Integrating the Healthcare Enterprise

IHE Patient Care Coordination (PCC) White Paper

A Data Access Framework using IHE Profiles

Draft for Public Comment

Revision 1.0

Date: March 28, 2014

Author: Nagesh (Dragon) Bashyam and Keith W. Boone

Email: pcc@ihe.net

Please verify you have the most recent version of this document. See here for Final versions and here for Public Comment versions.

Foreword

Integrating the Healthcare Enterprise (IHE) is an international initiative to promote the use of standards to achieve interoperability among health information technology (HIT) systems and effective use of electronic health records (EHRs). IHE provides a forum for care providers, HIT experts and other stakeholders in several clinical and operational domains to reach consensus on standards-based solutions to critical interoperability issues.

The primary output of IHE is system implementation guides, called IHE Profiles. IHE publishes each profile through a well-defined process of public review and trial implementation and gathers profiles that have reached final text status into an IHE Technical Frameworks.

This white paper is published on March 28,2014 for Public Comment. Comments are invited and can be submitted at http://www.ihe.net/PCC_Public_Comments. In order to be considered in development of the Final version of the white paper, comments must be received by April 27, 2014.

General information about IHE can be found at: www.ihe.net.

Information about the IHE Patient Care Coordination domain can be found at: ihe.net/IHE_Domains.

Information about the organization of IHE Technical Frameworks and Supplements and the process used to create them can be found at: http://ihe.net/IHE_Process and http://ihe.net/Profiles.

The current version of the IHE Patient Care Coordination Technical Framework can be found at: http://www.ihe.net/Technical_Frameworks/.

CONTENTS

1 Introduction

This document, the IHE PCC Data Access Framework using IHE Profiles White Paper, describes how to use existing IHE profiles to access data from Healthcare IT Systems.

1.1 Purpose of the Data Access Framework using IHE Profiles White Paper

The purpose of this white paper is to provide a framework of modular, substitutable and interoperable integration profiles that shows how IHE enables data access for a wide variety of use cases and can reduce integration costs by encouraging standards based integration both within and across enterprises.

1.2 Intended Audience

The intended audience of the IHE PCC Data Access Framework using IHE Profiles White Paper is:

- Organizations integrating healthcare IT solutions within their local IT departments
- Organizations connecting directly with other organizations to support sharing of healthcare data
- Organizations connecting via federated Health IT data access solutions to support exchange of healthcare data.
- Vendors developing Health IT solutions for the any of the above organizations
- Bodies selecting standards or developing information sharing policies supporting health information sharing.
- Experts involved in standards development

2 Introduction

Many countries are reaching a critical mass of Health IT systems (EHR Systems, EMRs, hospital information systems, medical record systems, data warehouses, etc.) that comply with data and vocabulary standards. The wide deployment of Health IT systems has created unique opportunities for providers, provider support teams, patients, public health agencies, healthcare professionals and organizations and others to access and use the patient data that is already collected during clinical workflows. This information may not be readily accessible through the applications to which the relevant party has access. Allowing access to this data can enable a provider to further analyze the collected data to understand a patient's overall health, the health of a provider's collective patient population, and use the data to power analytics applications and tools to take better care of patients and populations.

2.1 Problem Description

Data access can be accomplished via various mechanisms among which queries are one of the most widely used mechanisms many of which include the use of existing IHE profiles. Enabling data access using queries has to address the following challenges:

- Queries within and across enterprises
- Queries accessing individual patient data and queries accessing population data
- Support for multiple data models based on the type of data being accessed such as clinical data, provider data, and patient demographic data.
- Storage of data in multiple applications
- Securing the data being exchanged (both query inputs and query outputs)

A single integration profile typically does not address all of the outlined challenges above, however a framework of modular, substitutable, interoperable integration profiles shows how IHE enables data access for a wide variety of use cases and can reduce integration costs by encouraging standards based integration both within and across enterprises. This is further discussed in the scope statement.

2.2 Scope

The scope of the white paper is to describe a framework of IHE Integration Profiles and supporting standards that can support queries in a modular way, allowing for substitutions in a structured way to support greater levels of interoperability between systems.

Specifically the scope of the white paper involves:

- Developing a framework of profiles supporting data access that meets the requirements for common data use cases supporting individual patient and population health.
- Describing a means by which IHE profiles can support multiple means of access through a modular framework with substitutable components (profiles).

- Identifying the gaps in the existing profiles to meet the use cases and requirements outlined.
- Identifying a roadmap for future development of IHE profiles to meet those gaps and which supports emerging standards (e.g., HL7 FHIR).

The following aspects will not be in the scope of this white paper:

- Harmonization of existing profiles
- Development of new profiles
- Development of implementation guides

2.3 Approach

The organization of this white paper is based upon a variety of different Enterprise Architecture organizational frameworks, including RM-ODP, the Zachman Framework, and HL7's SAIF approach, as well as the organizational structure of IHE profiles. HL7 SAIF defines a framework that provides three specific perspectives across multiple architecture dimensions.

The three specific perspectives enable describing a healthcare information system at a:

- Conceptual level
- Platform Independent Level (Logical)
- Platform Specific Level (Implementation)

For each of the above perspectives the data access framework is defined using the following architecture viewpoints:

- Business dimension "Why" (Used to capture business requirements, policies etc.)
- Information dimension "What" (Used to capture the data model or content)
- Computational (Behavioral) dimension "How" (Used to capture behavior, collaboration, transactions)
- Engineering and Technology dimension "Where" (Used to capture the implementation aspects including platforms, standards, integration profiles etc.)

