenance, SVS defines a list of group descriptions to be associated with each
 value set, rather than combining all the keywords and groups from different organizations into a
 single list. The retrieval transaction searches all of these descriptions when doing a retrieval
 5705 based on group keyword or group OID.

An organization that is creating new groups can define a list of keywords and an OID for that
 group purpose. This group description can then be attached to each value set that should be a
 member of that group. If a value set needs to be removed from the group, then the attached
 description can be removed. This avoids accidental removal of keywords when multiple
 5710 organizations have used the same keyword

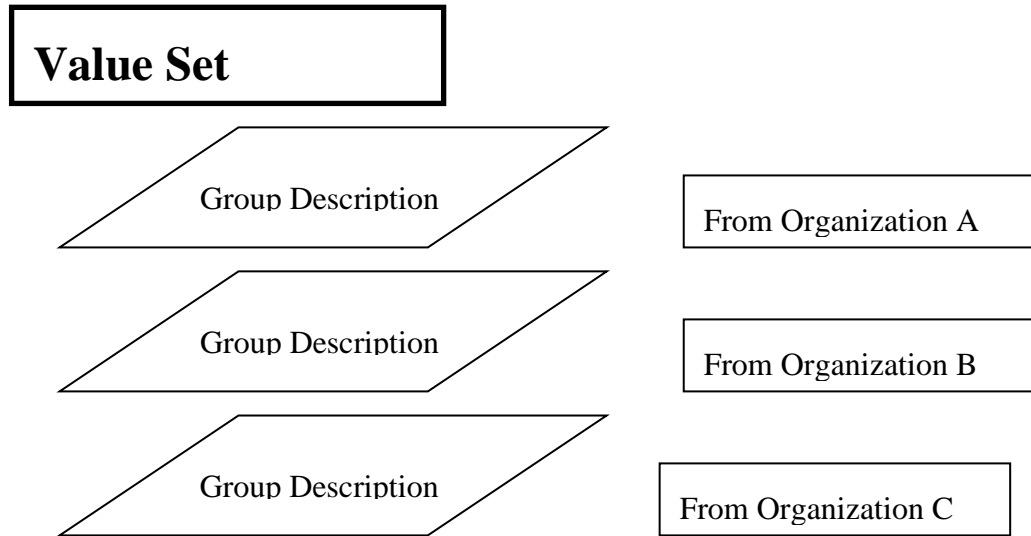


Figure 21.1.3.2-1: Group Descriptions for a Value Set

21.1.3.3 Value Set Descriptive Metadata

5715 A value set is described by metadata that includes the fields in Table 21.1.3.3.-1. For details on the metadata encoding, see ITI TF-2b: 3.60. Fields are mandatory or optional as shown in the table. Some of the metadata can be used as retrieval criteria for both the ITI-48 and ITI-60 transaction, some only for ITI-60, and some are only returned and cannot be used as retrieval criteria.

5720

Table 21.1.3.3-1: Value Set Metadata Summary

Metadata Element	Description	Optionality	Selection Criteria for Transactions
Id	This is the unique identifier of the value set	Mandatory	ITI-48, ITI-60
DisplayName	This is the name of the value set	Mandatory	ITI-60
Source	This is the source of the value set, identifying the originator or publisher of the information	Mandatory	ITI-60
Purpose	Brief description about the general purpose of the value set	Optional	ITI-60
Definition	A text definition describing how concepts in the value set were selected	Optional	ITI-60

Metadata Element	Description	Optionality	Selection Criteria for Transactions
Source URI	Most sources also have a URL or document URI that provides further details regarding the value set.	Optional	-
Version	A string identifying the specific version of the value set.	Mandatory	ITI-48
Status	Active, Inactive, local extensions	Mandatory	-
Type	This describes the type of the value set: <ul style="list-style-type: none"> • Intensional, • Extensional, or • Expanded Note: This is the type of the value set in the repository. Value set retrieval will return a value set expansion.	Mandatory	-
Binding	Static or Dynamic	Optional	-
Effective Date	The date when the value set is expected to be effective	Optional	ITI-60
Expiration Date	The date when the value set is no longer expected to be used	Optional	ITI-60
Creation Date	The date of creation of the value set	Optional	ITI-60
Revision Date	The date of revision of the value set	Optional	ITI-60
Groups	The identifiers and keywords of the groups that include this value set. A group may also have an OID assigned.	Optional	ITI-60

1. Status codes are determined by the Value Set developers. The suggested values shall be used if applicable.
2. The meaning of binding is not constrained by this Profile.

5725 Some of these metadata fields can be specified as part of the selection criteria for retrieve multiple value sets. All of the available metadata is returned as the results from a retrieve multiple value sets. Metadata is not returned for the ITI-48 transaction.

This profile does not specify how the value set repository is maintained, how new value sets are added, or how existing values sets are updated.

21.2 SVS Integration Profile Options

5730 Options that may be selected for this Integration Profile are listed in the **Table 21.2-1 Sharing Value Sets - Actors and Options** along with the actors to which they apply. Dependencies between options when applicable are specified in notes. Note that the Value Set Consumer shall implement at least one of the two bindings listed as options in the table. The Value Set Repository shall implement both bindings as specified in ITI TF-2b: 3.48.5.

5735

Table 21.2-1: Sharing Value Sets - Actors and Options

Actor	Options	Vol. & Section
Value Set Repository	No options defined	

Actor	Options	Vol. & Section
Value Set Consumer	HTTP binding (see note)	ITI TF-2b: 3.48.5, ITI TF-2b: 3.60.5.2
	SOAP binding (see note)	ITI TF-2b: 3.48.5, ITI TF-2b: 3.60.5.1
	Retrieve Multiple Value Sets	ITI TF-2b:3.60


Note: A Value Set Consumer must support either the HTTP binding, the SOAP binding or both bindings. The Value Set Repository must support both bindings.

5740 21.2.1 Retrieve Multiple Value Sets

Value Set Consumers that support the Retrieve Multiple Value Sets Option shall support the Retrieve Multiple Value Sets [ITI-60] transaction.

21.3 SVS Process Flow

5745 This section describes the process and information flow when a Value Set Consumer retrieves a Value Set from a Value Set Repository. There is no required order between the two transactions. The Value Set Consumer chooses whichever transactions and order are appropriate. The Value Set Consumer can use Retrieve Value Set ITI-48 to retrieve a single value set based upon a known value set OID. The Retrieve Multiple Value Sets ITI-60 can be used to retrieve all of the value sets that match a selection specification. The selection criteria for ITI-60 need not include a
5750 known value set OID.



5755

Figure 21.3-1: Basic Process Flow in SVS Profile

21.3.1 Overview of the entire process flow

5760 This profile describes functionality in the context of the larger system of anticipated actors involved in the creation and management of Value Sets.

The creation of a Value Set is out of scope of this profile. It will be addressed in a later cycle, once the basic infrastructure of this profile is in place. For definition purposes, creating a Value Set means the creation of a Value Set out of a Code System(s), or having the user proposing values that s/he uses in their own system.

5765 The complete process can be seen in Figure 21.3.1-1, Overview of process flows below, included for clarity's sake:

5770

5775

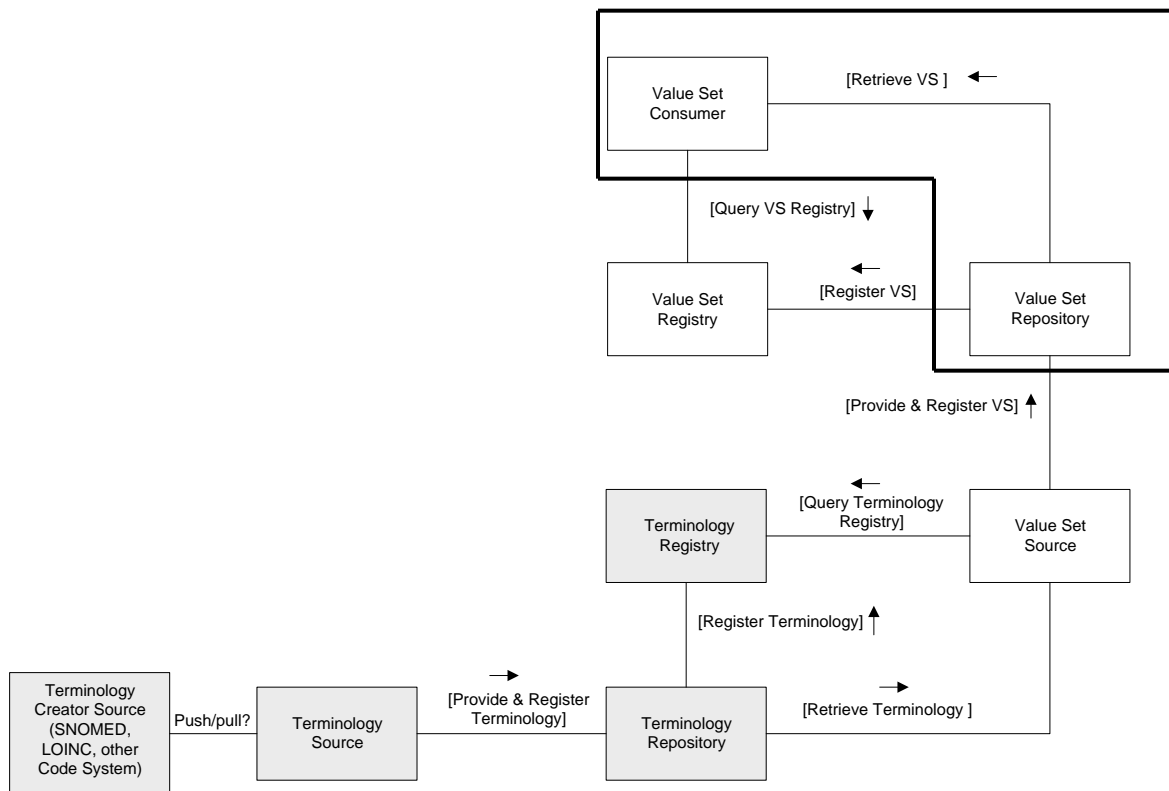


Figure 21.3.1-1: Overview of the process flow

5780

Figure 21.3.1-1 shows the Retrieve Value Set transaction in the context of the larger system of anticipated actors involved with the creation and management of Value Sets. This profile only addresses the actors and transactions outlined by the thick solid line.

5785

The SVS Profile addresses partly the semantic interoperability issue and assumes that a structure is already in place to provide the necessary context for the use of the Value Set.

While the representation of structure is out of scope of this profile, it must be recognized that it plays an important role in achieving semantic interoperability. The focus of the profile is to distribute a generalized and uniform nomenclature in order to populate the information model with the appropriate semantic content.

5790

21.3.2 Use Cases

The following use cases indicate how this profile might be used by various disciplines.

Note: All the tables present in the use cases are examples only. IHE will not be responsible for updating these tables.

5795 **21.3.2.1 Distributing a consistent nomenclature in an XDS Affinity Domain**

A common nomenclature is required in an XDS Affinity Domain for metadata elements such as classCode, confidentialityCode, eventCodeList, healthcareFacilityTypeCode, practiceSettingCode, and typeCode. More detailed information about a possible definition of an Affinity Domain can be found in the white paper **IHE IT Infrastructure Technical Committee White Paper - Template for XDS Affinity Domain Deployment Planning**, found on the **ihe.net**.

5800

21.3.2.1.1 Current state

The nomenclature used in the Affinity Domain is being entered into systems manually, a time consuming task, potentially leading to errors.

5805

21.3.2.1.2 Desired state

Each vendor's application would retrieve the necessary Value Sets used in a XDS Affinity Domain from a Value Set Repository, eliminating manual entry and improving accuracy.

21.3.2.2 Updating terminology codes for a medical and billing across systems

5810

Standardized coding systems are essential for health insurance programs to ensure that these claims are processed in an orderly and consistent manner.

The CPT is a uniform coding system consisting of descriptive terms and identifying codes that are used primarily to identify medical services and procedures performed by physicians and other health care professionals (HCP), for billing public or private health insurance programs.

21.3.2.2.1 Current state

5815

A patient is being referred by her PCP from a small healthcare facility to a large healthcare facility. She gets hospitalized and is being seen by a group of healthcare professionals: oncologists, laboratory practitioners, pharmacists, and nurses.

5820

The patient's record will contain medical information from different healthcare information edge systems, such as an Electronic Medical Record system (EMR), a Laboratory Information System (LIS), and a Radiology Information System (RIS).

All systems need up-to-date CPT codes so that seamless flow of encoded information results. Currently the update is achieved via application-specific processes on a system by system basis, which increases the risk of error when updating Value Sets in multiple systems.

5825

The Discharge Summary produced by the hospital lacks coded information about the care received due to the lack of a consistent and uniform nomenclature. The document is then published to a regional repository or saved on a portable media. The PCP can then retrieve it (via XDS or XDM, for example).

5830

Due to the full lack of encoding, two potentially undesirable outcomes can happen: either the correct billing information will not reach the provider, or the medical information is not machine processible and cannot be incorporated in other systems, with data mining being compromised.

21.3.2.2.2 Desired state

The hospital retrieves the significant CPT codes from the Value Set Repository so that all the applications are using the same nomenclature. This way, the medical and billing information will flow seamlessly, improving the quality of patient care.

5835 **21.3.2.3 Consistent Encoding Terms for anatomical regions in imaging**

21.3.2.3.1 Current state

5840 In hospital **A**, an imaging technologist is about to start a CT procedure. S/he chooses its protocol and estimates the body part s/he should be entering manually in the “*body part*” field present on the machine. The modality will over-ride the RIS information that the RIS administrator has configured for the CT exams, (or it might take the existing RIS information, depending on the vendor and on the implementation).

The study is sent to the healthcare facility **A** local PACS, and a manifest is sent to the XDS Repository **A**. Hospital **B** wishes to retrieve the study by querying the XDS Registry.

5845 Alternatively, the patient will bring the study performed in hospital **A** on a CD to be imported into the local system of hospital **B** via IRWF (Import Reconciliation Workflow).

The nomenclature used for “*body part*” in the RIS from hospital **A** is not consistent with the encoding chosen by the RIS in hospital **B**. The local PACS and RIS administrator need to place an order in the RIS, and manually reconcile the study so that it will have the same body part in order to ensure the same hanging protocols for the radiologists.

5850 **21.3.2.3.2 Desired state**

In hospital **A**, an imaging technologist is about to start a CT procedure. S/he chooses the correct “*body part*” from the latest Value Set Anatomical Regions downloaded from the Value Set Repository. The study is sent to the local PACS of healthcare facility **A**, and a manifest is sent to the XDS Repository **A**. Hospital **B** wishes to retrieve the study.

5855 Alternatively, the patient will bring the study performed in hospital **A** on a CD to be imported into the local system of hospital **B** via IRWF (Import Reconciliation Workflow). The nomenclature used for “*body part*” in the RIS from hospital **A** is consistent with the encoding chosen by the RIS in hospital **B** because hospital **B** has also downloaded the same Expanded Value Set from the Value Set Repository. The radiologist will see the images displayed according to the department’s hanging protocols.

5860 A set of flat list values that can be used for such purposes is DICOM Part 16, CID 4031 Common Anatomic Regions, of which an excerpt can be seen in Table 21.3.2.3.2-1: CID 4031 Common Anatomic Regions:

5865

Table 21.3.2.3.2-1: CID 4031 Common Anatomic Regions

Coding Scheme Designator (0008,0102)	Code Value (0008,0100)	Code Meaning (0008,0104)
SRT	T-D4000	Abdomen
SRT	R-FAB57	Abdomen and Pelvis
SRT	T-15420	Acromioclavicular joint
SRT	T-15750	Ankle joint
SRT	T-280A0	Apex of Lung
SRT	T-D8200	Arm
SRT	T-60610	Bile duct
SRT	T-74000	Bladder
SRT	T-04000	Breast
SRT	T-26000	Bronchus
SRT	T-12770	Calcaneus
SRT	T-11501	Cervical spine

Note: Excerpt from the Context ID 4031 Common Anatomic Regions, Type: Extensible Version 20061023, DICOM Part 16, OID 1.2.840.10008.6.1.308.

5870

21.3.2.4 Modification of a protocol code for a mammogram exam

Radiology departments or healthcare enterprises define local codes that are used in common by the systems in use, accordingly to the local policies and their workflow.

5875 According to the Mammography Acquisition Workflow Profile (MAWF) from the Radiology Technical Framework, codes are used for:

- scheduling and driving modality behavior (Requested Procedure, Reason for Requested Procedure and Scheduled Protocols)
- documenting the images and the workflow status: codes for Performed Procedure, Performed Protocols, Views, etc. enable displays to present images in adequate hanging protocols
- enabling radiological staff to track performed work or chose the right billing code.

5880

The MAWF Profile further states that a department or enterprise should define the code sets which are used by all of its systems in a common way, so that each relevant code set is available to each system with the same valid content. Each system needs to be configurable as to which code sets it uses. The lack of a common mechanism for distribution of code sets contributes to the development of local protocols like “*routine screening*”, “*magnification*”, “*CAD*”, that are understood by technologists or doctors, but could not be applied to another department or enterprise, nor by the modality in the scope of an automated error correction.

5885

5890 Moreover, those codes are subject to be modified, removed, declared obsolete, or simply dropped. This situation is confusing since the RIS list of protocol codes cannot be fully reliable anymore.

5895 Despite technical means defined in the *Scheduled Workflow* and *Mammography Acquisition Workflow Profiles*, variances in the way users and systems behave can lead to department inefficiencies, ambiguous data, special cases for automated billing, and less than optimal acquisition and reading environments.

21.3.2.4.1 Current state

5900 A patient comes in for a scheduled standard screening mammogram. While the acquisition is processed, a suspicious lump is detected, and additional views are required, taken by the technologist. A diagnostic mammogram is performed instead of the simple routine screening that was scheduled. This information must be then communicated to the RIS, in order to change the billing codes and implicitly change the hanging protocol for the radiologist. As it is, the technologist has to manually change the procedure.

The procedure code will have to be corrected in the RIS post-examination so that the correct information is captured, both for display and for billing purposes.

5905 **21.3.2.4.2 Desired state**

Changing a procedure code should be done directly from the modality, avoiding a subsequent intervention that can generate errors, misunderstandings, or discrepancies. SVS Profile provides the modality with a mechanism for accessing a uniformed, centralized and dedicated Expanded Value Set.

5910 An Expanded Value Set dedicated to mammography procedure codes is made available through the Value Set Repository.

The modality, acting as a Value Set Consumer, retrieves the Expanded Value Set commonly used by and defined for the mammography exams.

5915 The correct type of the exam is processed (or at least provides the technologist the ability to choose the right item from this list).

The list proposed is a flat list, and is pending approval in the DICOM standard.

Table 21.3.2.4.2-1: Codes for Mammography Procedures

Coding Scheme Designator (0008,0102)	Code Value (0008,0100)	Code Meaning (0008,0104)
IHERADTF	MAWF0001	Screening Mammography, bilateral
IHERADTF	MAWF0002	Screening Mammography, left
IHERADTF	MAWF0003	Screening Mammography, right
IHERADTF	MAWF0004	Diagnostic Mammography, bilateral
IHERADTF	MAWF0005	Diagnostic Mammography, left
IHERADTF	MAWF0006	Diagnostic Mammography, right
IHERADTF	MAWF0007	Mammary Ductogram, Single Duct, left
IHERADTF	MAWF0008	Mammary Ductogram, Single Duct, right
IHERADTF	MAWF0009	Mammary Ductogram, Multiple Ducts, left

Coding Scheme Designator (0008,0102)	Code Value (0008,0100)	Code Meaning (0008,0104)
IHERADTF	MAWF0010	Mammary Ductogram, Multiple Ducts, right
IHERADTF	MAWF0011	Mammogram for marker placement, left
IHERADTF	MAWF0012	Mammogram for marker placement, right
IHERADTF	MAWF0013	Needle Localization, Image Guided, Mammography, left
IHERADTF	MAWF0014	Needle Localization, Image Guided, Mammography, right
IHERADTF	MAWF0015	Stereotactic Biopsy, Image Guidance, left
IHERADTF	MAWF0016	Stereotactic Biopsy, Image Guidance, right
IHERADTF	MAWF0017	Breast Specimen Mammography, left
IHERADTF	MAWF0018	Breast Specimen Mammography, right
IHERADTF	MAWF0019	Quality Control, Mammography
IHERADTF	MAWF0020	Additional Mammography Views

5920

Note: These are provisional values, used as an example, whose inclusion in the DICOM Standard is currently requested (see RAD TF-1: B.2.ZA). IHE ITI is not responsible for updating these tables.

Table 21.3.2.4.2-2: Codes for Reasons for a Requested Procedure

Coding Scheme Designator (0008,0102)	Code Value (0008,0100)	Code Meaning (0008,0104)
Procedure type		
IHERADTF	MAWF0030	Recall for technical reasons
IHERADTF	MAWF0031	Recall for imaging findings
IHERADTF	MAWF0032	Recall for patient symptoms/ clinical findings
DCM	111416	Follow-up at short interval from prior study
SRT	R-42453	Screening (Note 1)
SRT	R-408C3	Diagnostic (Note 1)
SRT	A-04010	Implant (Note 1)

Note 1: These code values originate from DICOM CID 6061 (DICOM PS 3.16).

5925

Note: These are provisional values, whose inclusion in the DICOM Standard is currently requested (see RAD TF-1: B.2.ZA)

21.3.2.5 Distributing Value Sets from SDOs and other master sources

5930

There is a bidirectional relationship between the users of terminologies, codes, and value sets at one end, and the standards development organizations (SDOs) and other developers of terminologies, codes, and value sets. The following diagram shows the process by which terminologies and value sets flow up to the value set consumers. The users' comments and new requirements flow back down to the sources of information.

At the top of this diagram, the value set consumers retrieve values sets from a master value set repository that they need for particular purposes. This could be done with the ITI-48 transaction

5935 when the consumer is configured with specific OID values for specific purposes. Often, there is a need to retrieve a group of value sets that share a common purpose, such as all of the value sets needed to populate a particular kind of report. These retrievals are performed using the ITI-60 retrieve multiple value sets transaction.

5940 This master value set repository is subject to review and governance. The individual consumers have delegated responsibility for administering and maintaining the master value set repository to a coordinating organization. These organizations may be local, state, regional, or national organizations. They are typically not the developers of standard terminologies. The master repository organization serves an administrative and coordinating purpose to ensure that the releases of standard terminologies from SDOs do not interfere with daily operations of the value set consumers. They may also coordinate requests from value set consumers for new terminologies and value sets. There is a governance committee to coordinate these activities in both directions. These activities are important to the maintenance of the master value set repository. They are not described further as part of this profile.

5950 The terminology developers typically release new terminologies and value sets on a regular schedule or at times matching their process. These notifications may be via bulletins, electronic notification, and other processes. They are not covered as part of this profile. The governance committee may choose to use ITI-60 as their method of retrieving copies of the SDO value sets, if the SDO has established a value set repository as part of their distribution process.

Value Set Consumers

1. Search Vocabulary based on:
 - a) VS
 - b) Domain/ Group of VS
 - c) Resolved Group VS
 - d) Time Frame
 - e) Terminology
 - f) Vocabulary Version
2. Download results of search

Master VS Repository

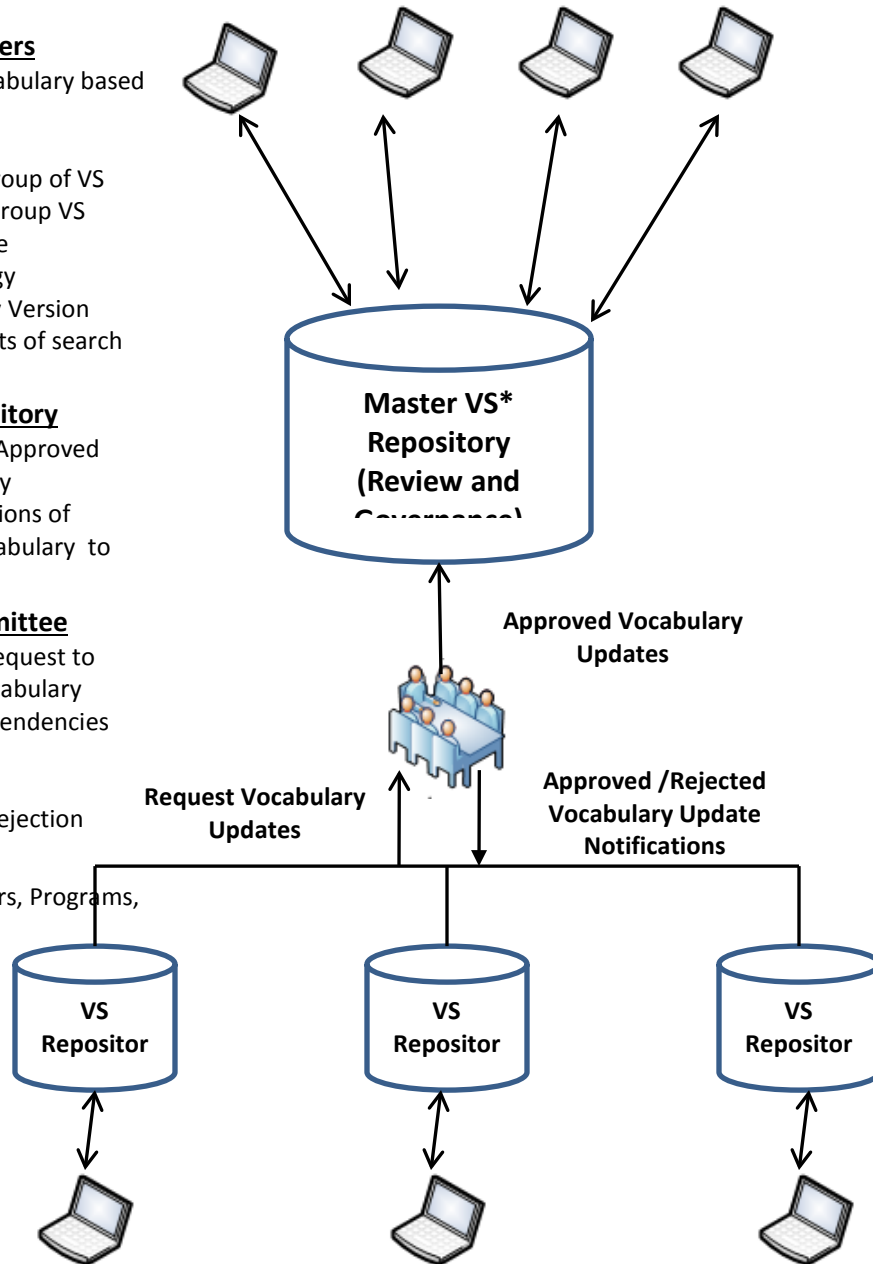
1. Received Approved Vocabulary
2. Send Notifications of Approved Vocabulary to Customers

Governance Committee

1. Received Request to Update Vocabulary
2. Review dependencies and impact
3. Provide Approval/Rejection

VS Source

Could be Consumers, Programs, SDO etc.



* Value Set

5955

Figure 21.3.2.5-1: Relationship between Users and Developers of Value Sets

5960 **21.3.2.6 Obtaining value sets based upon metadata**

There are often situations where notifications such as emails, bulletins, etc. contain descriptive information rather than a specific OID. Also, there are situations where potentially useful value sets must be found based upon only a description. An example of this kind of use is:

- 5965
1. A user needs all the value sets for stroke quality care measures from the US Joint Commission. These measures are identified by having a group name containing “stroke”. They plan to use this as the starting point for establishing triggers for decision support and data analytics application operating on data generated for the current year.
- 5970
1. The user interacts with a Value Set Consumer to request value sets that have a group that includes “stroke”, a source that includes “Joint Commission” or “JCAHO”, and that are effective for the current year.
 2. The Value Set Repository finds all the matching value sets and sends a response containing all the value sets and their descriptive metadata. Because there is also a European Joint Commission, this response includes some extras.
- 5975
3. Client uses the complete metadata to eliminate the extras that are not relevant to the purpose.

21.4 SVS Security Considerations

5980 The contents handled by the SVS Profile are not patient specific, so there are no risks to privacy. Some Expanded Value Sets are of little value to an attacker as they are public tables of non-critical information (e.g., Expanded Value Sets used for coding of body parts in medical exams). Other Expanded Value Sets might need protection against malicious modification or interception.

5985 The risks applicable to the SVS Profile are discussed in the table “*Risks associated with the profile SVS*” which is found on the IHE ftp site in *ftp://ftp.ihe.net/IT_Infrastructure/iheityr6-2008-2009/Technical_Cmte/Profile_Work/SharingValueSets/*. The nature of the Expanded Value Set exchange determines the type or risk that can incur. For example there can be integrity risks such as masquerade⁸, or the modification of Expanded Value Sets. Another possible type of risk would be at the privacy and confidentiality level such as the interception of an Expanded Value Set containing confidential data. The profile will allow mitigation of those risks when needed in the following manner:

- 5990
- A Value Sets Repository shall be grouped with an ATNA Secure Node or Secure Application. Since the Value Set Consumer is not required to be grouped with the Secure Node or Secure Application, the Value Set Repository shall support both secure and non-secure connections.
- 5995
- Value Set Repositories shall be able to restrict access to a specific Expanded Value Set to authorized and authenticated nodes, while allowing unauthenticated network queries to other Expanded Value Sets.

⁸A malicious server passing for the value set repository gives forged value sets.

6000

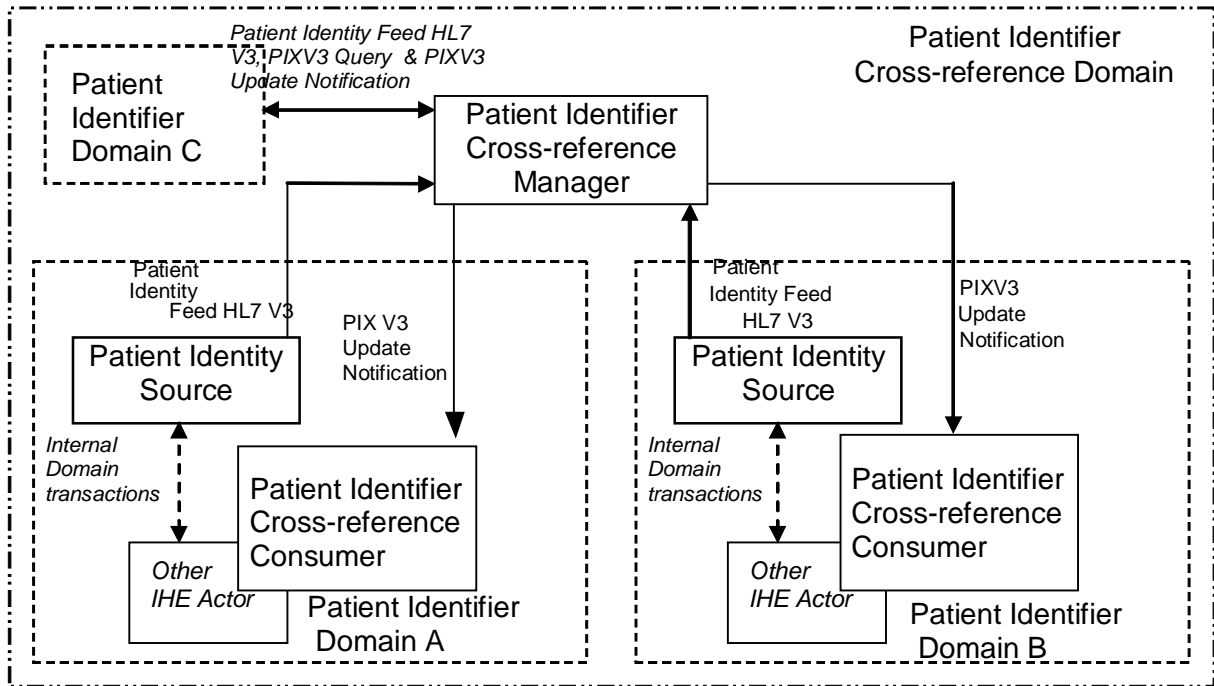
- Given the wide variety of systems that will be retrieving Expanded Value Sets (e.g., embedded medical device versus PACS) the profile does not mandate that the Value Set Consumer be grouped with an ATNA Secure Node or a Secure Application. Depending on local risk assessment, local policy may mandate such grouping.

22 Document-based Referral Request Integration Profile (DRR)

This profile has been retired in favor of use of the Cross-Enterprise Document Workflow (XDW) Profile.

6005 **23 Patient Identifier Cross-referencing HL7 V3 (PIXV3)**

The *Patient Identifier Cross-referencing HL7 V3 Integration Profile (PIXV3)* is targeted at cross-enterprise Patient Identifier Cross-reference Domains (as defined in ITI TF-1: 5) as well as healthcare enterprises with developed IT infrastructure. The discussion in ITI TF-1: 5 fully applies here, with the obvious adjustments to the referenced transactions.



6010

Figure 23-1: Process Flow with Patient Identifier Cross-referencing HL7 V3

23.1 Actors/Transactions

6015 The actors in this profile are the same as the actors defined in the PIX Profile (ITI TF-1: 5.1). Figure 23.1-1 shows the actors directly involved in the Patient Identifier Cross-referencing HL7 V3 Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other related profiles are not shown.

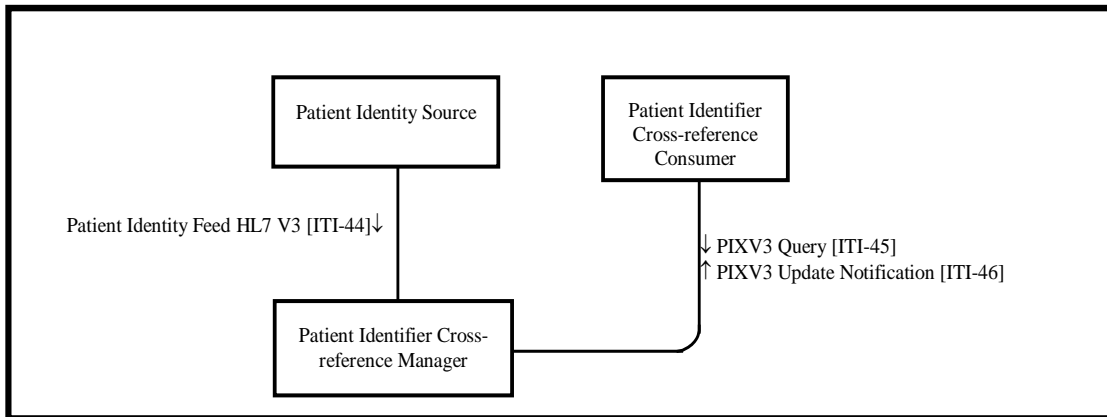


Figure 23.1-1: Patient Identifier Cross-referencing HL7 V3 Actor Diagram

6020 Table 23.1-1 lists the transactions for each actor directly involved in the Patient Identifier Cross-referencing Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in the ITI TF-1: 23.2.

6025

Table 23.1-1: Patient Identifier Cross-referencing HL7 V3 Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Patient Identity Source	Patient Identity Feed HL7 V3 [ITI-44]	R	ITI TF-2b: 3.44
Patient Identifier Cross-reference Consumer	PIXV3 Query [ITI-45]	R	ITI TF-2b: 3.45
	PIXV3 Update Notification [ITI-46]	O	ITI TF-2b: 3.46
Patient Identifier Cross-reference Manager	Patient Identity Feed HL7 V3 [ITI-44]	R	ITI TF-2b: 3.44
	PIXV3 Query [ITI-45]	R	ITI TF-2b: 3.45
	PIXV3 Update Notification [ITI-46]	R	ITI TF-2b: 3.46

The transactions in this profile directly correspond to the transactions used in the PIX Profile (ITI TF-1: 5) and provide the identical functionality. Table 23.1-2 describes this correspondence.

6030

Table 23.1-2: Transactions Correspondence between the PIX and PIXV3 profiles

Transactions in PIX	Section in Volume	Transactions in PIXV3	Section
Patient Identity Feed [ITI-8]	ITI TF-2a: 3.8	Patient Identity Feed HL7 V3 [ITI-44]	ITI TF-2b: 3.44
PIX Query [ITI-9]	ITI TF-2a: 3.9	PIXV3 Query [ITI-45]	ITI TF-2b: 3.45
PIX Update Notification [ITI-10]	ITI TF-2a: 3.10	PIXV3 Update Notification [ITI-46]	ITI TF-2b: 3.46

23.2 Patient Identifier Cross-referencing HL7 V3 Integration Profile Options

6035 Options that may be selected for this Integration Profile are listed in the Table 23.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 23.2-1: Patient Identifier Cross-referencing HL7 V3 - Actors and Options

Actor	Options	Vol. & Section
Patient Identity Source	Pediatric Demographics	ITI TF-1:23.2.1
Patient Identifier Cross-reference Manager	Pediatric Demographics	ITI TF-1:23.2.1
Patient Identifier Cross-reference Consumer	PIXV3 Update Notification Transaction	ITI TF-2b: 3.46

6040

23.2.1 Pediatric Demographics

6045 The experience of immunization registries and other public health population databases has shown that matching and linking patient records from different sources for the same individual person in environments with large proportions of pediatric records requires additional demographic data.

6050 In particular, distinguishing records for children who are twins, triplets, etc. – that is, avoiding false positive matches - may be difficult because much of the demographic data for the two individuals matches. For instance, twin children may have identical last names, parents, addresses, and dates of birth; their first names may be very similar, possibly differing by only one letter. It can be very difficult for a computer or even a human being to determine in this situation whether the slight first name difference points to two distinct individuals or just a typographical error in one of the records. Additional information is extremely helpful in making this determination.

6055 Pediatric Demographics makes use of the following six additional demographic fields to aid record matching in databases with many pediatric records.

Field	Reason for inclusion	Value
Mother's Maiden Name	Any information about the mother is helpful in making a match	Helps create true positive matches
Patient Home Telephone	A telecom helps match into the right household	Helps create true positive matches
Patient Multiple Birth Indicator	Indicates this person is a multiple – twin, triplet, etc.	Helps avoid false positive matches of multiples
Patient Birth Order	Distinguishes among those multiples.	Helps avoid false positive matches of multiples

Field	Reason for inclusion	Value
Last Update Date/Time, Last Update Facility	These fields, although not strictly demographic, can effectively substitute when multiple birth indicator and birth order are not collected. They indirectly provide visit information. Provider visits on the same day may likely indicate two children brought to a doctor together.	Helps avoid false positive matches of multiples

6060 Patient Identity Source actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction and shall provide values, when available, for the fields identified as Pediatric Demographics fields.

6065 Patient Identifier Cross-reference Manager actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction, and if values for one or more of the Pediatric Demographics fields are specified in the Patient Identity Management [ITI-30], they shall be considered as part of the matching algorithm of the PIX Manager.

Pediatric Demographics are defined as all of the following:

- Mother's Maiden Name
- Patient Home Telephone
- Patient Multiple Birth Indicator
- 6070 • Patient Birth Order
- Last Update Date/Time
- Last Update Facility

23.3 Patient Identifier Cross-referencing HL7 V3 Integration Profile Process Flows

6075 Sections ITI TF-1: 5.3.1 and ITI TF-1: 5.3.2 describe use cases that this profile addresses. Figures 5.3-1 and 5.3-2 also apply with the changes to the corresponding PIXV3 transactions as specified in Table 23.1-2.

23.4 Relationship between the PIXV3 Integration Profile and eMPI

The discussion in ITI TF-1: 5.4 fully applies to this profile.

6080 23.5 Patient Identifier Communication Requirement

The patient identifier in HL7 V3 messages is represented by the II data type. This data type has two components: a root, and an extension. For compatibility with the use of patient identifiers in profiles using HL7 V2 messages, and with the specification of the patient identifier in the XDS Profile, the patient identifier SHALL be represented as a root and an extension, where the root is

6085 an appropriately assigned OID. The direct correspondence between the II data type and the HL7 Version 2.5 CX data type (used in field PID-3) is shown in ITI TF-2x: Appendix R.

23.6 Security Considerations

6090 The implementer of this profile is advised that many risks cannot be mitigated by the IHE profile and instead the responsibility for mitigation is transferred to the vendor, and occasionally to the operational environment.

In order to address identified security risks:

- All actors in PIXV3 should be grouped with a Consistent Time (CT) Profile - Time Client Actor. This grouping will assure that all systems have a consistent time clock to assure a consistent timestamp for audit logging.
- 6095 • All actors in PIXV3 should be grouped with an Audit Trail and Node Authentication (ATNA) Profile - Secure Node Actor or ATNA Secure Application Actor. This grouping will assure that only highly trusted systems can communicate and that all changes are recorded in the audit log.
- 6100 • All actors in PIXV3 should be grouped with an XUA X-Service User or X-Service Provider Actor as appropriate. This grouping will enable service side access control and more detailed audit logging.
- All actors in PIXV3 should be grouped with the appropriate actor from the Enterprise User Authentication (EUA) Profile to enable single sign-on inside an enterprise by facilitating one name per user for participating devices and software.

6105 **24 Patient Demographics Query HL7 V3 (PDQV3)**

The *Patient Demographics Query HL7 V3 Integration Profile (PDQV3)* provides ways for multiple distributed applications to query a patient information server for a list of patients, based on user-defined search criteria, and retrieve a patient's demographic information directly into the application. The discussion and use cases in ITI TF-1: 8 fully apply here, with the obvious adjustments to the referenced transactions.

6110

24.1 Actors/Transactions

The actors in this profile are the same as the actors defined in the PDQ Profile (ITI TF-1: 8.1).

Table 24.1-1: Patient Demographics Query HL7 V3 Integration Profile - Actors and Transactions

6115

Actors	Transactions	Optionality	Section
Patient Demographics Consumer	Patient Demographics Query HL7 V3	R	ITI TF-2b: 3.47
Patient Demographics Supplier	Patient Demographics Query HL7 V3	R	ITI TF-2b: 3.47

The transaction in this profile directly corresponds to one of the transactions used in the PDQ Profile (ITI TF-1: 8) and provide the identical functionality. Table 24.1-2 describes this correspondence. Note that unlike the PDQ Profile there is no transaction which corresponds to the Patient Demographics and Visit query (ITI-22).

6120

Table 24.1-2: Transactions Correspondence between the PDQ and PDQV3 profiles

Transactions in PDQ	Section in Volume	Transactions in PDQV3	Section in Volume
Patient Demographics Query [ITI-21]	ITI TF-2: 3.21	Patient Demographics Query HL7 V3 [ITI-47]	ITI TF-2b: 3.47

24.2 Patient Demographics Query HL7 V3 Integration Profile Options

Options that may be selected for this Integration Profile are listed in the Table 24.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

6125

Table 24.2-1: Patient Demographics Query HL7 V3 - Actors and Options

Actor	Options	Vol. & Section
Patient Demographics Consumer	Continuation Option Pediatric Demographics	ITI TF-1: 24.2.2

Actor	Options	Vol. & Section
Patient Demographics Supplier	Continuation Option Pediatric Demographics	ITI TF-1: 24.2.2

6130

Support of continuations is described in transaction ITI-47. This option allows the Patient Demographics Consumer to get the full set of responses in several increments, as opposed to in one single response.

24.2.1 Continuation

6135

Support of continuations is described in transaction ITI-47. This option allows the Patient Demographics Consumer to get the full set of responses in several increments, as opposed to in one single response.

24.2.2 Pediatric Demographics Option

6140

The experience of immunization registries and other public health population databases has shown that matching and linking patient records from different sources for the same individual person in environments with large proportions of pediatric records requires additional demographic data.

6145

In particular, distinguishing records for children who are twins, triplets, etc. – that is, avoiding false positive matches - may be difficult because much of the demographic data for the two individuals matches. For instance, twin children may have identical last names, parents, addresses, and dates of birth; their first names may be very similar, possibly differing by only one letter. It can be very difficult for a computer or even a human being to determine in this situation whether the slight first name difference points to two distinct individuals or just a typographical error in one of the records. Additional information is extremely helpful in making this determination.

6150

Pediatric Demographics makes use of the following six additional demographic fields to aid record matching in databases with many pediatric records.

Field	Reason for inclusion	Value
Mother's Maiden Name	Any information about the mother is helpful in making a match	Helps create true positive matches
Patient Home Telephone	A telecom helps match into the right household	Helps create true positive matches
Patient Multiple Birth Indicator	Indicates this person is a multiple – twin, triplet, etc.	Helps avoid false positive matches of multiples
Patient Birth Order	Distinguishes among those multiples.	Helps avoid false positive matches of multiples

Field	Reason for inclusion	Value
Last Update Date/Time, Last Update Facility	These fields, although not strictly demographic, can effectively substitute when multiple birth indicator and birth order are not collected. They indirectly provide visit information. Provider visits on the same day may likely indicate two children brought to a doctor together.	Helps avoid false positive matches of multiples

6155 Patient Identity Source actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction and shall provide values, when available, for the fields identified as Pediatric Demographics fields.

6160 Patient Identifier Cross-reference Manager actors which support the Pediatric Demographics Option are required to support the Patient Identity Management [ITI-30] transaction, and if values for one or more of the Pediatric Demographics fields are specified in the Patient Identity Management [ITI-30], they shall be considered as part of the matching algorithm of the PIX Manager.

Pediatric Demographics are defined as all of the following:

- Mother’s Maiden Name
- 6165 • Patient Home Telephone
- Patient Multiple Birth Indicator
- Patient Birth Order
- Last Update Date/Time
- Last Update Facility

6170 **24.3 Patient Demographics Query HL7 V3 Process Flow**

ITI TF-1: 8.3 describes use cases that this profile addresses. Figure 8.3-1 also applies to this profile with the changes to the corresponding PDQV3 transactions as specified in Table 24.1-2, and omitting transaction ITI-22, which has no correspondence in this profile.

24.3.1 Combined Use of PDQV3 with other IHE Workflow Profiles

6175 In addition to the discussion in ITI TF-1: 8.3.1, the use of web services as the transport in the transactions in this profile makes it well suited in cases where other web services-based profiles are used, like XDS.b and PIXV3.

24.3.2 Supplier Data Configuration

6180 The Patient Demographics Supplier provides demographics information about possible matches to the parameters of the query. As described in ITI TF-2x: Appendix M, while it is possible for

the supplier to have demographics information from multiple domains, only a single set of demographics shall be returned by the supplier.

6185 If the supplier holds information for a single Patient ID domain, it shall provide the demographics information from that domain. In the case where the supplier holds demographics information from multiple Patient ID domains, the determination of which set of information to return must be based on the ID values for the Receiver's Device and Organization classes of the query transmission wrapper (the equivalent of MSH-5 and MSH-6 in the HL7 Version 2.5 corresponding message).

24.4 Intentionally Left Blank

6190 24.5 Security Considerations

The implementer of this profile is advised that many risks cannot be mitigated by the IHE profile and instead the responsibility for mitigation is transferred to the vendor, and occasionally to the operational environment.

In order to address identified security risks:

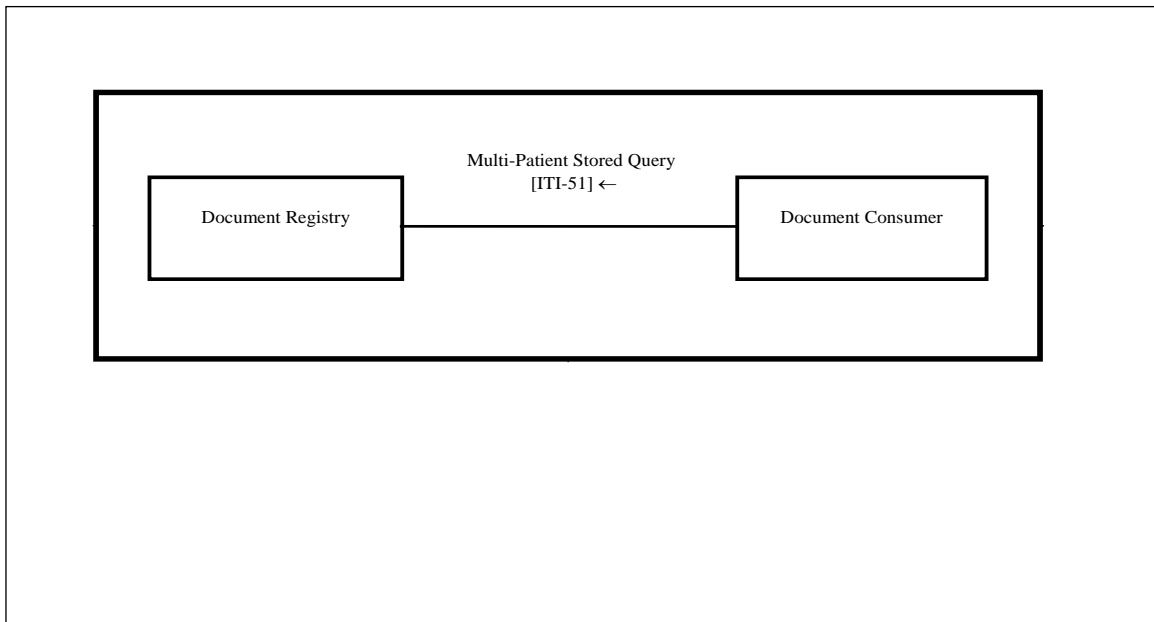
- 6195 • All actors in PDQV3 should be grouped with a Consistent Time (CT) Profile - Time Client Actor. This grouping will assure that all systems have a consistent time clock to assure a consistent timestamp for audit logging.
- 6200 • All actors in PDQV3 should be grouped with an Audit Trail and Node Authentication (ATNA) Profile - Secure Node Actor or ATNA Secure Application Actor. This grouping will assure that only highly trusted systems can communicate and that all changes are recorded in the audit log.
- All actors in PDQV3 should be grouped with an XUA X-Service User or X-Service Provider Actor as appropriate. This grouping will enable service side access control and more detailed audit logging.
- 6205 • All actors in PDQV3 should be grouped with the appropriate actor from the Enterprise User Authentication (EUA) Profile to enable single sign-on inside an enterprise by facilitating one name per user for participating devices and software.

25 Multi-Patient Queries (MPQ)

6210 The Multi-Patient Queries Profile defines a mechanism to enable aggregated queries to a Document Registry based on certain criteria needed by areas related to data analysis, such as quality accreditation of health care practitioners or health care facilities, clinical research trial data collection or population health monitoring.

25.1 Actors/ Transactions

6215 Figure 25.1-1 shows the actors directly involved in the MPQ Integration Profile in a solely XDS Affinity Domain and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in other related profiles, etc. are not necessarily shown.



6230 **Figure 25.1-1: Multi-Patient Queries Actor Diagram**

Table 25.1-1 lists the transactions for each actor directly involved in the Multi-Patient Query Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in ITI TF-1: 25.2.

6235

Table 25.1-1: Multi-Patient Queries Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section in Vol. 2
Document Registry	Multi-Patient Stored Query [ITI-51]	R	ITI TF-2b: 3.51
Document Consumer	Multi-Patient Stored Query [ITI-51]	R	ITI TF-2b: 3.51

25.2 Multi-Patient Query Integration Profile Options

Options that may be selected for this Integration Profile are listed in Table 25.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 25.2-1: MPQ - Actors and Options

Actor	Options	Vol. & Section
Document Registry	Asynchronous Web Services Exchange	ITI TF-1: 25.2.2
Document Consumer	Asynchronous Web Services Exchange	ITI TF-1: 25.2.2

25.2.2 Asynchronous Web Services Exchange Option

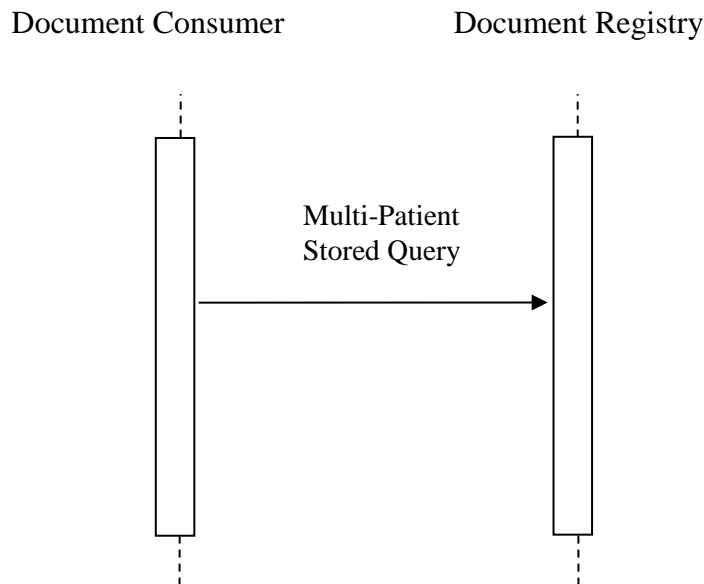
Actors that support this option shall support the following:

1. Document Consumer Actor shall support Asynchronous Web Services Exchange for the Multi-Patient Stored Query [ITI-51] transaction
2. Document Registry Actor shall support Asynchronous Web Services Exchange for the Multi-Patient Stored Query [ITI-51] transaction

Use of Synchronous or Asynchronous Web Services Exchange is dictated by the individual install environment and policies. Refer to section ITI TF-2x: V.5 Synchronous and Asynchronous Web Services Exchange for an explanation of Asynchronous Web Services Exchange.

25.3 MPQ Process Flow

This section describes the process and information flow when a Document Consumer will query a Document Registry.



6260

Figure 25.3-1: Basic Process Flow in Multi-Patient Queries Profile

25.4 Use Cases

25.4.1 Multi-Patient Query used in Public Health

Current Situation

6265 The emergency department at Hospital A is treating patient B for certain symptoms, which are indicative of a reportable condition (such as A1H1), according to already established guidelines from an official public health agency. The symptoms mandate the use of a pre-determined value set for the XDS metadata *eventCodeList*. This can be a combination of the eventCodeList and observation such as “*influenza*” and “*possible A1H1*”. Hospital A sends any type of document capturing this information such as a Discharge Summary, an ED Encounter Summary (EDES),
6270 or in a larger sense any document intended for this purpose, using an XDS.b Provide and Register transaction to the local XDS repository, as well as a report to the appropriate public health agency P, using mechanisms which are outside the scope of this supplement.

6275 After reviewing the report, the public health agency P determines that a review of recent patients’ encounters with similar symptoms is necessary. Unfortunately, the XDS Document Registry only accepts patient specific queries, as currently defined in the Stored Query

transaction. The public health agency P needs to obtain a list of patients with the appropriate symptoms from the healthcare providers.

6280 Hospital A queries the local Document Registry for other Document Entries containing the same event code. Since it is not possible to query for multiple patients in one operation, a query is initiated for each patient known to the Document Registry. This is very time consuming and may not be very accurate.

Desirable Situation

6285 The emergency department at Hospital A is treating patient B for certain symptoms, which are indicative of a reportable condition (such as A1H1), according to already established guidelines from an official public health agency. The symptoms mandate the use of a pre-determined value set for the XDS metadata eventCodeList. This can be a combination of the eventCodeList and observation such as “influenza” and “possible A1H1”. Hospital A sends any type of document capturing this information such as a Discharge Summary, an ED Encounter Summary (EDES),
6290 or in a larger sense any document intended for this purpose, using an XDS.b Provide and Register transaction to the local XDS repository, as well as a report to the appropriate public health agency P, using mechanisms which are outside the scope of this supplement.

6295 After reviewing the report, the public health agency P determines that a review of recent patients’ encounters with similar symptoms is necessary. Using Multi-Patient Queries, the health care provider is able to provide in a timely and accurate fashion all the documents with the having the same pre-determined value in the eventCodeList XDS metadata to the public health agency P. The public health agency is able to initiate an appropriate response and hence to contain a possible outbreak of the A1H1.

25.4.1.1 Post-factual and semi-real time reporting

6300 There are needs to aggregate data so that a pattern can emerge, but the patients’ identities need not to be known. For example, CDC (The Center for Disease Control and Prevention) or the InVS in France would like to know how many case of A1H1 are present at a national level at one point in time. In this case, there is no need to identify the patient, and unless other data is necessary to establish a trend (such as age, for example); an aggregated query on the metadata
6305 *eventCodeList* is sufficient using the *ObjectRefs* query. In this case irreversible pseudonymization or anonymization can be used since the data is employed statistically to generate a trend. This is the simplest case of implementing policies regarding security and privacy.

6310 There are other cases where statistical analysis in semi-real time is desired, such as an aggregated query at a district level to do profiling by region in times of an influenza epidemic. Again, this is a situation where the patient’s identity is not needed, but the number of cases and perhaps certain parameters such as the date. In order to be able to perform the aggregated queries, there has to be a minimum data set as per HIPAA recommendations.

25.4.1.2 Detailed queries

6315 If more scrutiny is needed, such as in patient safety (reporting to FDA a patient safety issue concerning medications, medical equipment malfunction, or surgical procedures), or population health monitoring such as the real-time control of an outbreak), detailed queries can be used.

6320 If in the Stored Query the *LeafClass* are specified the metadata of the document or of the folder (including the document ID and Repository ID) is returned. According to policies, these metadata can be pseudonymized or not.

For the multi-patient queries for detailed use, depending on the need, the policies regarding patient's privacy are different.

25.4.2 Technical Use Cases

6325 The output of a Multi-Patient Query can be in one of two forms: a list of opaque identifiers, each identifying a matching document (assuming that the query targets Document Entries and not Folders or Submission Sets); or full metadata where all details known in metadata are returned.

25.4.2.1 Opaque Identifiers

6330 Opaque identifiers, known in XDS as ObjectRefs, are useful to: discover the number of matches in the registry and then possibly to later retrieve the full metadata for the matching registry content. Applications that need only statistics (counts) can count the returned identifiers. Note that these identifiers represent documents (for example) that match the query and not patients. A single patient could have multiple matching documents.

25.4.2.2 Full Metadata

6335 A Multi-Patient Query can return full metadata, known as LeafClass in XDS. This metadata includes Patient Ids and patient demographics from potentially multiple patients so it is difficult to protect yet must be protected. Because of this sensitivity this type of return result would likely be only allowed by very highly trusted systems and thus this query is likely not to be available as widely as others.

25.5 Security Considerations

6340 This profile applies the same ATNA grouping to protect against the typical XDS identified risks. This profile may be grouped with XUA to further provide authentication of the user of the result.

6345 The new security and privacy considerations arise because this profile allows for a single query to result in multiple patients XDS metadata to be returned in one transaction. Although the XDS metadata is not high-grade health data it is still identifiable health information and thus needs to be protected. The combination of multiple patient's protected information in the same result results in a more difficult task to assure that the intended recipient has all the authorizations necessary for the intended use. In classical XDS queries the query request/response is constrained to a single patient and therefore the access control decision can be done across the whole transaction.

- 6350 This profile allows for two different types of return result. The ObjectRef result can be used to limit the exposure as this result will return only opaque identifiers of the matching documents. It is expected that this result would be more widely allowed. The Document Consumer can still obtain the full metadata but must use the classic XDS queries on an object-by-object basis thus allowing for transactions that are constrained to a single patient. This additional set of
- 6355 transactions to retrieve the metadata may be unnecessary when the system doing the query is authorized to use the LeafClass response. For example when the querying system is known as a system that will protect the information to the same degree. Where it is known that this querying system will apply the appropriate access control prior to ultimate use or disclosure.

26 Reserved

6360 **27 XCPD Integration Profile**

The Cross-Community Patient Discovery (XCPD) Profile supports the means to locate communities that hold patient relevant health data and the translation of patient identifiers across communities holding the same patient’s data. A community is defined as a group of facilities/enterprises that have agreed to work together using a common set of policies for the purpose of sharing health information within the community via an established mechanism. Facilities/enterprises may host any type of healthcare application such as EHR, PHR, etc. A community is identifiable by a globally unique id called the homeCommunityId. Membership of a facility/enterprise in one community does not preclude it from being a member in another community. Such communities may be XDS Affinity Domains which define document sharing using the XDS Profile or any other communities, no matter what their internal sharing structure.

6365

6370

27.1 Actors/ Transactions

Figure 27.1-1 shows the actors directly involved in the XCPD Integration Profile and the relevant transactions between them. Other actors that may be indirectly involved due to their participation in XCA, XDS, PIX, PDQ, etc. are not necessarily shown.

6375

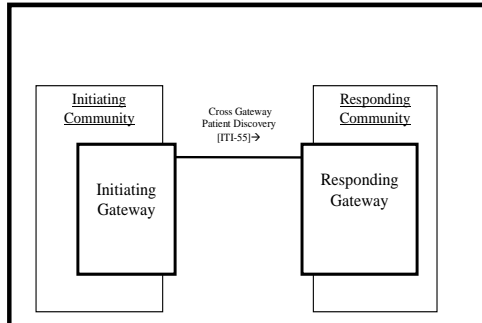


Figure 27.1-1: XCPD Actor Diagram

Table 27.1-1 lists the transactions for each actor directly involved in the XCPD Profile. In order to claim support of this Integration Profile, an implementation must perform the required transactions (labeled “R”). Transactions labeled “O” are optional. A complete list of options defined by this Integration Profile and that implementations may choose to support is listed in Section 27.2.

6380

Table 27.1-1: XCPD Integration Profile - Actors and Transactions

Actors	Transactions	Optionality	Section
Initiating Gateway	Cross Gateway Patient Discovery [ITI-55]	R	ITI TF-2b:3.55
Responding Gateway	Cross Gateway Patient Discovery [ITI-55]	R	ITI TF-2b:3.55

6385

27.1.1 Actors**27.1.1.1 Initiating Gateway**

6390 The Initiating Gateway supports all outgoing inter-community communications. XCPD uses this actor to initiate the Cross Gateway Patient Discovery [ITI-55]. The Initiating Gateway is required to support synchronous transaction messaging and may declare an option to support Asynchronous Web Services Exchange. Choosing Asynchronous Web Services Exchange will allow the Initiating Gateway to support workflows which scale to large numbers of communities because Asynchronous Web Services Exchange allows for more efficient handling of latency and scale.

6395 **27.1.1.2 Responding Gateway**

6400 The Responding Gateway supports all incoming inter-community communications. XCPD uses this actor to receive the Cross Gateway Patient Discovery [ITI-55]. The Responding Gateway is required to support Asynchronous Web Services Exchange on all implemented transactions. This allows the Initiating Gateway to choose the best of the two messaging patterns (synchronous or asynchronous) that fit the needs of the workflow. Support for Asynchronous Web Services Exchange allows for workflows which scale to large numbers of communities because it can handle latency and scale more efficiently.

27.1.2 Transactions**27.1.2.1 Cross Gateway Patient Discovery [ITI-55]**

6405 The Cross Gateway Patient Discovery transaction supports the ability for Initiating Gateways and Responding Gateways to discover mutually known patients. This transaction assumes an environment where patient data is well described and high quality demographic data is available.

Because the transaction supports the mutual discovery of patients it can be seen as having dual purposes.

- 6410
- To support a query by the Initiating Gateway requesting a demographically matching patient from within the Responding Gateway's community.
 - To support a feed to Responding Gateway announcing that the patient is known by the Initiating Gateway's community.

6415 This dual nature of the transaction is chosen for scalability purposes, as demographic matching algorithms are expensive on a large scale and once a match is identified it is important that both the initiating and responding sides of the transaction can use the results of that successful match.

The Cross Gateway Patient Discovery transaction has several modes, useful in different environments:

- 6420 • Demographic Query only mode – in this mode only the demographics of the patient are included in the request. The initiating community does not have, or does not choose to specify, a patient identifier for use by the Responding Gateway.
- Demographic Query and Feed – in this mode both the demographic and initiating community identifier are included in the request.
 - 6425 • Shared/national Patient Identifier Query and Feed – in this mode only a shared/national identifier is specified. Demographics are not necessary because matching can be done on the identifier alone.

This transaction can be used synchronously and asynchronously.

27.2 XCPD Integration Profile Options

6430 Options that may be selected for this Integration Profile are listed in Table 27.2-1 along with the Actors to which they apply. Dependencies between options when applicable are specified in notes.

Table 27.2-1: XCPD - Actors and Options

Actor	Options	Vol & Section
Initiating Gateway	<i>Asynchronous Web Services Exchange</i>	ITI TF-1: 27.2.1
	<i>Deferred Response</i>	ITI TF-1:27.2.2
Responding Gateway	<i>Deferred Response</i>	ITI TF-1:27.2.2

27.2.1 Asynchronous Web Services Exchange Option

6435 Initiating Gateways which support Asynchronous Web Services Exchange shall support Asynchronous Web Services Exchange on the Cross Gateway Patient Discovery [ITI-55]. Asynchronous processing is necessary to support scaling to large numbers of communities because Asynchronous Web Services Exchange allows for more efficient handling of latency and scale.

27.2.2 Deferred Response Option

6440 Responding Gateways which support the Deferred Response Option shall support Deferred Response as described in ITI TF-2b: 3.55.6.2 on the Cross Gateway Patient Discovery [ITI-55] transaction.

6445 Initiating Gateways which support the Deferred Response Option shall support Deferred Response as described in ITI TF-2b: 3.55.6.2 on the Cross Gateway Patient Discovery [ITI-55] transaction.

The Deferred Response Option reflects the more detailed understanding and feedback from implementers regarding processing that may result in significant delay. The existing

Asynchronous Web Services Exchange Option can support some scenarios with delayed response but not environments where the delay in responding may be as much as days or weeks. 6450 These cases require a mechanism that is managed by the application and which supports recovery through system restart. Deferred Response mode provides applications with such capability. In doing so it also adds responsibilities to the application, in particular for managing message correlation, creating application level acknowledgements and determining where to send a Deferred Response message. The new flexibility allowed by the Deferred Response 6455 Option is deemed worthy of these additional requirements on the application. For more information about Deferred Response and Asynchronous messaging in general see http://wiki.ihe.net/index.php?title=Asynchronous_Messaging.

27.3 XCPD Process Flow

27.3.1 Use Cases

6460 This section lists the use cases considered in developing this profile.

Use Case: Multiple primary residences

This use case describes the situation where a patient maintains more than one principal residence. Generally the principal residences are not geographically close so the medical data generated while in each residence would be created by separate institutions.

6465 A common example of this use case is what is described in the United States as the Snow Bird. This is a person who maintains two residences, one in the northern part of the USA for use during the hot summer months and one in the south for the colder winter. If a patient lives in Florida in the winter and in New York in the summer, this patient will likely have medical records in both places which need to be shared. If the patient is managing a long term medical 6470 condition, like diabetes, it will be important as she moves from New York to Florida and back that the background and related testing associated with management of the medical condition is readily accessible to the local physician.

Use Case: Border towns

6475 A variation of the multiple primary residence use case involves a patient who lives on the border between two communities or works and lives at some distance. A patient who lives in Longwy, France, which is on the border between the France and Belgium, may access health facilities both in France and Belgium. If that patient works in the European Investment Fund, he might also access health organizations in Luxembourg. All of these disparate areas, although closer geographically than the two French cities Longwy and Toulouse would probably hold patient 6480 information in separate domains thereby requiring the same kind of cross domain sharing as the snowbird case described above.

Use Case: Patient Move

A patient moves from one region to a different, remote region. The new region needs to access records from the patients prior location.

6485 **Use Case: Vacationer**

A patient is traveling and goes to the hospital. The hospital needs to access records from organizations in the patient's region of residence. Upon return from vacation, organizations in the region of residence need to access records from the remote hospital.

Use Case: Regional coalition

6490 Several medical facilities, related by region or other purpose, form an alliance to share medical data in an ongoing and integrated way. Patients across these facilities see seamless integration of medical records.

Use Case: Specialized treatment in different region

6495 The patient travels for specialized treatment and upon return requests the provider access records associated with the specialized treatment. Patient knows the city or region of the treatment but not the specific facility.

Use Case: Patient changes last name

The patient changes her last name as a result of marriage.

27.3.2 Detailed Interactions (Informative)

6500 This section details some of the likely interactions when using XDS, XCA and XCPD to share health data across communities.

Due to the peer-to-peer nature of XCPD interactions the Initiating Gateway must determine which communities to contact. XCPD does not address this question directly. Some possible approaches are:

6505 The patient carries an identification card which points to a Patient Medical Home. The Initiating Gateway uses the identification of the Patient Medical Home to determine the correct Responding Gateway to contact.

6510 A nation supports a registry of communities which can be searched using keys like homeCommunityId, regions within the nation, states, cities or other mechanisms to narrow a set of communities. When seeing a patient with a known home address a community may search the registry for communities near that home address and contact those first.

A nation would have a set of recognized communities that is accessible either through the registry or via configuration.

6515 Prior to any point-to-point sharing of protected health data the initiating and responding communities will need to agree to policies regarding the use of the data being shared. Those agreements may be:

Peer-to-Peer – each community negotiates with each other community in a one-to-one method.

6520 Regional – a regional government or regional organization creates an agreement that all within the region may agree to and by doing so join sharing within that region.

Across Region – A super regional (e.g., national) brings together sharing across regional sets of communities.

6525 Across Nation – At some point in the future nations may develop agreements for sharing across national boundaries, and as such will enable regions to share at a much larger scale.

27.3.2.1 Illustration of use of Transactions (Informative)

6530 The following sections illustrate scenarios for use of XCPD transactions. The illustrations assume grouping of XCPD Initiating and Responding Gateways and XCA Initiating and Responding Gateways⁹ for simplicity of the diagram and that each community is an XDS Affinity Domain. The interactions depicted in the diagrams are examples of use of the transactions, not requirements. The XCPD transactions are designed to enable a variety of behaviors. Implementers are free to choose some or none of the behaviors described.

Scenario # 1: Peer-to-peer

6535 Figure 27.3.2.1-1 shows the transactions involved in sharing healthcare data for one patient among three communities. Details on each interaction follow the diagram. “I&R Gateway” indicates the grouping of Initiating and Responding Gateways.

⁹ IHE defines an actor in the context of a profile, so an XCA Initiating Gateway does not have to be grouped with an XCPD Initiating Gateway and likewise for an XCA Responding Gateway and XCPD Responding Gateway. In practice these actors may customarily be grouped but it is not a requirement of IHE.

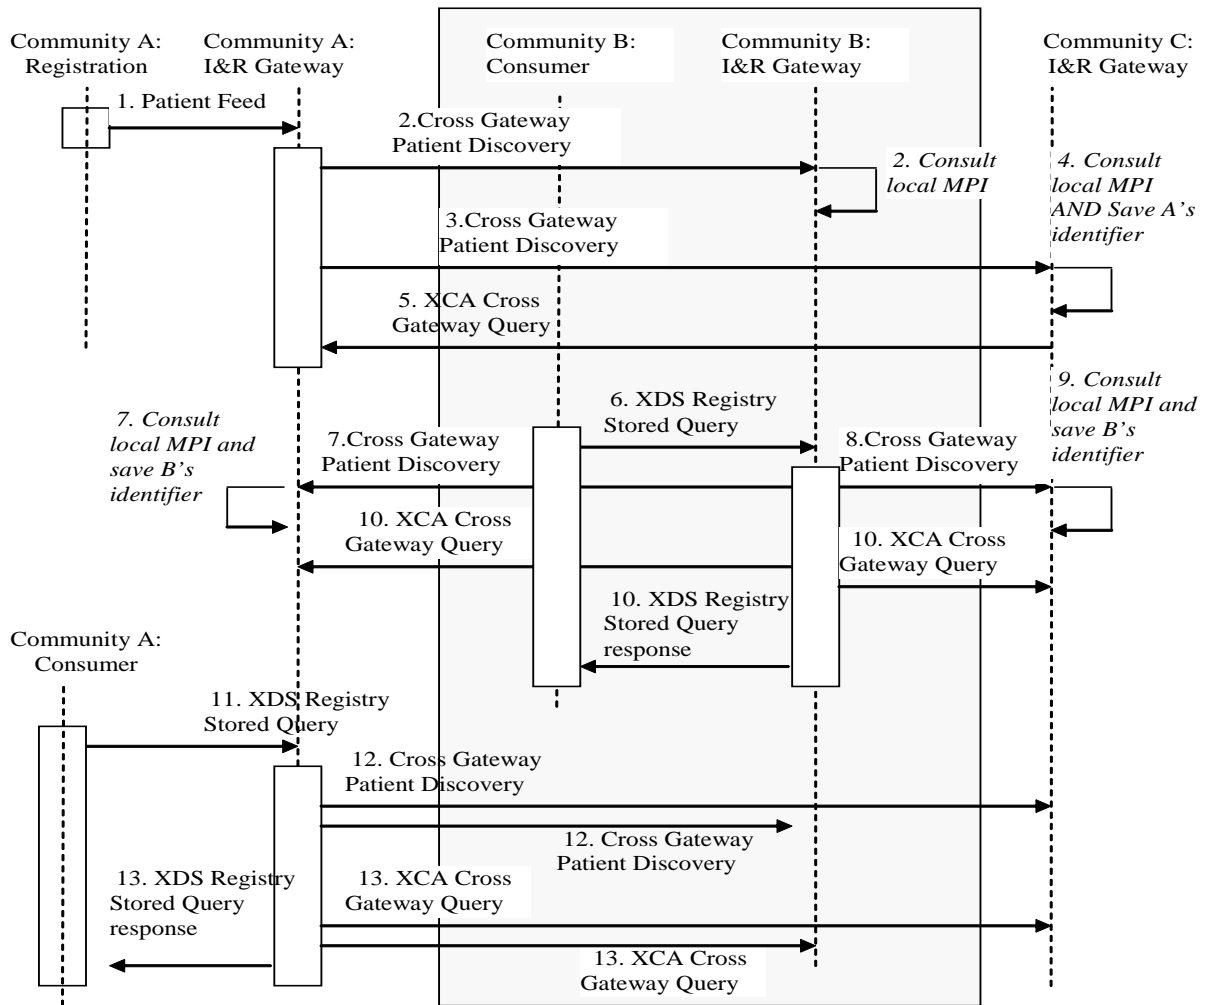


Figure 27.3.2.1-1: Detailed Interactions Peer-to-peer (Informative)

- 6540 • [1] The patient registers within Community A and a Patient Feed is sent to the Gateway.
- [2] The Gateway uses the Cross Gateway Patient Discovery transaction to determine if this patient is known in community B. Community B consults with its local MPI and, finding no match, responds with no matches, indicating the patient is not known in community B.
- [3] The Gateway uses the Cross Gateway Patient Discovery transaction to determine if this patient is known in community C. Community C responds with one match including the patient identifier in C.
- 6545 • [4] Community C consults with its local MPI and finds a match. It saves the identifier designated on the Cross Gateway Patient Discovery transaction as Community A's identifier for this patient.
- 6550 • [5] Community C pre-loads data for this patient by sending a XCA Cross Gateway Query to community A.

- [6] This patient is seen, for the first time, within an organization in community B which subsequently requests data about this patient and sends an XDS Registry Stored Query to its local Gateway.
- 6555 • [7] The Gateway uses the Cross Gateway Patient Discovery transaction to determine if this patient is known in community A. Community A consults its local MPI and responds with one match including the patient identifier in A. Community A saves the identifier designated on the Cross Gateway Patient Discovery transaction as community B's identifier for this patient.
- 6560 • [8] The Gateway uses the Cross Gateway Patient Discovery transaction to determine if this patient is known in Community C. Community C responds with one match including the patient identifier in C.
- [9] Community C consults with its local MPI and finds a match. It saves the identifier designated on the Cross Gateway Patient Discovery transaction as community B's identifier for this patient.
- 6565 • [10] The community B gateway sends a XCA Cross Gateway Query to both Community A and C because both responded positively to the Cross Gateway Patient Discovery transaction. Both responses are combined by the community B gateway and returned to the organization which originated the XDS Registry Stored Query in step [6].
- 6570 • [11] An organization in community A requests data about this patient and sends an XDS Registry Stored Query to its local Gateway.
- [12] The Gateway has saved the community B identifier in step [7] and Community A identifier in step [3]. But this query may happen days or weeks or years later. To verify the correlation of the identifier, community A's gateway repeats the Cross Gateway Patient Discovery to both B and C. B and C will consult their local MPI to identify the match.
- 6575 • [13] Having verified the correlation, community A sends an XCA Cross Gateway Query to both community B and C and combines the responses in order to respond to the XDS Registry Stored Query from step [11].

27.3.2.2 Hierarchical use of Cross Gateway Patient Discovery (Informative)

6580 The Cross Gateway Patient Discovery transaction has been designed to support the use of a Responding Gateway representing multiple organizations where each has its own patient identification domain and there is no patient identification domain in common across the organizations represented.

6585 Figure 27.3.2.2-1 shows the flow when an Initiating Gateway interacts with a Responding Gateway which is representing multiple organizations and reflecting those organizations in its Cross Gateway Patient Discovery response. The Responding Gateway interacts, using a community internal mechanism or any standard mechanism (e.g., XCPD, PDQ), with each of the Master Patient Indices (MPI) to find a match for the specified demographics. The figure shows the Responding Gateway interacting with three MPI's and discovering three matching patient identifiers: X in the A domain, Y in the B domain and Z in the C domain. All three of the patient identifiers are returned in the Cross Gateway Patient Discovery response in a way that indicates that they represent three different domains, rather than three matches from one domain (which is

6590

6595 also possible and coded in a different way see ITI TF-2b: 3.55.4.2.2.4). In order for the Initiating Gateway to access all data it must issue three separate XCA Cross Gateway Query transactions, each with one of the patient identifiers returned from the Cross Gateway Discovery transaction.

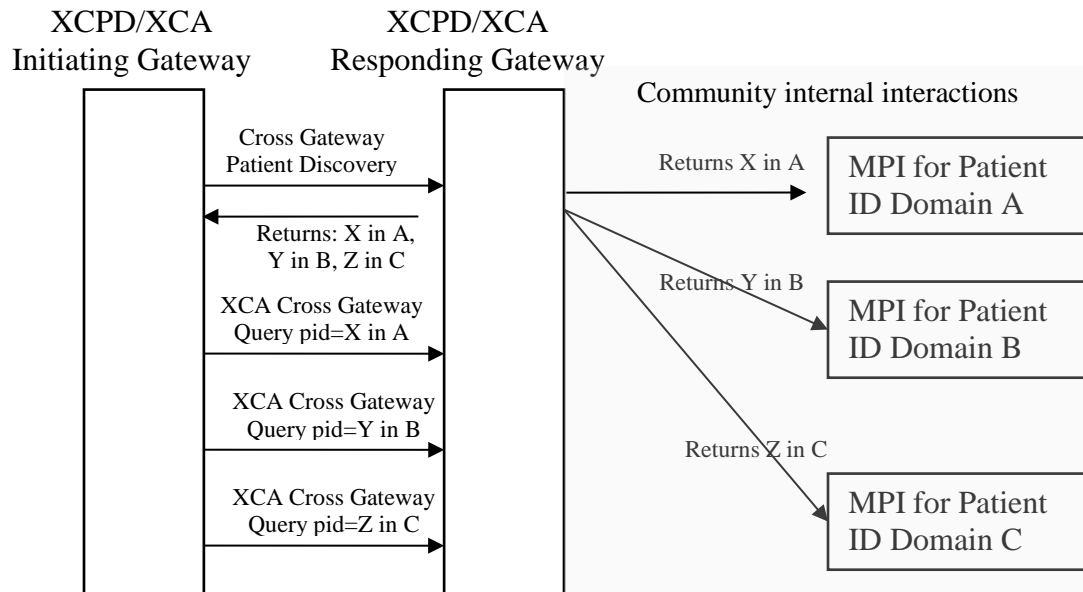


Figure 27.3.2.2-1: Hierarchical use of Cross Gateway Query (Informative)

6600 **27.4 XCPD Security Considerations**

The risk analysis for XCPD enumerates assets, threats, and mitigations. The complete risk data is stored and maintained in a central location. The complete risk data is stored and available from IHE¹⁰.

6605 The purpose of this risk assessment is to notify vendors of some of the risks that they are advised to consider in implementing XCPD actors. For general IHE risks and threats please see ITI TF-1: Appendix L. The vendor is also advised that many risks cannot be mitigated by the IHE profile and instead the responsibility for mitigation is transferred to the vendor, and occasionally to the XDS Affinity Domain and enterprises. In these instances, IHE fulfills its responsibility to notify affected parties through the following section.

6610 **27.4.1 Requirements/Recommendations**

The following mitigations shall be implemented by all XCPD actors. These mitigations moderate all high impact risks.

¹⁰ The risk analysis data may be found at: ftp://ftp.ihe.net/IT_Infrastructure/iheityr7-2009-2010/Technical_Cmte/Profile_Work/XCPD/XCPD_Risk_assessment_and_mitigation_table.xls

All actors in XCPD shall be grouped with an ATNA Secure Node Actor (or ATNA Secure Application) and a CT Time Client Actor.

6615 As a consequence of grouping with ATNA Secure Node or Secure Application both incoming and outgoing messages will be via a secure communication channel, including all asynchronous response messages.

The following mitigations are transferred to the vendors, XDS Affinity Domains, and enterprises.

6620 Network protection services are recommended to be sufficient to guard against denial of service attacks on all service interfaces.

A process that reviews audit records and acts on inappropriate actions is recommended.

It is recommended that service interfaces be implemented with a good design to guard against corruption and denial of service attacks

6625 **27.4.2 Policy Choices**

Policy choices will not be addressed by this profile. Each community may have different policies. The profile has been designed with this fact in mind and an understanding of enough variety of policies so that any reasonable policy can be implemented without violating the profile.

6630

Appendix A: Actor Descriptions

6635 Actors are information systems or components of information systems that produce, manage, or act on information associated with operational activities in the enterprise. The following are definitions of actors used in the IHE IT Infrastructure Integration Profiles:

Audit Repository – This actor provides a repository for audit events. IHE does not specify what analysis and reporting features should be implemented for an audit repository.

Client Authentication Agent – Provides local management of user authentication.

6640 **Context Manager** – This actor serves as a broker for the communication between two or more context participant actors (either Patient Context Participant or User Context Participant). It supports the passing of the user and patient subjects.

Display – A system that can request specific information or documents from an Information Source and display them.

6645 **Document Source** – The Document Source Actor is the producer and publisher of documents. It is responsible for sending documents to a Document Repository Actor. It also supplies metadata to the Document Repository Actor for subsequent registration of the documents with the Document Registry Actor.

6650 **Document Consumer** – The Document Consumer Actor queries for document metadata meeting certain criteria, and may retrieve selected documents.

Document Recipient – This actor receives a set of documents sent by another actor. Typically this document set will be made available to the intended recipient who will choose to either view it or integrate it into a Health Record.

6655 **Document Registry** – The Document Registry Actor maintains metadata about each registered document in a document entry. This includes a link to the Document in the Repository where it is stored. The Document Registry responds to queries from Document Consumer actors about documents meeting specific criteria. It also enforces some healthcare specific technical policies at the time of document registration.

6660 **Document Repository** – The Document Repository is responsible for both the persistent storage of these documents as well as for their registration with the appropriate Document Registry. It assigns a URI to documents for subsequent retrieval by a Document Consumer.

DNS Server – This actor has authoritative location information.

Form Archiver – the actor responsible for receiving form instance data for archival purposes.

6665 **Form Filler** – the actor responsible for retrieving a form from a Form Manager, and for submitting form instance data to a Form Receiver. The Form Filler may optionally be responsible for retrieving clarifications information from a Form Manager.

Form Manager – the actor that supplies a form based upon a request that supplies a form identification. The Form Manager also supplies clarification information.

- 6670 **Form Processor** – an integrated Form Manager and Form Receiver, supporting all of the transactions and options of those actors. The Form Processor constructs forms such that form data is submitted back to itself.
- Form Receiver** – the actor that receives form instance data.
- Initiating Gateway** - supports all outgoing inter-community communications.
- 6675 **Information Source** – A system that responds to requests for specific information or documents and returns ready for presentation information to be displays on the requesting actor.
- Kerberos Authentication Server** – Provides central authentication of enterprise users.
- Kerberized Server** – Receives user authentication information for further use by the service that contains this actor
- 6680 **On-Demand Document Source** - The On-Demand Document Source is the producer and publisher of On-Demand Documents. This actor registers On-Demand Document Entries with the Document Registry and responds to Retrieve Document Set transactions with a document reflecting current information for the entry requested.
- 6685 **Patient Context Participant** – This actor participates in a shared context environment by both setting the patient context and responding to context changes as communicated by the Context Manager Actor. This actor shall respond to all patient context changes. This actor shall set the patient context, if the application containing this actor has patient selection capability.
- Patient Demographics Consumer** – A system that uses demographic information provided by the Patient Demographics Supplier about a patient.
- 6690 **Patient Demographics Supplier** – A system responsible for adding, updating and maintaining demographics about a patient, and additional information such as related persons (primary caregiver, guarantor, next of kin, etc.). It supplies new and updated information to the Patient Demographics Consumer.
- 6695 **Patient Encounter Source** – A system responsible for adding, updating and maintaining encounter information about a patient. It supplies new and updated information to the Patient Encounter Consumer.
- Patient Encounter Consumer** – A system that uses patient encounter information provided by the Patient Encounter Source about a patient.
- 6700 **Patient Identifier Cross-reference Consumer** – This actor allows a system in a Patient Identifier Domain to determine the identification of a patient in a different Patient Identifier Domain by using the services of a Patient Identifier Cross-Reference Manager Actor.
- Patient Identifier Cross-reference Manager** – Serves a well-defined set of Patient Identifier Domains. Based on information provided in each Patient Identifier Domain by a Patient Identification Source Actor, it manages the cross-referencing of patient identifiers across Patient Identifier Domains.
- 6705 **Patient Identity Source** – The Patient Identity Source Actor is a provider of unique identifier for each patient and maintains a collection of identity traits. Each Patient Identifier Domain requires this Actor to assign patient identities and to notify other Actors (e.g., a Patient Identifier

Cross-reference Manager or a Document Registry Actor) of all events related to patient identification (creation, update, merge, etc.).

6710 **Personnel White Pages Consumer** – This actor has a use for information that can be found in the Personnel White Pages Directory.

Personnel White Pages Directory – This actor has authoritative Personnel White Pages information on the human workforce members of the enterprise.

6715 **Portable Media Creator** – This actor assembles the content of the media and writes it to the physical medium. A priori this document set is extracted from an Electronic Healthcare Record (EHR) or a Personal Health Record (PHR) system.

6720 **Portable Media Importer** – This actor reads and displays the information contained on the media, allows the user to select information, and store any or all of the elements. Typically this document will be integrated into an Electronic Healthcare Record (EHR) or Personal Health Record (PHR) and can then process the instances.

Responding Gateway – supports all incoming inter-community communications.

6725 **Secure Application** – This actor is a subset of Secure Node where the security capabilities (user authentication, secure communications, security audit recording, and security policy enforcement) are required to cover the product (e.g., Web Applications) identified in the IHE Integration Statement and not necessarily the whole node.

Secure Node – The presence of this actor on a system means that all of the other actors and other non-IHE software comply with the IHE rules for user authentication, secure communications, security audit recording and security policies.

6730 **Time Client** – Establishes time synchronization with one or more Time Servers using the NTP protocol and either the NTP or SNTP algorithms. Maintains the local computer system clock synchronization with UTC based on synchronization with the Time Servers.

Time Server – Provides NTP time services to Time Clients. It is either directly synchronized to a UTC master clock (e.g., satellite time signal) or is synchronized by being grouped with a Time Client to other Time Server(s).

6735 **User Context Participant** – Receives notification of user context changes and follows them for the application that contains it.

Value Set Consumer – an actor who may

1. Retrieve an Expanded Value Set based on its OID and possibly its version if the latter is available.
2. Retrieve metadata described a Value Set, together with the expanded value set that corresponds to its current meaning.

Value Set Repository – actor whose role is to provide Expanded Value Sets and metadata describing these value sets.

X-Service Provider – System providing a service that needs an X-User Assertion.

6745 **X-Service User** – System making a services request of an X-Service Provider.

Appendix B: Transaction Descriptions

Transactions are interactions between actors that transfer the required information through standards-based messages. The following are brief descriptions of the transactions defined by IHE.

- 6750 1. **Maintain Time:** NTP transactions used to maintain time synchronization.
2. **Get User Authentication:** The Client Authentication Agent requests user authentication from the Kerberos Authentication Server. When the user is authenticated, the Kerberos Authentication Server returns a Ticket Granting Ticket (TGT) to optimize future activity.
3. **Get Service Ticket:** Obtain a ticket using Kerberos protocol for use with a service.
- 6755 4. **Kerberized Communication:** The Kerberized Communication transaction is an aspect of the connection between a local client and a remote server.
5. **Join Context:** Allows a Context Participant Actor to locate and establish communication with the Context Manager Actor.
6. **Change Context:** Includes all messages required to initiate and finalize a context change transaction:
6760
 - Initiation of a context change request from the instigating participant actor
 - Delivery of survey results to instigating actor and display of associated replies
 - Communication of context change decision to the Context Manager Actor
7. **Leave Context:** Allows Context Participant Actor to notify the Context manager Actor that it is breaking off communication.
6765
8. **Patient Identity Feed:** Allows a Patient Identity Source Actor to notify a Patient Identifier Cross-Reference Manager Actor of all events related to patient identification (creation, update, merge, etc.).
9. **PIX Query:** This transaction allows a Patient Identifier Cross-reference Consumer to find out the identification of a patient in different Patient Identifier Domains by using the services of a Patient Identifier Cross-reference Manager Actor.
6770
10. **PIX Update Notification:** Allows a Patient Identifier Cross-reference Consumer to be notified by the Patient Identifier Cross-reference Manager Actor of changes to the identification of all patients in Patient Identifier Domains the Consumer is interested in.
- 6775 11. **Retrieve Specific Information for Display:** A request issued by a display system for specific information related to a patient returned in a ready for presentation information format.
12. **Retrieve Document for Display:** A display system requests an instance of a uniquely identified persistent document under custodianship by an information source and receives its content ready for presentation.
6780

13. **Follow Context:** Accounts for all messages required to propagate a context change to a responding participant actor:
- Survey of all other Context Participant Actors by the Context Manager Actor and display by the instigating Participant Actor of any associated replies
- 6785
- Notification of context change result from the Context manager Actor to the Context Participant Actors
 - Retrieval of the context data by the Context Participant Actors
14. Intentionally Left Blank
15. Intentionally Left Blank
- 6790
16. Intentionally Left Blank
17. Intentionally Left Blank
18. **Registry Stored Query:** The Registry Stored Query transaction is issued by the Document Consumer Actor on behalf of a care provider (EHR-CR) to a Document Registry. The Document Registry Actor searches the registry to locate documents that meet the provider's specified query criteria. It will return a list of document entries that contain metadata found to meet the specified criteria including the locations and identifier of each corresponding document in one or more Document Repositories.
- 6795
19. **Node Authentication:** This transaction is embedded within all network communications activity. All DICOM, HL7, and HTML connections shall comply with the IHE specification for bi-directional authentication and authorization of communications of Protected Healthcare Information (PHI). IHE does not specify how other protocols that transfer PHI shall perform bi-directional authentication and authorization, but requires that other protocols perform such authentication and authorization.
- 6800
20. **Record Audit Event:** The delivery of an audit event description from any secure node to the Audit Repository.
- 6805
21. **Patient Demographics Query:** Look up and return patient demographic information in a single patient demographics source, based upon matches with full or partial demographic information entered by the user.
22. **Patient Demographics and Visit Query:** Look up and return patient demographic and visit information in a single patient demographics source, based upon matches with full or partial demographic/visit information entered by the user.
- 6810
23. **Find Personnel White Pages:** This transaction will find the LDAP Directory by querying the DNS.
24. **Query Personnel White Pages:** This transaction provides for read-only access to the Personnel White Pages directory.
- 6815
25. **Reserved for Send Notification**
26. **Reserved for Receive Notifications**

- 27. **Reserved for** Send Acknowledgement
- 28. **Reserved for** Receive Acknowledgement
- 6820 29. **Reserved for** Cross Enterprise User Authentication
- 30. **Patient Identity Management:** The Patient Demographics Supplier registers or updates a patient and forwards the demographic information (i.e., all information directly related to the patient, such as ID, address, next of kin, guarantor, etc.) to other systems implementing the Patient Demographics Consumer Actor.
- 6825 31. **Patient Encounter Management:** The Patient Encounter Supplier registers or updates an encounter (inpatient, outpatient, pre-admit, etc.) and forwards the information to other systems implementing the Patient Encounter Consumer Actor. This information will include the patient's location and care providers for a particular (usually current) encounter.
- 6830 32. **Distribute Document Set on Media:** A source actor (Portable Media Creator) writes a set of documents on an interchange media. The media is physically transported to another actor (Portable Media Importer) which then imports the document set, or sent as a ZIP attachment via Email. The media can also be provided to a patient or a referring physician for web-based viewing.
- 6835 33. placeholder
- 34. **Retrieve Form:** This transaction retrieves the requested form from a Form Manager.
- 35. **Submit Form:** This transaction submits form instance data, in XML format, to a Form Receiver.
- 6840 36. **Archive Form:** This transaction supplies the form instance data, in XML format, to a Form Archiver.
- 37. **Retrieve Clarifications:** This transaction retrieves a set of clarifications from a Form Manager.
- 38. **Cross Gateway Query:** send a query from one community to another in order to identify where healthcare information satisfying specific restraints is located.
- 6845 39. **Cross Gateway Retrieve:** request the retrieval of a specific set of healthcare information (a document or documents) from a remote location.
- 40. **Provide X-User Assertion:** This transaction provides a trustable user assertion from the service user to the service provider
- 6850 41. **Provide and Register Document Set-b:** A Document Source Actor initiates the Provide and Register Document Set-b transaction. For each document in the submitted set, the Document Source Actor provides both the documents as an opaque octet stream and the corresponding metadata to the Document Repository. The Document Repository is responsible to persistently store these documents, and to register them in the Document Registry using the Register Document Set-b transaction by forwarding the document metadata received from the Document Source Actor.
- 6855

- 6860 42. **Register Document Set-b:** A Document Repository Actor initiates the Register Document Set-b transaction. This transaction allows a Document Repository Actor to register one or more documents in a Document Registry, by supplying metadata about each document to be registered. This document metadata will be used to create XDS Submission Set, XDS Document, and potentially XDS Folder Entries in the registry. The Document Registry Actor ensures that document metadata is valid before allowing documents to be registered. If one or more documents fail the metadata validation, the Register Document Set-b transaction fails as a whole.
- 6865 43. **Retrieve Document Set:** A Document Consumer Actor initiates the Retrieve Document Set transaction. The Document Repository will return the set of documents that was specified by the Document Consumer.
- 6870 44. **Patient Identity Feed HL7 V3:** Allows a Patient Identity Source Actor to notify a Patient Identifier Cross-Reference Manager Actor of all events related to patient identification (creation, update, merge, etc.) using HL7 V3.
- 6875 45. **PIX V3 Query:** Using HL7 V3, this transaction allows a Patient Identifier Cross-reference Consumer to find out the identification of a patient in different Patient Identifier Domains by using the services of a Patient Identifier Cross-reference Manager Actor.
- 6880 46. **PIX V3 Update Notification:** Using HL7 V3, this allows a Patient Identifier Cross-reference Consumer to be notified by the Patient Identifier Cross-reference Manager Actor of changes to the identification of all patients in Patient Identifier Domains the Consumer is interested in.
- 6885 47. **Patient Demographics Query HL7 V3:** Using HL7 V3, look up and return patient demographic information in a single patient demographics source, based upon matches with full or partial demographic information entered by the user.
48. placeholder
49. placeholder
- 6890 50. **Multi-Patient Stored Query:** A Document Consumer Actor issues a Multi-Patient Stored Query to a Document Registry to locate documents that meet the user's specified query criteria. The Document Registry returns a list of document entries pertaining to multiple patients found to meet the specified criteria, including the locations and identifier of each corresponding document in one or more Document Repository.
51. **Retrieve Value Set:** The Value Set Consumer retrieves an Expanded Value Set from the Value Set Repository.
52. **Retrieve Multiple Value Sets:** The Value Set Consumer retrieves multiple value sets from the Value Set Repository. These retrieved value sets have metadata that matches retrieval selection criteria. The retrieved sets provide the full metadata describing the value set, and an expanded value set representation for that value set.
53. placeholder

- 6895 54. placeholder
55. **Cross Gateway Patient Discovery**: supports the ability to discover patients mutually know across communities.
56. placeholder
57. placeholder
- 6900 58. placeholder
59. placeholder
60. placeholder
61. **Register On-Demand Document Entry**: The Register On-Demand Document Entry transaction is used by the On-Demand Document Source to register one or more On-Demand Document Entries in the Document Registry.
- 6905

Appendix C: IHE Integration Statements

6910 IHE Integration Statements are documents prepared and published by vendors to describe the conformance of their products with the IHE Technical Framework. They identify the specific IHE capabilities a given product supports in terms of IHE actors and integration profiles (described in ITI TF-1:2).

6915 Users familiar with these concepts can use Integration Statements to determine what level of integration a vendor asserts a product supports with complementary systems and what clinical and operational benefits such integration might provide. Integration Statements are intended to be used in conjunction with statements of conformance to specific standards (e.g., HL7, IETF, DICOM, W3C, etc.).

6920 IHE provides a process for vendors to test their implementations of IHE actors and integration profiles. The IHE testing process, culminating in a multi-party interactive testing event called the Connect-a-thon, provides vendors with valuable feedback and provides a baseline indication of the conformance of their implementations. The process is not intended to independently evaluate, or ensure, product compliance. In publishing the results of the Connect-a-thon and facilitating access to vendors' IHE Integration Statements, IHE and its sponsoring organizations are in no way attesting to the accuracy or validity of any vendor's IHE Integration Statements or any other claims by vendors regarding their products.

6925 **IMPORTANT -- PLEASE NOTE: Vendors have sole responsibility for the accuracy and validity of their IHE Integration Statements.** Vendors' Integration Statements are made available through IHE simply for consideration by parties seeking information about the integration capabilities of particular products. IHE and its sponsoring organizations have not evaluated or approved any IHE Integration Statement or any related product, and IHE and its
6930 sponsoring organizations shall have no liability or responsibility to any party for any claims or damages, whether direct, indirect, incidental or consequential, including but not limited to business interruption and loss of revenue, arising from any use of, or reliance upon, any IHE Integration Statement.

C.1 Structure and Content of an IHE Integration Statement

6935 An IHE Integration Statement for a product shall include:

1. The Vendor Name
2. The Product Name (as used in the commercial context) to which the IHE Integration Statement applies.
3. The Product Version to which the IHE Integration Statement applies.
- 6940 4. A publication date and optionally a revision designation for the IHE Integration Statement.
5. The following statement: "This product implements all transactions required in the IHE Technical Framework to support the IHE Integration Profiles, Actors and Options listed below:"

- 6945 6. A list of IHE Integration Profiles supported by the product and, for each Integration Profile, a list of IHE Actors supported. For each integration profile/actor combination, one or more of the options defined in the IHE Technical Framework may also be stated. Profiles, Actors and Options shall use the names defined by the IHE Technical Framework Volume 1. (Note: The vendor may also elect to indicate the version number of the Technical Framework referenced for each Integration Profile.)
- 6950

Note that implementation of the integration profile implies implementation of all required transactions for an actor as well as selected options.

The statement shall also include references and/or internet links to the following information:

- 6955
1. Specific internet address (or universal resource locator [URL]) where the vendor's Integration Statements are posted
 2. URL where the vendor's standards conformance statements (e.g., HL7, DICOM, etc.) relevant to the IHE transactions implemented by the product are posted.
 3. URL of the IHE Initiative's web page for general IHE information <http://www.ihe.net/>.

6960 An IHE Integration Statement is not intended to promote or advertise aspects of a product not directly related to its implementation of IHE capabilities.

C.2 Format of an IHE Integration Statement

Each Integration Statement shall follow the format shown below. Vendors may add a cover page and any necessary additional information in accordance with their product documentation policies.

6965

IHE Integration Statement		Date	12 Oct 2003
Vendor	Product Name	Version	
Any Medical Systems Co.	IntegrateRecord	V2.3	
This product implements all transactions required in the IHE Technical Framework to support the IHE Integration Profiles, Actors and Options listed below:			
Integration Profiles Implemented	Actors Implemented	Options Implemented	
Retrieve Information for Display	Information Source	none	
Enterprise User Authentication	Kerberized Server	none	
Patient Identity Cross-referencing	Patient Identifier Cross-reference Consumer	PIX Update Notification	
Internet address for vendor's IHE information:		www.anymedicalsystemsco.com/ihe	
Links to Standards Conformance Statements for the Implementation			
HL7	www.anymedicalsystemsco.com/hl7		
Links to general information on IHE			
North America: www.ihe.net	In Europe: www.ihe-europe.org	In Japan: www.jira-net.or.jp/ihe-j	

Appendix D: User Authentication Techniques - Passwords, Biometrics, and Tokens

6970 Authentication techniques are based on one or more of three factors: Something you know, something you are, or something you have. There are many different authentication techniques in use today. The technologies supporting these techniques are not well standardized. There are also excellent security reasons to avoid specifying any single set of technologies for authentication use.

6975 The Kerberos protocol was originally defined to work with any user authentication technique. Kerberos has been shown to support a wide variety of authentication technologies. These include various forms of tokens and biometric technologies. Specific implementations of these technologies often include proprietary components. There is often a pair of proprietary components added – one at the user workstation and a matching component at the authentication server. Once the user authentication is complete, the subsequent Kerberos transactions are the same.

6980 These extensions are not yet standardized. The IHE specification for the use of Kerberos does not prevent the use of these extensions at a specific site, nor does it ensure that the extensions will work.

6985 The Kerberos system specified for the Enterprise User Authentication utilizes a challenge response system together with a username and password system to authenticate the user. The minimal support of passwords provides a standardized baseline for the IHE “Enterprise User Authentication”. Kerberos enables enforcement of a central password policy which facilitates stronger passwords. Such password policies are beyond the scope of IHE. Kerberos does not prevent the use of weak passwords. The password strength policy must be chosen and enforced by the site security administration.

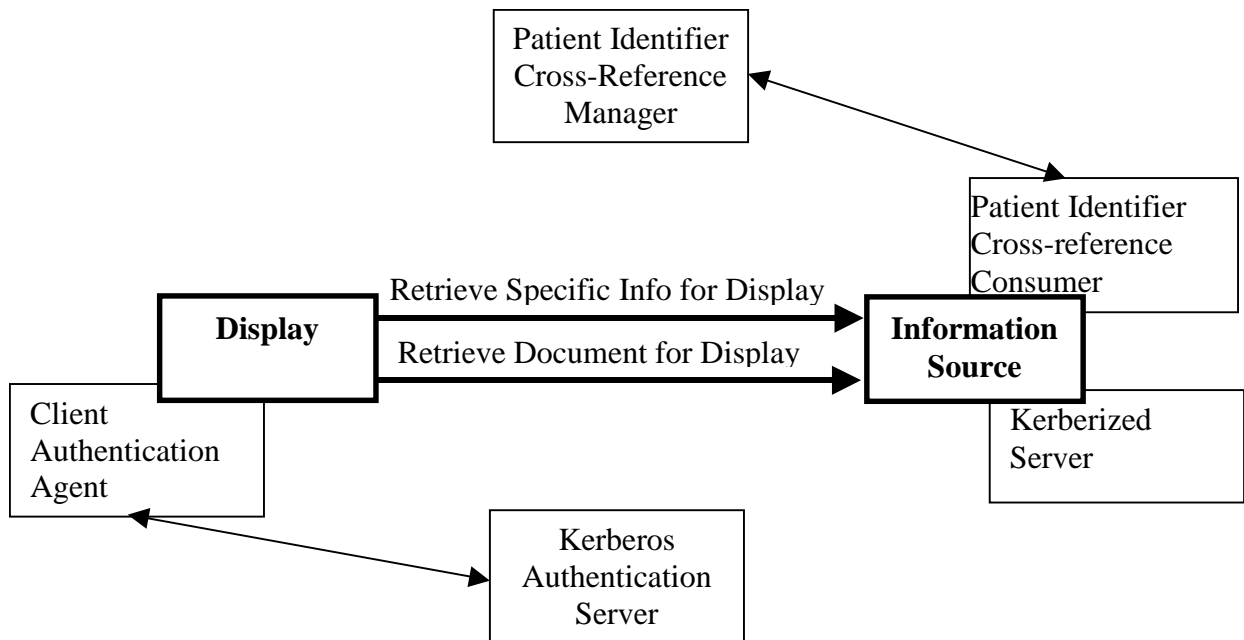
6990 **Appendix E: Cross Profile Considerations**

E.1 Combined use of RID, EUA and PIX Integration Profiles

When used alone, the Retrieve Information for Display Integration Profile assumes that the Patient Identifier Domain is the same for both the Display and the Information Source Actors. Furthermore, any user authentication on the Information Source is not addressed explicitly. This Appendix discusses combination of the Retrieve Information for Display Integration Profile with other IHE Integration Profiles to address these two problems.

When used in conjunction with the Patient Identifier Cross-referencing Integration Profile, implementations of the Retrieve Information for Display Integration Profile shall take into account that the Information Source Actor may need to map Patient IDs from different identifier domains to the one used in its own domain. The combined use of these Integration Profiles is achieved by grouping the Information Source and the Patient Identifier Cross-reference Consumer Actors. This is depicted in Figure E-1.

Similarly, the Information Source Actor may perform certain access control functions based on the requesting user authentication performed by the actors implementing the Enterprise User Authentication Integration Profile. The combined use of these Integration Profiles is achieved by grouping the Display Actor with the Client Authentication Agent Actor and the Information Source Actor with the Kerberized Server Actor. This is also shown in Figure E-1.



7010 **Figure E-1: Combined use of actors implementing multiple Integration Profiles**

E.2 XDS Integration with RID

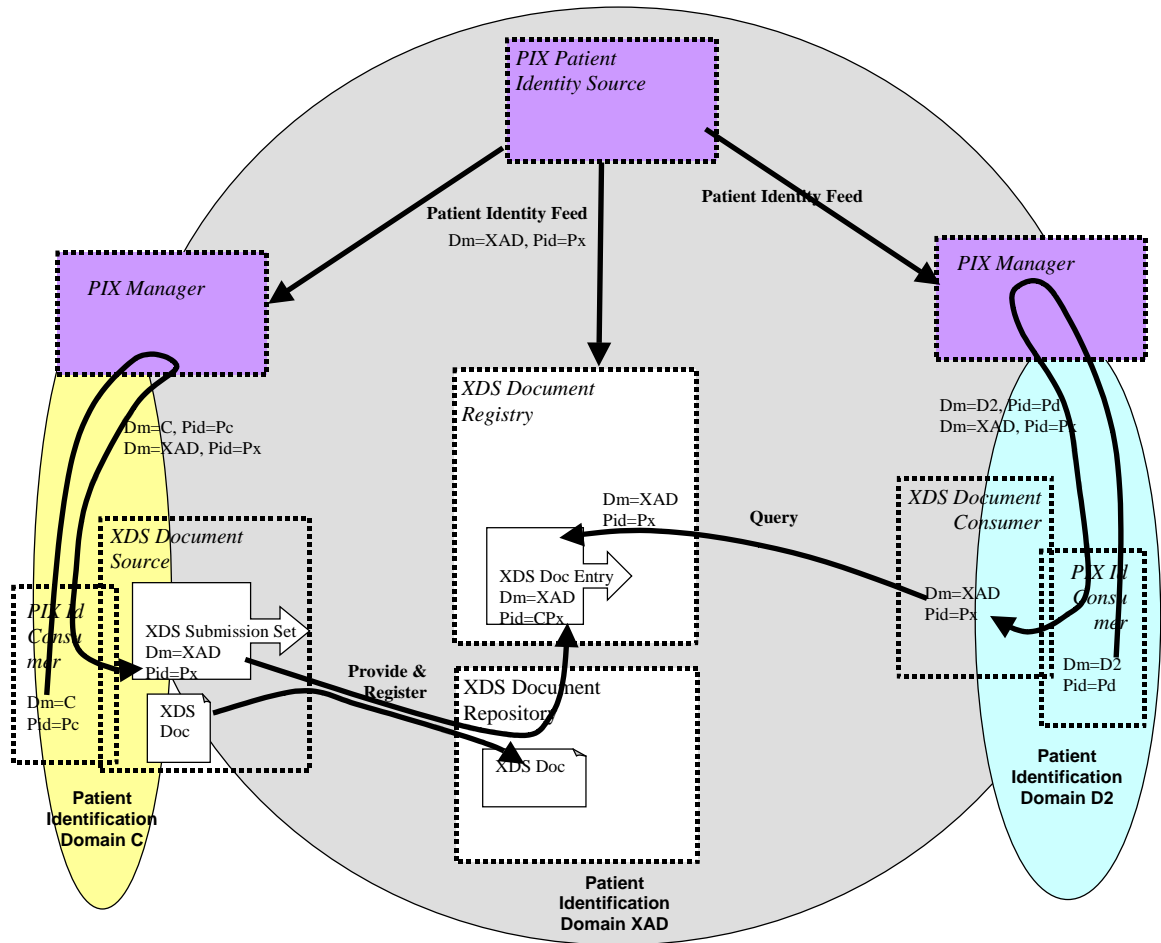
7015 The RID Retrieve Document for Display transaction [ITI-12] is compatible with the XDS.a Retrieve Document transaction [ITI-17]. Thus, an RID Information Source implementing the Retrieve Document for Display transaction can be used to implement the XDS.a Retrieve Document transaction. In this instance, the RID Information Source must be a Secure Node [see ATNA].

Note: The XDS.a profile is now deprecated. RID is not compatible with XDS.b Retrieve Document Set transaction.

7020 **E.3 XDS Integration with PIX**

All Patient IDs managed in the XDS transactions (either in XAD-Pid Domain or in an EHR-CR Domain) shall include the related Patient Domain ID (OID of the Assigning Authority) associated with the patient ID. It is recommended that this unambiguous patient identification be used with Patient IDs within the Documents also.

7025 Because XDS is Document content neutral, there is no verification by the XDS Repository that the Patient IDs included inside the documents are consistent with the patient IDs managed by the Registry in the document entry related to that document.



7030 **Figure E.3-1: XDS Affinity Domain with patient ID cross-referencing with IHE PIX Managers**

7035 Figure C.6-1 depicts an example of an XDS Affinity Domain with a Patient Identifier Domain (called XAD) and two EHR-CRs where the cross-referencing is performed by Patient Identifier Cross Referencing Managers internal to both the Document Source and the Document Consumer Domains (called C and D2 respectively).

7040 A Document Source may choose to perform the cross-referencing of its own patient IDs in that of the XAD-Pid Domain by leveraging the IHE PIX Integration Profile (See figure). The Patient ID Feed Transaction from the XAD Patient ID Source may be used to provide input to the Patient Identifier Cross-Referencing Manager used by the Document Source. The PIX manager may either be internal to the EHR-CRs or be shared across the XDS Affinity Domain.

E.4 XDS Integration with PWP

The XDS Document Source Actor in the XDS Integration Profile may choose to utilize the PWP Query Personnel White Pages [ITI-24] transaction to obtain information needed to fill the

7045 authorPerson and legalAuthenticatorName fields for the Register Document Set – b [ITI-42] and Provide & Register Document Set – b [ITI-41] transactions.

7050 The Personnel White Pages transaction defines, in ITI TF-2a: 3.24.4.1.2.3.1, a “cn” attribute with “lang-x-ihe” that contains the information in the HL7 XCN (extended composite ID number and name for persons) format for personal information. These fields are optional in the PWP Integration Profile. A care delivery organization may choose to populate these fields in their Personnel White Pages Directory and utilize the ITI-24 transaction to support its XDS activities. This is not a required dependency, but is a possible reason to group a Document Source Actor with a Personnel White Pages Consumer Actor.

7055 The PWP Integration Profile only provides the personnel information. Organizational information must be obtained via other means, e.g., extending the LDAP directory with organizational objects.

E.5 XDS Integration with PDQ

7060 The Patient Demographics Query (PDQ) Integration Profile may be used in conjunction with the XDS Integration Profile to provide a lookup for XDS Affinity Domain Patient Identifiers to XDS Document Consumer and Document Source Actor. In this case a Patient Demographics Supplier Actor needs to be grouped with the XDS Patient Identifier Source Actor on one hand, and on the other hand a Patient Demographics Consumer Actor needs to be grouped with the Document Source/Consumer where one may want to query based on local patient traits and obtain a pick-list of candidate Patient Ids in the XAD Patient Identifier Domain. This offers a simpler solution than the use of the PIX Integration Profile.

E.6 XDM Integration with XDS, Content Integration Profiles, PIX, and DSG

7065 The XDM Profile does not constrain the document types or purposes. Content Integration Profiles may impose such constraints.

7070 The XDM Profile does not address the issue of patient reconciliation. The PIX and PDQ functionality might be available to a Portable Media Importer, but the XDM Profile does not require it. If there is no PIX or PDQ available to the Portable Media Importer, some other method for performing the necessary coercion of patient identifiers must be provided. This might be manual for Portable Media Importers that are intended for very small sites.

7075 The Cross-enterprise Document Media Interchange (XDM) Integration Profile may be used in conjunction with the DSG Integration Profile to provide for the digital signature of the documents content and of the XDS metadata.

E.7 XDM/XDR Distinction

7080 Both XDR and XDM describe the exchange of a set of patients' documents. They are relevant in situations where XDS is not yet implemented or available at one of the participating organizations or where point-to-point (versus sharing through a registry) interaction is desired.

7085 XDM is applicable in situations where the information receiver is an individual who will manually interpret or examine the data and associated documents. XDM allows for one exchange which contains documents relating to multiple patients and can be used in situations where no continuous networking capability is available on one or both of the participating healthcare providers.

7090 XDR is applicable in situation where the information exchanged is going to an automated application or robust system capable of automated storage or process of documents relative to one patient. XDR requires continuous networking capability between the healthcare providers exchanging data.

E.8 XDR Integration with XDS, Content Integration Profiles, PIX, and DSG

E.8.1 XDR Integration with XDS

The XDR Profile and the XDS Profile are both similar and complementary.

7095 Both profiles are document content neutral, conveying documents without modification and managing the same set of metadata. They both enable a Document Source to transmit a set of documents to another IHE Actor, using a HTTP based on-line mode or a SMTP based off-line mode.

7100 But they do differ in some important ways. XDS is a centralized profile with "servers" (Registry and possibly Repository) and "clients" (Source and Consumer). XDR is more symmetrical (Source and Recipient).

If the Document Source and the Document Recipient belong to the same Affinity Domain, the metadata shall respect the rules defined for this Affinity Domain (patientId, assigning Authority, encoding schemes...).

7105 In case the actors are not all part of the same Affinity Domain, the following options should be considered:

- If the Document Source is sending the Document Set to a Document Repository while sending it to the Document Recipient(s) at the same "time" (consequent network on-line messages), then the rules defined in the Affinity Domain which includes the Document Source and the Document Repository shall be used.
- In the other cases:
- If the Document Source and the Document Recipient(s) have agreed on the rules to use (for example using a "regional patientId" accessible through a PIX compliant server), then these rules shall be used.

7115

Note: it is highly recommended to define such “mutual agreement”

- If not, the rules available on the Document Source side will be used, and so, the Document Recipient has to transpose the patientId and the codes to follow its local rules.

7120

The Cross-enterprise Document Reliable Interchange (XDR) Integration Profile may be used in conjunction with the XDS Integration Profile to provide both a cross-enterprise sharing capability and a targeted sending of a set of documents to one or more specific receivers. This is illustrated in the figure below, when the Document Source Actor supports both the XDS and the XDR Integration Profiles.

Flexible Infrastructure: Sharing, Reliable Interchange, and Media

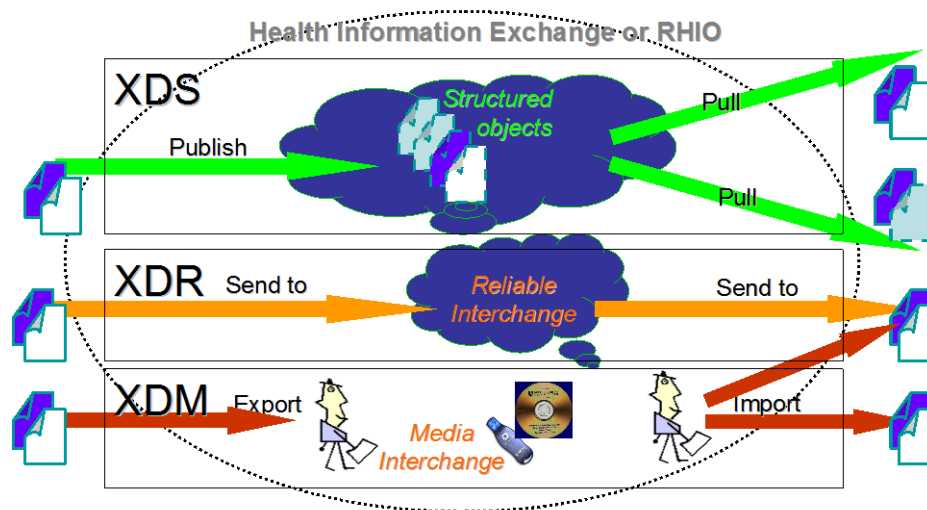


Figure E.6-1: Associated XDS, XDR, and XDM Profiles

7125

E.8.2 XDR Integration with XDS Content Integration Profiles

The Cross-enterprise Document Reliable Interchange (XDR) Integration Profile is intended to be used in conjunction with any number of XDS Content Integration Profiles to provide an interoperable specification for the content of the documents interchanged. Any XDS Content Integration Profile such as XDS-SD for scanned documents (See IT Infrastructure Technical Framework) or XDS-MS for medical summaries (See Patient Care Coordination Technical Framework) are examples of document content Integration Profiles that may be integrated along with XDR. One should note that although these Content Integration Profiles are called XDS-Scan or XDS-MS, the use of the XDS in their name does not imply that their use is restricted to XDS. It is equally intended for XDR, for point-to-point interchange.

7130

7135

E.8.3 XDR Integration with PIX

7140 The Cross-enterprise Document Reliable Interchange (XDR) Integration Profile may be used in
 conjunction with the PIX Integration Profile to provide the cross-referencing or linkage of the
 patient identifier used by the Document Source with that of the Document Receiver. This
 Integrated use requires the grouping of the XDR Document Source and of the XDR Document
 7145 Receiver with a PIX Patient Identity Source, so that the PIX Manager is fed with the patient
 identities in the Document Source and the Document Receiver identification domains. In
 addition, the Document Recipient shall be grouped with a PIX Patient Identifier Consumer
 Actor, so that when a patient Identifier is received in the XDS Document Metadata of the XDR
 Provide and Register Document Set transaction, it may invoke the services of the PIX Manager
 Actor to cross-reference the received patient identifier to a patient identifier of the Document
 Recipient Identification Domain.

E.9 XCA Integration with XDS and non-XDS communities

This section is informative and suggests some potential configurations that may be used by a
 community. The following types of community are described:

- 7150 • An XDS Affinity Domain
- A non-XDS Affinity Domain
- A collection of XDS Affinity Domains
- A collection of non-XDS Affinity Domains
- An XDS Affinity Domain with a “transparent” Gateway

7155 E.9.1 An XDS Affinity Domain

In the example below, the responding community is an XDS Affinity Domain which is served by
 a Responding Gateway.

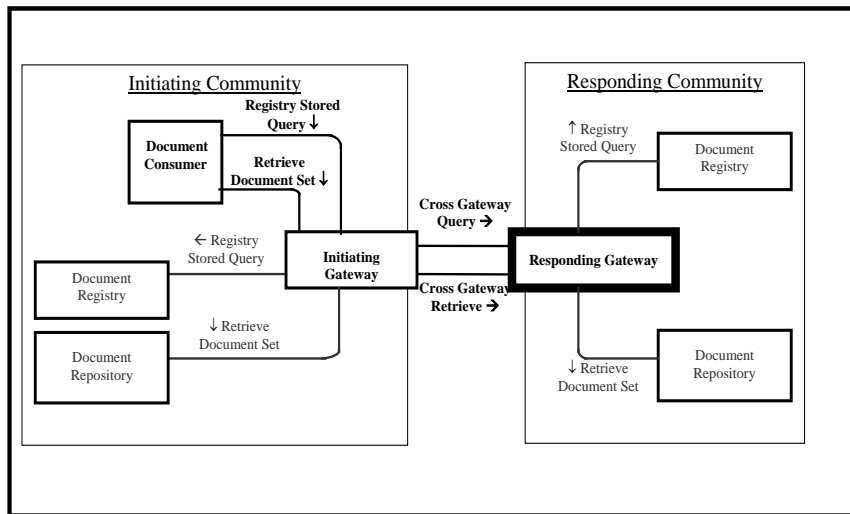


Figure E.9.1-1: XDS Affinity Domain

7160 E.9.2 A Non-XDS Affinity Domain

In the example below, the responding community is served by a Responding Gateway. However, within this community, there is no XDS Document Registry or Repositories. A proprietary mechanism is used by the Responding Gateway to gather data for the response to the Cross Gateway Query and Cross Gateway Retrieve transactions.

7165

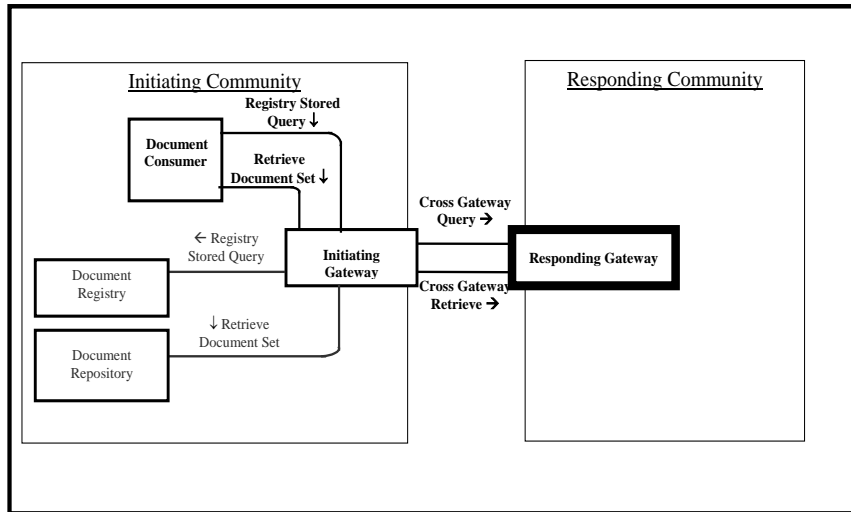


Figure E.9.2-1: Non-XDS Affinity Domain

E.9.3 A Collection of XDS Affinity Domains (informative)

7170 In the example below, one Responding Gateway is serving two communities. Each one of these communities is an XDS Affinity Domain served by its own Responding Gateway; these two Responding Gateways are hidden from the initiating community.

This example is informative only. The XCA Profile does not specifically support this configuration and does not address all the considerations of such a configuration.

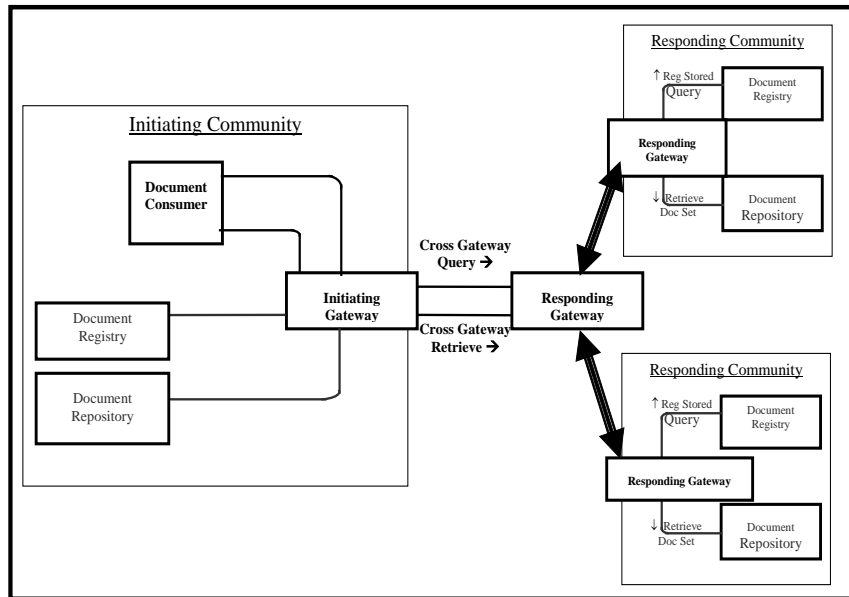


Figure E.9.3-1: Collection of XDS Affinity Domains

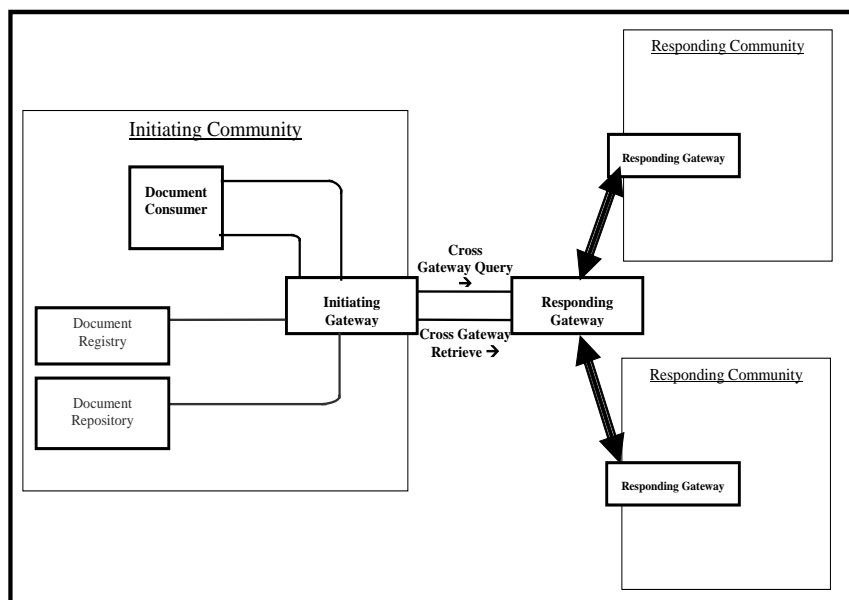
7175

E.9.4 A Collection of Non-XDS Affinity Domains (informative)

In the example below, one Responding Gateway is serving two communities. Each one of these communities is a non-XDS Affinity Domain served by its own Responding Gateway; these two Responding Gateways are hidden from the initiating Community.

7180

This example is informative only. This profile does not specifically support this configuration and does not address all the considerations of such a configuration.



7185

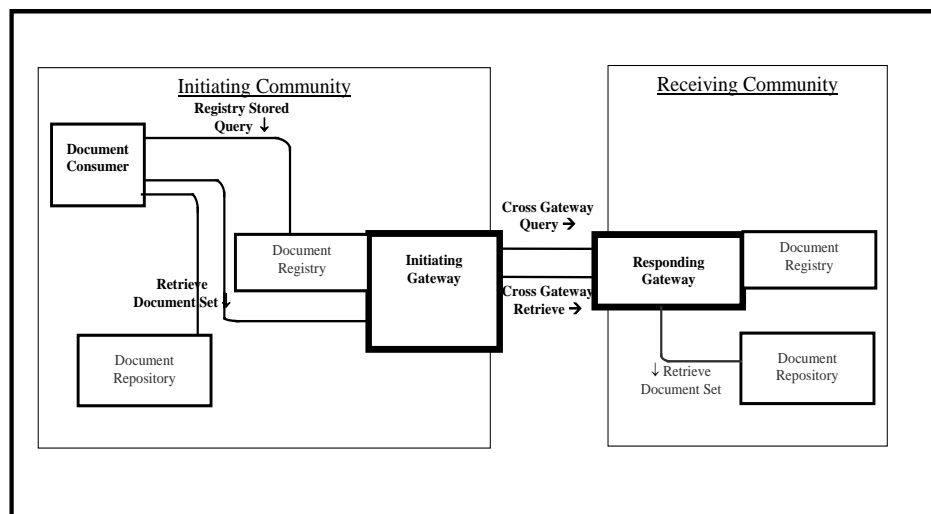
Figure E.9.4-1: Collection of Non-XDS Affinity Domains

E.9.5 An XDS Affinity Domain with a “Transparent” XCA Gateway (informative)

7190 In the example below, the initiating community is an XDS Affinity Domain where the Initiating
 7195 Gateway is grouped with the XDS Affinity Domain Document Registry. Thus the Document
 Consumer interacts with one system to retrieve both local and non-local data. This is called
 “transparent Gateway” as the Document Consumers do not see the cross-domain communication
 explicitly, but it is hidden by the Domain Registry and a Proxy Repository. Configuration would
 be needed to instruct the Document Consumer to interact with the Initiating Gateway when a
 non-local repository identifier was found in the metadata. In this way the Document Consumer
 interacts with the Initiating Gateway as a Proxy repository.

This diagram also shows a Responding Gateway grouped with a Document Registry.

This example is informative only. This profile does not specifically support this configuration
 and does not address all the considerations of such a configuration



7200 Figure E.9.5-1: An XDS Affinity Domain with a “Transparent” XCA Gateway

E.10 XCA and Patient Identification Management

7205 This section describes two models for resolving the patient identity in a cross-community
 exchange environment. As the XCA Profile is not intended to address patient identification
 management, it is therefore necessary to combine XCA with appropriate identification
 management Integration Profiles. This section is informative and describes only two possible
 ways to resolve patient identification relying on the existing two IHE Integration Profiles in this
 domain, Patient Identifier Cross-Referencing (PIX) and Patient Demographics Query (PDQ).
 7210 The description in this section is only at a high level and more details (not covered here) are
 necessary for implementation of these models. Other models for patient identification exist and
 will not be described in this section. Future work by the IHE IT Infrastructure Technical
 committee may support more sophisticated approaches.

E.10.1 Patient Identification using PIX

7215 The following diagram describes a mechanism for managing patient identities where there is
topmost PIX which cross references between communities A, B and C. This diagram assumes
that a Responding or Initiating Gateway for each community interacts in order to drive a patient
identity feed to the topmost PIX. The diagram does not include processing on the remote
communities (B and C) to respond to the query request. The topmost PIX is not defined in this
example, but can be assumed to be a PIX Manager, or equivalent, which is accessible to all
7220 communities.

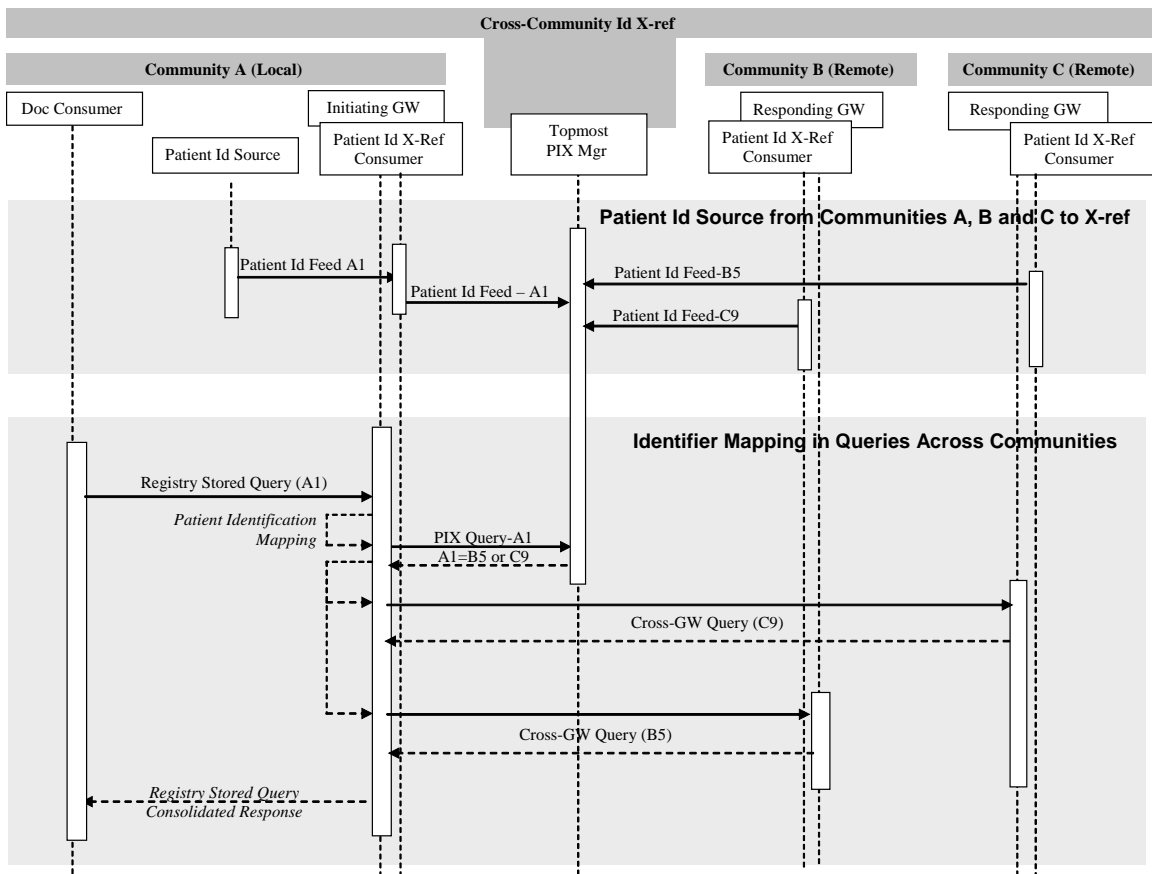


Figure E.10.1-1: Patient Identification using PIX

E.10.2 Patient Identification using PDQ

7225 The following diagram describes one approach to patient identification in a cross-community
exchange where there is no entity which can cross reference between local and remote
identifiers. Note that interactions among entities in remote communities (B & C) are not detailed
in this diagram.

7230

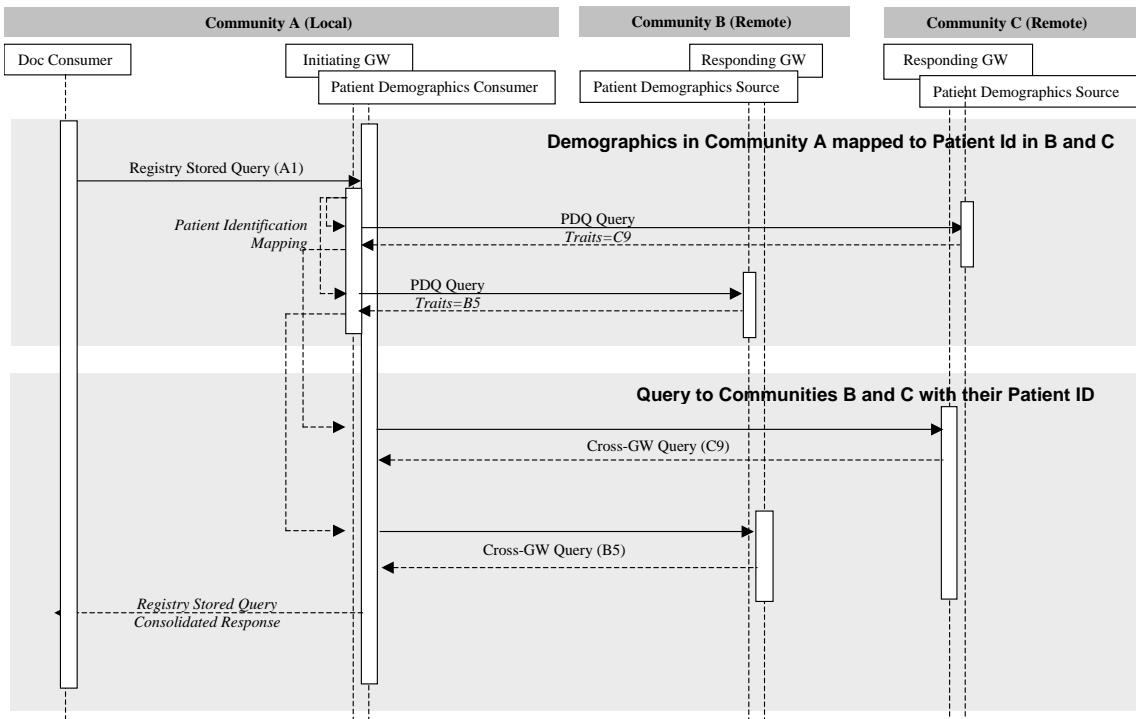


Figure E.10.2-1: Patient Identification using PDQ

7235 This diagram present a basic approach relying on the existing IHE Patient Demographics Query
 (PDQ) Integration Profile by the Initiating and Responding Gateways, where the Responding
 Gateway respond to queries with patient demographics traits for potential patients in the
 community it serves, thus allowing Initiating Gateways to obtain the patient Id to use in the
 Cross Gateway Query. The result of this transaction would be a) zero, indicating the patient does
 not have records at that community b) one, indicating the gateway was able to uniquely identify
 the patient. In the case of a) or b) the transaction is complete and does not require human
 7240 intervention. If multiple results are returned this requires human intervention to resolve.

7245 This approach requires a significant number of policy decisions to be in place, coordinated with
 privacy consent in cross-community environment that are well beyond the scope of the combined
 use of PDQ and XCA presented in Figure E.10.2-1. In addition, the integration of a large number
 of communities with a large number of non-overlapping patient populations is likely to require
 addressing significant scaling issues in allowing Responding Gateways to process the requests
 for identity resolution.

Future IHE work in this area may offer more sophisticated integration profiles that could be
 combined with XCA.

7250 **Appendix F: Request to Standards Development Organizations**

This Appendix is blank.

Appendix G: Security Considerations

G.1 Cross Profile Considerations

7255 IHE compliant systems usually process private healthcare information. This is subject to national privacy regulations, and possibly other state and contractual requirements. The IHE Infrastructure profiles do not fully define the security mechanisms necessary to protect this information. The Enterprise User Authentication Profile provides one component of this solution.

IHE assumes that actors will be installed on nodes with the following characteristics:

- 7260 • Each node has a security policy and procedure that applies to its operation. This is assumed to be part of the healthcare enterprise security policy.
- Any user (human, or application process) external to the node boundaries is submitted to an access control procedure in which the user/application will be authenticated.
- All required audit trail events are captured and recorded.

7265 The profiles in this framework assume the following environment:

- Physical Security Environment
 - 7270 • The equipment is assumed to be located in a physically protected and actively monitored area. This is normally the case with modality equipment because of other patient safety, privacy, and operational concerns. Similarly, the HIS systems and various archives are normally protected. Equipment like PACS workstations is sometimes placed in unprotected areas, but it is usually located where hospital staff monitors and limit access. It assumes that the threat of equipment modification is protected against by means of the physical security mechanisms.
 - 7275 • The network equipment that connects the computers is also assumed to be physically protected against unauthorized connections and unauthorized modifications. In the treatment areas of most hospitals the network equipment is in ceilings, cableways, locked cabinets, and other protected areas. There is usually staff present to monitor that no unauthorized activity is taking place.
 - 7280 • Local procedures and operations will be in place to ensure that the physical security assumptions are valid for other areas of the hospital, such as administrative offices, that may be at greater risk.
 - 7285 • Remote locations, especially home offices, are not physically protected. Other means will be used to provide equivalent protection. This may include the use of technology such as VPN connections or HTTPS encryption. Use of encryption or VPN is not a complete replacement for physical security but may be part of an overall protection system.

- The home computer that is used for both personal and professional purposes is difficult to protect. It will be protected from inadvertent modification by malicious software or its use will be prohibited.

7290

- Network Security Environment
 - In addition to the physical security of the network, there will be protection against network access by unsupervised systems. This is typically provided by mechanisms such as firewalls and VPNs.

The threat profile is assumed to be:

7295

- Accidental and inadvertent misuse
- Individual abuse for personal gain, malice, revenge, or curiosity. The abusers are assumed to have only limited access to the underlying systems and software. They are not expert at the internal structure of the systems.
- Random untargeted abuse, such as from an Internet hacker.

7300

The threat profile also assumes that the following threats are either not present or otherwise protected.

- Individual abuse by a system administrator, system developer, or other expert.
- Military or hostile government action
- Organized criminal attack

7305

IHE addresses only those security requirements related to IT systems within the scope of IHE healthcare applications. It does not address security requirements for defending against network attacks, virus infection, etc.

7310

IHE does not mandate the use of encryption because the performance impact of current encryption algorithms is excessive. Most hospital networks provide adequate security through physical and procedural mechanisms. The additional performance penalty for encryption is not justified for these networks. The profiles permit the use of encryption so that it can be used as part of an overall security plan.

Appendix H: Intentionally Left Blank

7315 **Appendix I: Intentionally Left Blank**

Appendix J: Content and Format of XDS Documents

7320 The XDS Integration Profile purposely leaves a number of policies up to the XDS Affinity Domain to decide, including the structure and format of the content of XDS Documents to be shared, the mapping of content metadata into the XDS Document Registry, the coding of XDS Document metadata, the events that trigger an XDS Submission Request, and the policies concerning the use of XDS Folders to facilitate sharing.

7325 It is important to recognize that until sufficient experience has been gained in cross-enterprise document sharing, it is not possible to establish common or even best practices in the use of the XDS Integration Profile. IHE has therefore chosen to abstain to make recommendations in these topics at this time.

7330 IHE also recognizes that there will be a need for content-oriented integration profiles to be used in cooperation with this Integration Profile. It is expected that in the future the various IHE Domains (Patient Care Coordination, Cardiology, Laboratory, Radiology, IT Infrastructure, etc.) will produce IHE Integration Profiles refining the use of XDS within the domain. These various content-oriented integration profiles may rely on XDS, but would further constrain the forms of documents to be shared, or the uses of XDS features such as Folders and Submission Sets, et cetera.

Content Neutrality

7335 XDS is content neutral. It neither prescribes nor prohibits the format, content, structure or representation of documents that can be retrieved from an XDS Document Repository. For the XDS Integration Profile to have immediate value to an XDS Affinity Domain, it must be able to adapt to the documents that are present and available from its members. Thus, prohibitions on content would only serve to limit the utility and adoption of the XDS Integration Profile. Similarly, XDS Affinity Domains must be able to adapt to emerging standards, which cannot be
7340 enumerated in any list of prescribed content formats.

IHE strongly recommends that XDS Affinity Domains adopt rules that require documents to comply with widely accepted standards where possible (e.g., HL7 CDA, CEN ENV 13606, ASTM CCR, and DICOM Composite Object).

Document Headers and Metadata

7345 Because XDS is content neutral, XDS cannot validate metadata contained within the body of an XDS document against the metadata supplied to the XDS Document Registry. XDS Affinity shall therefore select content where IHE has defined Integration Profiles, or until that point, the XDS Affinity Domains shall carefully define how the attributes in the XDS Document Registry are filled.

7350 Metadata and the Patient Record

Although metadata in the document header may be duplicated in the XDS Document Registry, the XDS Document Registry metadata has a particular role in term of being part of the legal medical record stored. It is definitively not part of the clinical record as managed by the XDS Document Repositories where documents reside. Furthermore, XDS does not provide for

- 7355 transactions to “sign” or legally authenticate the content of an XDS Submission Set (See IHE Document Digital Signature Content Profile- DSG), although it offers the ability to track its author, if the XDS Affinity Domain so desires to enforce it. The contents of XDS Folders are tracked, through the Submission Sets that contributed to placing document references in folders. However, the existence of document metadata in the registry and the potential medical acts
- 7360 involved in creating an XDS Submission Set or XDS Folder may make the contents of the XDS Document Registry part of the patient’s legal medical record. It will be up to individual XDS Affinity Domains to decide how to address the issues involved with these clinical acts and to resolve them in accord with common sense, acceptable medical practices, and local regulations.

7365 **Appendix K: XDS Concept Details**

K.1 XDS Document Concept

An XDS Document is the smallest unit of information that may be provided to a Document Repository Actor and be registered as an entry in the Document Registry Actor.

7370 An XDS Document is a composition of clinical information that contains observations and services for the purpose of exchange with the following characteristics: Persistence, Stewardship, Potential for Authentication, and Wholeness. These characteristics are defined in the HL7 Clinical Document Architecture Release 1 specification.

7375 An XDS Document may be human and/or application readable. In either cases, it shall comply with a published standard defining its structure, content and encoding. IHE intends to define content-oriented Integration Profiles relying on such content standards to be used in conjunction with XDS.

Furthermore:

1. When submitted for sharing, an XDS Document shall be provided to the Document Repository Actor as an octet stream with an associated MIME type.
- 7380 2. When retrieved through the Retrieve Document transaction, an XDS Document shall be unchanged from the octet stream that was submitted (full fidelity repository).

Note: An XDS Document may be a MIME multipart document (e.g., an HL7 CDA as its first part followed by attachments as files). The first part of the multi-part contains the primary part of the document; other parts are direct attachments to the primary part. The Document Repository handles this multi-part data set as an “opaque entity”. The Document Repository does not need to analyze or process its multi-part structure nor the content of any parts in the context of the XDS Integration Profile.

Note: An XDS Document may be retrieved using alternate methods using document specific retrieval methods. Such optional capabilities are not provided in the current specification of XDS, but are possibly candidates for addition as future options this Integration Profile.
- 7390 3. An XDS Document shall be associated with metadata defined by the Document Source. This metadata information shall be placed by the XDS Registry Actor in an XDS Document Entry, and is used for query purposes by XDS Consumer Actors.
- 7395 4. The XDS Integration Profile manages XDS Documents as a single unit of information, it does not provide mechanisms to access portions of an XDS Document. Only the Document Sources or Document Consumers have access to the internal information of the XDS Document.
- 7400 5. An XDS Document is globally uniquely identified, so that no two XDS Documents with different content shall bear the same Unique Identifier. This identifier is unique across all XDS Affinity Domains, which allows potential merger of XDS Document Repositories from different domains, or exchange of XDS Documents between Clinical Affinity Domains, if so desired.
6. The XDS Document Registry Actor shall maintain a single document entry for each XDS Document stored in a Document Repository Actor. Duplicate copies of the same XDS

- 7405 Document (with the same unique identifier) may be stored and registered. Registration of an XDS Document with the same unique identifier but a different content is rejected.
- 7410 7. This Integration Profile specifies the metadata required for each XDS document registered in the Document Registry. It is the responsibility of the Document Source to ensure that the XDS Document metadata reflects the actual content of the associated XDS Document. Neither the Document Repository nor the Document Registry checks this consistency.
- 7415 8. The Document Source maintains the following responsibilities over the XDS Documents it has registered:
- a. It has rights to change the status of any of these Documents from “approved” to “deprecated” or to delete them outright.
 - b. It has rights to submit an XDS Document with a “Parent Relationship” of replacement (“RPLC”) for one of its previously submitted document¹¹.

7420 XDS Affinity Domains should have policies and procedures to provide patient access to these operations where necessary. For example, in certain regions, patients may request the removal of documents from the EHR-LR. The Registry and Repositories implementations should be ready to support these local operations although there are no IHE transactions defined at this time.

K.2 Concept of an XDS Affinity Domain

An XDS Affinity Domain is made of a well-defined set of Document Repositories and Document Consumers that have agreed to share the clinical documents. An XDS Affinity Domain has a number of properties defined:

- 7425 1. An XDS Affinity Domain does not deliver care. Only the EHR-CRs belonging to an XDS Affinity Domain as Document Sources and Consumers do.
- 7430 2. An XDS Affinity Domain is managed by a single Document Registry Actor.
- Note: A distributed registry approach will be considered as a future and separate Integration Profile. For Document Source and Document Consumer Actors, the perception of a single Document Registry Actor hides the complexity of a distributed registry.
- 7435 3. It includes any number of Document Repository Actors (a distributed configuration is the default, however, a centralized configuration with a grouped Registry/Repository is also supported).
4. It contains an explicit list of Document Consumer and Document Repository actors that participate in document sharing. The addition of a Document Repository or Document Consumer Actor is an administrative task that requires involvement of authorities maintaining the Registry and Repositories.

¹¹ For example, in DICOM, where the document identity does not change even though its internal patient metadata may have been updated, the Document Source would submit an updated DICOM Document as a replacement for the existing one.

5. There is a chain of trust established between the users (healthcare staff) in each EHR-CR and the XDS Affinity Domain.
- 7440 6. Document Repositories and Document Consumers may belong to more than one XDS Affinity Domain and share the same or different documents. This is an implementation strategy and will not be further described.
- 7445 7. The XDS Affinity Domain supports a primary Patient Identification Domain that is used by the Document Source and Consumers to communicate with the Document Registry. When Document Sources and Consumers in the XDS Affinity Domain belong to different Patient Identifier Registration Domains, the Document Source and Consumers must cross-reference their own Patient Identifier Registration Domains to that of the Registry. They may use the IHE Patient Identifier Cross-referencing Integration Profile, the IHE Patient Demographics Query Integration Profile or other XDS Affinity Domain specific mechanisms for cross-referencing (See ITI TF-1: E.3 and E.5).
- 7450 8. A Document Source may only contribute documents with Document Codes and Health Facility Codes that draw from a Vocabulary Value Set that is approved by the XDS Affinity Domain.

K.3 Other Principles of XDS

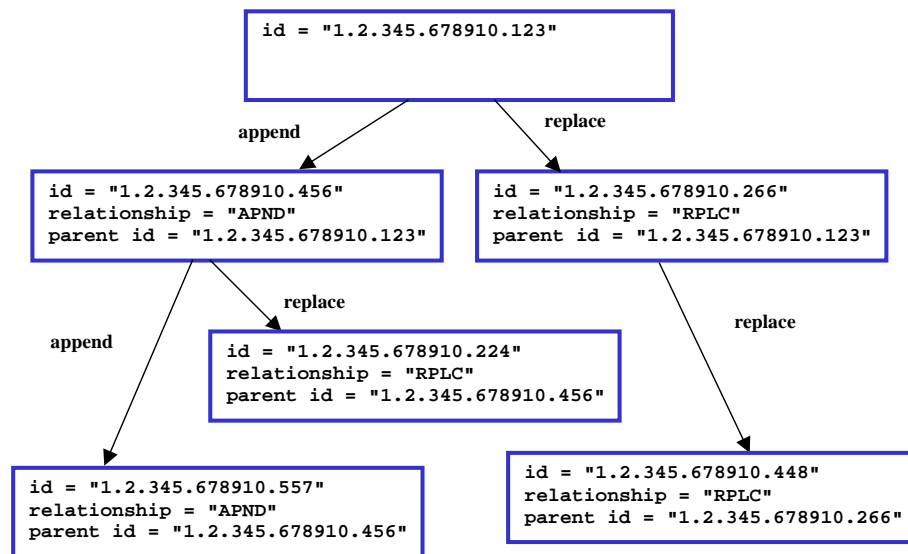
- 7455 The XDS Integration Profile has been designed with the following limitations and principles:
1. A Document may contain references to other documents in its content which are not under the management of the XDS Document Registry. Such references may be available to the EHR-CR that registered the document that includes the reference. It is beyond the scope of XDS to provide access to such documents internal to the EHR-CR.
 - 7460 2. The XDS Repositories are not expected to perform any processing or translations on document content. Processing and translation are the responsibility of a Source EHR-CR or Consumer EHR-CR. The analysis, cross-document combination and presentation of document content are outside the scope of the XDS Integration Profile and its actors.
 - 7465 3. The custodianship for the clinical information contained in a registered document remains with the Source Actor of the EHR-CR. The EHR-LR offers only a “shared space” under the responsibility of each contributing EHR-CR. Through XDS, replacement or deletion of documents in the EHR-LR may only be initiated by the corresponding EHR-CR Source.
 - 7470 4. When an XDS Document that has already been registered in the XDS Registry of an XDS Affinity Domain is resubmitted as if it was a new XDS Document with the same Document Unique identifier, this “duplicate submission” is detected by the Repository and/or Registry based on the fact that the XDS Document Unique Identifier already exists in a Document Entry. The submission request to which that resubmitted Document belongs shall be rejected in the case where the identifiers match but the actual content differs (detected by use of a hash key computed by the Document Repository at the time of submission).
 - 7475

K.4 Document Identification

7480 In order to reduce the number of unique identifiers associated with an XDS Document, the globally unique Document Id assigned by the document source and the unique XDS Document Id used by the Repository are the same. It is strongly recommended to limit the use of the Document Entry UUID created per ebRS in order to reference the document entry for referencing internally to the encoding of the IHE transactions operations, and to encourage the use of the globally unique Document Id for all external operations (e.g., links maintained in data bases internal to the Document source Actor, links within documents, etc.).

7485 The XDS Document Entry includes two separate attributes: an XDSDocument.uniqueId and XDSDocument.repositoryUniqueId. The Document Unique ID is a location independent identifier. As the result of XDS Document migration from one XDS Document Repository to another one within an XDS Affinity Domain, the repositoryUniqueId would be changed, but not the Document unique ID.

7490 K.5 Example of Document Relationship



Adapted from HL7 CDA Release 2

Figure K.5-1: Example of Document Relationships

7495 These relationships are illustrated in the above figure. Typical scenarios are a simple replacement (e.g., XDSDocument.id "1.2.345.678910.266" replacing XDSDocument.id "1.2.345.678910.123") and a simple addendum (e.g., XDSDocument.id "1.2.345.678910.456" appends XDSDocument.id "1.2.345.678910.123"). More complex scenarios that might be anticipated include:

- 7500
1. Replacement of an addendum (e.g., XDSDocument.id "1.2.345.678910.224" replaces XDSDocument.id "1.2.345.678910.456", which itself is an addendum to XDSDocument.id "1.2.345.678910.123") - expected behavior would be to render the

replacement as the addendum (e.g., render XDSDocument.id "1.2.345.678910.224" as the addendum to XDSDocument.id "1.2.345.678910.123");

7505

2. Addendum to a replaced document (e.g., XDSDocument.id "1.2.345.678910.456" appends XDSDocument.id "1.2.345.678910.123", which has been replaced by XDSDocument.id "1.2.345.678910.266") - expected behavior would be to render the addendum along with the replacement (e.g., render XDSDocument.id "1.2.345.678910.456" as an addendum to XDSDocument.id "1.2.345.678910.266").

Appendix L: XDS Affinity Domain Definition Checklist

7510 The concept of an XDS Affinity Domain is defined in ITI TF-1:10 and ITI TF-1: Appendix K. ITI TF-1: Appendix L originally provided an informative checklist for the key policies that need to be addressed in order to deploy an EHR-LR document sharing environment for an XDS Affinity Domain. However, it was recognized that this checklist was incomplete as it did not deal with many necessary XDS Affinity Domain deployment issues. In order to address these

7515 shortcomings, a new “Template for XDS Affinity Domain Deployment Planning” White Paper has been created: http://www.ihe.net/Resources/Technical_Frameworks/#iti

It takes the form of a template rather than a checklist because it acts more as an outline for all the issues that should be considered, rather than a checklist to be used to verify the correctness of a particular implementation. This new template can be used when defining policies for either an

7520 individual XDS Affinity Domain, or multiple XDS Affinity Domains within a particular nation or region.

Here is a summary of the topics defined in the new “Template for XDS Affinity Domain Deployment Planning”:

- Organizational Rules
- 7525 • Structure, Roles, Transparency, Legal Considerations and Enforcement
- Operational Rules
 - Service Level Agreements, Daily Governance, Configuration Management, Data Retention, Archive, and Backup, and Disaster Recovery
- Membership Rules
- 7530 • Acceptance, Types of Membership, Membership Policies
- Connectivity to the XDS Affinity Domain from External Systems
 - System Architecture
 - Global Architecture, Affinity Domain Actors, Transaction Support
 - Terminology and Content
- 7535 • Refinement of Metadata and Content Attribute Use
 - Patient Privacy and Consent
 - Access and Use, Patient consent, and Override Guidelines
 - Technical Security
- 7540 • Authorization, Role Management, User/Role Authentication, Node Authentication, Certificates Management, Information Access Security, Information Integrity, Updates, and Maintenance Policies, Secure Audit Trails, Consistent Time, Audit Checks, and Risk analysis

Appendix M: Cross-Enterprise Document Sharing and IHE Roadmap

7545 The IHE Cross-Enterprise Document Sharing Integration Profile is part of a family of IHE Integration Profiles grouped in a number of domain-specific Technical Frameworks Patient Care Coordination, Cardiology, Laboratory, Radiology, IT Infrastructure, etc.). XDS is a central foundation for Cross-Enterprise interoperability that may be combined with a number of the existing IHE Integration Profiles (See ITI TF-1: Appendix E). However a number of new IHE Integration Profiles need to be developed, pending the availability of the relevant base standards.

7550 M.1 Document Content Integration Profiles for XDS

It is expected that the various IHE Domains (Cardiology, Laboratory, Radiology, IT Infrastructure, etc.) will produce new IHE Integration Profiles addressing the content of the documents that need to be shared. These various “content-oriented” Integration Profiles will rely on the XDS Integration Profile for managing the registration, discovery and access processes in a common manner.

7555

Such an effort is underway with the IHE Patient Care Coordination Domain for medical summaries used in referrals and discharge summaries and other document types. See www.ihe.net.

M.2 Cross-Enterprise Dynamic Information Sharing

7560 The management of dynamic information (non-document-oriented) such as allergy lists, medication lists, problem lists, etc. is not addressed by XDS. However, a means to access this information in a structured form and to manage updates to such dynamic clinical information is a candidate for a specific Integration Profile.

M.3 Collaborative Workflow Process Management

7565 There is a wide array of shared care delivery collaborative processes such as the placing and tracking of orders (e.g., drug prescriptions, radiology orders, etc.) for which XDS provides only a partial solution (the creation of the patient record with the resulting persistent artifacts). XDS offers a critical infrastructure for ePrescribing and eReferral in that it can ensure that the various providers share access to orders, prescriptions, dispensations, and results. The means to interoperate on the command/control part of these collaborative workflow processes is a candidate for specific Integration Profiles in the future.

7570

M.4 Security and Privacy Management

7575 The operation of any XDS Affinity Domain will require that a proper security model be put in place. It is expected that a range of security models should be possible. Although the XDS Integration Profile is not intended to include nor require any specific security model, it is required that XDS implementers group XDS Actors with actors from the IHE Audit Trail and Node Authentication and will need an Access Control capability that operates in such a cross-enterprise environment. Specific IHE Integration Profiles complementary to XDS are available (e.g., Cross-Enterprise User Authentication, Document Digital Signature, etc.).

7580 **M.5 Federation of XDS Affinity Domains**

7585 XDS is an effective means to establish XDS Affinity Domains that include care delivery organizations at any level, local, regional or national. However, the establishment of independent but consistently XDS Affinity Domains will call for their federation, as patients expect their records to follow them as they move from region to region, or country to country. IHE foresees a need for transferring information from one XDS Affinity Domain to another, or to allow access from one XDS Affinity Domain to documents managed in other XDS Affinity Domains. XDS has been designed with this extension in mind. The Cross-Community Access (XCA) Integration Profile that complements XDS provides this function.

Appendix N: Intentionally Left Blank

7590 **Appendix O: Intentionally Left Blank**

Appendix P: Privacy Access Policies (Informative)

This Appendix provides information about when consent could be automated and consequently when BPPC could be used. Privacy consent can be summarized as: "I agree on my personal data being disclosed to someone under specific conditions".

7595 Conditions are based on various factor(s) for example:

- type of person the data is disclosed to;
- type of data disclosed;
- type of access (normal access, emergency access...);
- security level in which the disclosure takes place (weak authentication vs. strong authentication);
- type of purpose for which the data is disclosed;
- timeframe (period of validity of the consent, window of disclosure...);

7600

BPPC could be used when conditions can be described with a limited number of factors and when the factors can be defined and be easily interpreted by a Document Consumer implementing the Basic Patient Privacy Enforcement Option.

7605

The XDS Affinity Domain Privacy Consent Policies could result in various actions, for example:

- limitation of the display of the existence of specific documents to the users of a Document Consumer
- limitation of the access to specific documents by the users of a Document Consumer
- display of a warning note (either concerning this access or to inform that further disclosure is not allowed, limited to some defined population, needed further consent...)
- collection of new consent (oral consent, patient authentication, electronically signed consent, paper consent...)

7610

7615 **P.1 Consents in a sensitivity labeled and role based access control environment**

One possible implementation may have a collection of policies and sensitivity markers that would form an access control matrix. An example simple access control matrix is shown in the table below.

7620

Table P-1: Sample Access Control Policies

Sensitivity	Billing Information	Administrative Information	Dietary Restrictions	General Clinical Information	Sensitive Clinical Information	Research Information	Mediated by Direct Care Provider
Functional Role							
Administrative Staff	X	X					
Dietary Staff		X	X				
General Care Provider		X	X	X			
Direct Care Provider		X	X	X	X		X
Emergency Care Provider		X	X	X	X		X
Researcher						X	
Patient or Legal Representative	X	X	X	X	X		

Each instance of the matrix results in a single Patient Privacy Policy. This vocabulary must then be configured in the XDS Affinity Domain. Thus configuring each application in the XDS Affinity Domain to recognize for each Patient Privacy Policy identified, and which sensitivity (confidentialityCode); what types of accesses are allowed. Using the example above, the Patient Privacy Policy might look like.

7625

Table P-2: Patient Privacy Policies When Expressed by Document Sensitivity

Privacy Consent Policy	Description
Billing Information	May be accessed by administrative staff and the patient or their legal representative.
Administrative Information	May be accessed by administrative or dietary staff or general, direct or emergency care providers, the patient or their legal representative.
Dietary Restrictions	May be accessed by dietary staff, general, direct or emergency care providers, the patient or their legal representative.
General Clinical Information	May be accessed by general, direct or emergency care providers, the patient or their legal representative.
Sensitive Information	May be accessed by direct or emergency care providers, the patient or their legal representative.
Research Information	May be accessed by researchers.

Privacy Consent Policy	Description
Mediated by Direct Care Provider	May be accessed by direct or emergency care providers.

7630 Other divisions of the access control matrix are possible, so long as a Patient Privacy Policy covers each layout of the matrix.

The following list of references is provided as good references to understand the terms and concepts presented here. These references are not required by this profile.

- ISO/TS 21298 "Health informatics – Functional and structural roles".
- 7635 • ISO/TS 22600 "Health Informatics – Privilege Management and Access Controls".
- CEN prEN 13606-4 "Health informatics — Electronic health record communication — Part 4: Security requirements and distribution rules"

P.3 Possible checklist for implementations

General (before anything else)

- 7640 • Granularity of confidentiality implementation:
 - Granularity of document: all documents, document type, each document.
 - Granularity of user: all users, user type, each type.
- Depth of confidentiality implementation:
 - 7645 • Is the existence (metadata) about a document that can't be read by the user shown in a list of available documents for this patient?
 - Is the user informed there are / might be not shown documents and how much?
 - Is there the possibility to manage different depth of confidentiality depending on users or document type?
- How to identify users, documents and policy?
- 7650 • Does confidentiality management spread through further use (once the document is downloaded by a user)

While implementing

- Definition of default codes depending on site / hardware, document type, author, patient...
- Implementing options:
 - 7655 • possibility of a list to choose from and how the list is constituted (out of all the possible value, out of the value acknowledged by patient...)
 - possibility to change default codes prior to publication

- possibility to use different format depending on the confidentiality policy (only non-downloadable image, pdf, word...)
- 7660
- Later modification of policy (possible directly when requesting a document or have to be validated before)

Prior to publication

- What elements should be checked before publication:
 - existence of a policy
- 7665
- existence of the policy used
 - existence of a consent for that policy
 - What additional information should be given (general consent policy, patient's specific consent policy...?)

Prior to allowing access to a document

- 7670
- What elements should be checked before publication:
 - accessing user role
 - existence of the policy used vs. accessing user
 - Specific accesses and impact on confidentiality policy:
 - emergency (specific policy, short cut of confidentiality policy...)
- 7675
- break glass
 - What additional information should be given (general consent policy, patient' specific consent policy...)

P.4 Potential obligations

7680 Possible things that the BPPC policies might include are not fully known at this time. The following is a list that has been discovered through use by researchers, health information exchanges, and vendors. The following are some thoughts of things that might be orchestrated by BPPC Policies.

General

- 7685
1. Is the existence (metadata) about a document that can't be read by the user shown in a list of available documents for this patient
 2. Map local role codes into some Affinity Domain defined role codes

Prior to implementation

- 7690
3. the specific Document Source is configured with one site specific “normal” code to publish all of that Document Source documents against. For example an automatic blood-pressure device being used by one specific patient.

4. prompt user for the code to apply to the document (drop-down-list)
5. document-type based codes

Prior to publication

6. validate that the code to be published against has been acknowledged
- 7695 7. support for a XDS Affinity Domain Patient Privacy Policy that forbids the publication and/or use of documents in the XDS Affinity Domain (aka Opt-Out).

Prior to allowing access to a document

8. should documents with unrecognized codes be shown?
9. prompt the user with some site defined text "do you really want to do this?"
- 7700 10. allow the user to review the base consent policy
11. allow the user to review the patient's specific Patient Privacy Policy Acknowledgement Documents
12. allow the user to override a consent block (break-glass)
- 7705 13. require that a new consent be acquired from the patient before using the documents in the XDS Affinity Domain
14. support for a XDS Affinity Domain Patient Privacy Policy that forbids the publication and/or use of documents in the XDS Affinity Domain (aka Opt-Out).
15. validate that the code on the document has been acknowledged
- 7710 16. confidentialityCode that would indicate that the Document can only be viewed, it cannot be incorporated or copied.
17. use of this document shall result in an ATNA emergency access audit event

P.5 Dynamic Use Models

- It has also been suggested that documents should simply be published with the expected codes, and that only on use of a document that ALL current Patient Privacy Policy Acknowledgements are evaluated against with the code on the document. In this way revocation is more dynamic.
- 7715

GLOSSARY

- 7720 **Actor:** An entity within a use case diagram that can perform an action within a use case diagram. Possible actions are creation or consumption of a message
- ADT:** Admit, Discharge & Transfer.
- Care Delivery Organization:** A Care Delivery Organization refers to a broad variety of healthcare facilities: private practice, nursing home, ambulatory clinic, acute care in-patient facility, hospitals etc.
- 7725 **CCOW:** ANSI certified technology neutral specification for the Health Level Seven Context Management Architecture (CMA). This architecture enables multiple applications to be automatically coordinated and synchronized in clinically meaningful ways at the point of use. The architecture specified in this document establishes the basis for bringing interoperability among healthcare applications to point-of-use devices, such as a personal computer that serves as
- 7730 a clinical desktop.
- Community:** A community is defined as a group of facilities/enterprises that have agreed to work together using a common set of policies for the purpose of sharing health information via an established mechanism. Facilities/enterprises may host any type of healthcare application such as EHR, PHR, etc. A community is identifiable by a globally unique id called the
- 7735 homeCommunityId. Membership of a facility/enterprise in one community does not preclude it from being a member in another community. Such communities may be XDS Affinity Domains which define document sharing using the XDS Profile or any other communities, no matter what their internal sharing structure.
- Confidentiality Code:** A value from the value-set that indicates the sensitivity and/or
- 7740 confidentiality of an object (e.g., Document). This may be from the HL7 defined vocabulary or extension defined in the Security/Privacy Domain. The Confidentiality Code is used during access control to indicate the type of object as viewed by the security and/or privacy policies.
- Context Management Registry:** An HTTP technology specific service defined by the HL7
- 7745 Context Management “CCOW” Standard to locate an instance of a context manager servicing a specific desktop.
- Context Session:** A collection of participant applications that are sharing context on one or more subjects.
- CDA:** Clinical Document Architecture (specified by HL7).
- CT:** Consistent Time Integration Profile.
- 7750 **XDS Affinity Domain:** A group of healthcare enterprises that have agreed to work together using a common set of policies and which share a common infrastructure of repositories and a registry.
- Directory:** A book containing the names and residences of the inhabitants of any place, or of classes of them; an address book; as, a business directory.

- 7755 **EHR-CR:** An EHR-CR or Care-delivery Record abstracts the patient information managed by the IT system or set of systems of a Care Delivery Organization, which may support a broad variety of healthcare facilities: private practice, nursing home, ambulatory clinic, acute care in-patient facility, etc.
- 7760 **EHR-LR:** The documents shared by the EHR-CR and tracked by the Registry form a Longitudinal Record for the patients that received care among the EHR-CRs of the XDS Affinity Domain. This is known as the EHR-LR.
- eMPI:** Enterprise Master Patient Index.
- Encounter:** An interaction between a patient and care provider(s) for the purpose of providing healthcare-related service(s). Healthcare services include health assessment.
- 7765 *Examples: outpatient visit to multiple departments, home health support (including physical therapy), inpatient hospital stay, emergency room visit, field visit (e.g., traffic accident), office visit, occupational therapy, telephone call.*
- EUA:** Enterprise User Authentication Integration Profile.
- 7770 **Expanded Value Set** – a set of concept representations that were in effect at a specific time for a particular version of a Value Set. See Value Set. The Value Set and the Expanded Value Set concepts are similar to the programming concepts of Class and Instance of Class. This may also be called a value set resolution or resolved value set.
- Expected Actions:** Actions which should occur as the result of a trigger event.
- Extensional Value Set** – a set of concepts that is specified in terms of a list of concepts.
- 7775 *(The definitions are based on HL7 Vocabulary terms).*
- Globally Unique Identifier (GUID):** An identifier of an entity, such as persistent document, that has been generated by an algorithm guaranteeing its global uniqueness. For more information please see ITI TF-2x: Appendix B discussion of OIDs and UUIDs.
- HIMSS:** Healthcare Information and Management Systems Society.
- 7780 **HIS:** Hospital Information System.
- homeCommunityId:** A globally unique identifier for a community. It is used in XCA to obtain the Web Services endpoint of services that provide access to data in that community.
- IETF:** Internet Engineering Task Force
- IHE:** Integrating the Healthcare Enterprise.
- 7785 **inetOrgPerson:** The inetOrgPerson [RFC 2798] object class is a general purpose object class that holds attributes about people. The attributes it holds were chosen to accommodate information requirements found in typical Internet and Intranet directory service deployments. The inetOrgPerson object class is designed to be used within directory services based on the LDAP v3 [RFC 2251] and the X.500 family of protocols, and it should be useful in other
- 7790 contexts as well.

Intensional Value Set – a set of concepts that is specified in terms of the “intension” of use, for example “all concepts that are children of this node in a tree of concepts”. Intensional value sets often have some kind of algorithmic basis for selection of concepts.

Interaction Diagram: A diagram that depicts data flow and sequencing of events.

7795 **IT:** Information Technology.

JPEG: Joint Photographic Experts Group.

KDC: Key Distribution Center (the Kerberos server that issues Ticket Granting Tickets and service tickets. See RFC1510).

7800 **LDAP:** Lightweight Directory Access Protocol is designed to provide access to directories supporting the X.500 models, while not incurring the resource requirements of the X.500 Directory Access Protocol (DAP). This protocol is specifically targeted at management applications and browser applications that provide read/write interactive access to directories. When used with a directory supporting the X.500 protocols, it is intended to be a complement to the X.500 DAP.

7805 **Local Authentication:** In the ATNA Profile the term “local authentication” means that the user identification, authentication, and authorization method is chosen by the local system administration and does not necessarily comply with any IHE profile. It may be a local username password system, a secure token system, or any other system that is considered acceptable by the local security administration.

7810 **Movement:** An event describing a change of the situation of the patient in the context of the encounter. This concept encompasses changes such as transfers of patient location, change of patient class, new attending doctor, new consulting doctor, new encounter starting, encounter closing, etc. The concept of Movement is a superset of the concept of “Transfer”.

MPI: Master Patient Index.

7815 **MRN:** Medicare Record Number.

NEMA: National Electrical Manufacturers Association.

NTP: Network Time Protocol. This is the standard Internet protocol for synchronizing computer clocks. The web site <http://www.ntp.org> provides extensive background documentation at the introductory and expert level on how to synchronize computers.

7820 **OID:** Object Identifier. (See also 'Globally Unique Identifier').). For more information please see ITI TF-2x: Appendix B discussion of OIDs and UUIDs.

PACS: Picture Archive and Communication System.

7825 **Patient:** (When used in the context of ATNA) RFC-3881 defines the means of identifying the person who is a patient. The patient information in audit event records corresponds to the information available to identify a patient at the time the audit record was generated, and does not reflect later updates (e.g., patient reconciliation).

PatientID: (When used in the context of ATNA) A free text that holds the system-internal patient identifier being unique within that system domain. The patient identifier domain is that

7830 assigned to the system that generated the audit event record. The patient information in audit event records corresponds to the information available to identify a patient at the time the audit record was generated, and does not reflect later updates (e.g., patient reconciliation).

7835 **Patient Identifier Cross-reference Domain:** Consists of a set of Patient Identifier Domains known and managed by a Patient Identifier Cross-reference Manager Actor. The Patient Identifier Cross-reference Manager Actor is responsible for providing lists of “alias” identifiers from different Patient Identifier Domains.

7840 **Patient Identifier Domain:** A single system or a set of interconnected systems that all share a common identification scheme for patients. Such a scheme includes: (1) a single identifier-issuing authority, (2) an assignment process of an identifier to a patient, (3) a permanent record of issued patient identifiers with associated traits, and (4) a maintenance process over time. The goal of Patient Identification is to reduce errors.

Patient Mapping Agent: The CCOW defined component that provides for the mapping of patient identifiers across disparate patient identity domains.

7845 **Patient Privacy Policy Acknowledgement Document:** A document that follows the BPPC Content Profile and captures the act of the patient acknowledging a specific Patient Privacy Policy Domain defined Patient Privacy Policy.

7850 **Patient Privacy Policy:** A Patient Privacy Policy further explains appropriate use of data/documents in a way that provides choices to the patient. The BPPC Profile places no requirements on the content of these policies nor the method used to develop these policies (See ITI TF-1: Appendix P for some guidance on developing these policies). A Patient Privacy Policy will identify who has access to information, and what information is governed by the policy (e.g., under what conditions will a document be marked as containing that type of information). The Patient Privacy Policy may be a consent policy, dissent policy, authorization policy, etc.

7855 **Patient Privacy Policy Identifier:** A Patient Privacy Policy Domain assigned identifier (OID) that uniquely identifies the Affinity Domain: Patient Privacy Policy. There is one unique identifier (OID) for each Privacy Policy within the Patient Privacy Policy Domain.

Patient Privacy Policy Domain: The Domain for which the Patient Privacy Policies apply. When using XDS this would likely be equivalent to the XDS Affinity Domain.

Patient Subject: The PSA defined subject that supports sharing the currently selected patient identifier amongst disparate applications running on the desktop.

7860 **PDF:** Portable Document Format.

7865 **Personnel White Pages:** Information on human workforce members within the authority of the PWP directory. This information has broad use among many clinical and non-clinical applications across the healthcare enterprise. The information can be used to enhance the clinical workflow (contact information), enhance the user interface (user friendly names and titles), and ensure identity.

PIX: Patient Identifier Cross-referencing Integration Profile.

PMA: Patient Mapping Agent component as defined by CCOW.

- 7870 **Principal:** An end user, an application, a machine, or any other type of entity that may act as a requester in a transaction. A principal is typically represented in a transaction with a digital identity and the principal may have multiple valid digital identities to use with different transactions
- Process Flow Diagram:** A graphical illustration of the flow of processes and interactions among the actors involved in a particular example.
- PSA:** Patient-Synchronized Applications Integration Profile.
- 7875 **RID:** Retrieve Information for Display Integration Profile.
- RIS:** Radiology Information System.
- Role:** The actions of an actor in a use case.
- RSNA:** Radiological Society of North America.
- Scope:** A brief description of the transaction.
- 7880 **Secure Domain:** A network, hardware systems, secure nodes, and physical environment for which a single set of security policies is defined and enforced for access to its addressable objects.
- Secure Node:** A network-addressable system that conforms to a secure domain's access policies and management. A secure node often supports IHE actors.
- 7885 **SNTP:** Simple Network Time Protocol. This is a reduced accuracy version of NTP. The protocol fields are the same, but the data values and algorithms used are greatly reduced accuracy so that it can be implemented on limited capacity systems.
- Submission Set:** A set of XDS documents registered together to a Document Repository concerning information related to one care event of a single patient, provided by an EHR system.
- 7890 **SUID:** The Study Instance UID from a DICOM SOP instance, or collection of SOP instances.
- TGT:** Ticket Granting Ticket. The initial credentials that verify that the user has been authenticated. It is used to avoid repeated user authentication events and as a token to request access to services.
- 7895 **Trigger Event:** An event such as the reception of a message or completion of a process, which causes another action to occur.
- UID:** Unique Identifier (See also Globally Unique Identifier).
- Universal ID:** Unique identifier over time within the UID type. Each UID must belong to one of specifically enumerated species. Universal ID must follow syntactic rules of its scheme.
- Use Case:** A graphical depiction of the actors and operation of a system.
- 7900 **Username:** A sequence of characters, different from a password, that is used as identification and is required when logging on to a multi-user computer system, LAN, bulletin board system, or online service. Also called user ID, or uid.

User Assertion: A set of claims about an authenticated principal (user, application, system...) that is issued by an identity provider

7905 **User Subject:** The PSA defined subject that supports sharing the user identity of the currently logged in to the applications on the desktop.

UTC: Universal Coordinated Time. This is the replacement for GMT. It defines a reference time base that is internationally recognized and supported.

7910 **Value Set:** A uniquely identifiable set of valid concept representations where any concept representation can be tested to determine whether or not it is a member of the value set. A value set may be a simple flat list of concept codes drawn from a single code system, or it might be an unbounded hierarchical set of possibly post-coordinated expressions drawn from multiple code systems. Also known as a list of valid concept codes.

7915 A valid concept is a concept that would be logically representative of the Value Set that it belongs to, for example for the Value Set “Colours of the rainbow”, “yellow” would be a valid concept.

Wet Signature: Ink on paper signature.

X-Assertion Provider: This is a SAML Identity Provider (IDP) or WS-Trust Security Token Service (STS), and is not further specified by IHE.

7920 **XDS Affinity Domain Policy:** XDS Affinity Domain Policy that clearly defines the appropriate uses of the XDS Affinity Domain. Within this policy is a defined set of acceptable use Privacy Consent Policies that are published and understood.

7925 **XDS Document:** An XDS Document is the smallest unit of information that may be provided to a Document Repository and registered in a Document Registry. An XDS Document may contain simple text, formatted text (e.g., HL7 CDA Release 1), images (e.g., DICOM) or structured and vocabulary coded clinical information (e.g., CDA Release 2, CCR), or may be made up of a mixture of the above types of content.

7930 **XDS Folder:** An XDS Folder allows document sources to group the documents they submit with other related documents. What constitutes a Folder and the vocabulary associated with the specific Folders used by an EHR-CR is decided by an agreement between the care delivery organization members of an XDS Affinity Domain.

XUA: Cross-Enterprise User Assertion Integration Profile