These dimensions and perspectives are derived from similar dimensions in the RM-ODM framework which consists of enterprise viewpoint, information viewpoint, computational viewpoint, engineering and technology viewpoints as shown in Figure 2.3-1 below.

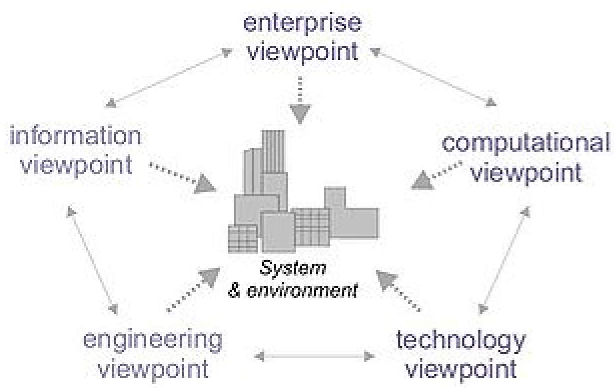

Figure 2.3-1: RM-ODP viewpoints applied to a System and Environment[1]

An example of the SAIF approach as applied to IHE profiles appears below in Figure 2.3-2 An HL7 SAIF Approach applied to the IHE Process.

	Perspectives	**Business Requirements**	**Information Models**	**Behavioral Models**	**Engineering & Technology**
Specifications	*Conceptual*	Scope, Problem Statement, Story Boards, Use Cases, Scenarios	Conceptual data models and metadata models.	Actors and Transactions, Use Case Roles, Trigger Events	Systems and their limitations, List of candidate standards
	Platform Independent	Authorities and Alignments, Core Principles, Priority user stories	Mapping of conceptual models to existing metadata and clinical data specifications	Process Flow, Referenced Standards, Interactions	Narrow down the candidate Standards to specific standards based on use cases and user stories
	Platform Specific		Message Semantics, Templates, Schema	Expected Actions and Behaviors, WSDLs	APIs, Protocols, Transforms, Schematron

Figure 2.3-2: An HL7 SAIF Approach applied to the IHE Process

IHE profiles are traditionally organized into two main parts, identified by the volume of an IHE Domain's Technical Framework in which they appear. Recent modifications to the technical framework structure over the last few years have subdivided Volume 2 into multiple parts, separating content modules (Volume 3) from transaction specifications (Volume 2). Volume 4 contains modifications to the implementable transactions and content based on regional business requirements (e.g., ensuring conformance to regional policy).

[1] Image © 2008 by Marcel Douwe Dekker reproduced with permission. Retrieved from http://en.wikipedia.org/wiki/File:RM-ODP_viewpoints.jpg

2.3.1 IHE Perspectives

The approach used by this white paper combines the Conceptual and Platform Independent level of SAIF and/or RM-ODP into a single Conceptual level. The conceptual level of specification in IHE corresponds to content found in Volume 1 of IHE technical frameworks. Implementable content is found in Volume 2 through 4.

2.3.2 Viewpoints

The viewpoints expressed in various enterprise frameworks are also expressed in IHE technical frameworks. These are summarized in table 2.3.2-1 below and described in more detail in the sections that follow.

Table 2.3.2-1: Specification Content

Viewpoints	Perspectives	
	Conceptual	**Implementable**
Business Requirements	Volume 1: Scope, Problem Statement, Use Cases and Scenarios	Volume 4
Information Models	Volume 1: Concepts (Conceptual Models)	Volume 2: Message Semantics Volume 3: Content Modules Volume 2/3 Appendixes: Schema
Behavior Models	Volume 1: Process Flows, Actors and Transactions Volume 2: Use Case Roles, Trigger Events	Volume 2: Interactions, Expected Actions and Behaviors Volume 2 Appendixes: WSDLs Appendix V: Web Services for IHE Transactions
Engineering	Project Scope Statement: Information Systems	Volume 2: Specific Standards, APIs, Protocols and Transforms Volume 3: Content Modules
Technology	N/A	N/A

Note in the table above that IHE does not provide any specific requirements at the conceptual or implementable perspective for the technology viewpoint. This means that any technology or platform could be used to implement the profiles.

3 Conceptual

Most of the conceptual content of an IHE profile is found in Volume 1 content of a domain's technical framework.

3.1 Business Requirements

This section captures the business requirements including any policy constraints that need to be addressed by this white paper. The business requirements driving the development of the various IHE profiles are found in the profile descriptions in Volume 1.

For the data access framework, the following business requirements have been identified based on user stories outlined in the Appendix.

Table 3.1-1: Data Access Framework Business Requirements

Req#	Requirement Text	User Story	Query Name
1	Get all clinical summary documents produced locally and those received from other healthcare facilities for a single patient so that the provider can analyze the patients overall health.	B.1	Find Patient Identifier for Patient Demographics Find Document(s) based on Patient Identifier Find Document(s) based on Patient Demographics Get Document(s) based on Document Identifier
2	A patient requests his care team provider(s) to provide him/her with all their medical documents while preparing to move from one state to another state.	B.2	Queries identified as part of Req #1 above.
3	Gastroenterologist sets up queries so that he can be alerted based on specific fasting glucose values for a patient.	B.3	Get clinical data for a patient based on discrete data elements
4	Gastroenterologist queries their EHR system to retrieve all documents for a patient including sensitive records	B.3	Queries identified as part of Req #1 above.
5	Gastroenterologist collects the patient consent to disclose sensitive records to prepare for a referral and authorizes the Endocrinologist to be able to query sensitive information.	B.3	Capture Patient Consent Supply and Consume User Assertions.
6	An Authorized Endocrinologist is allowed to access sensitive records during a referral process.	B.3	Supply User Assertions Queries identified as part of Req #1 above.
7	Researchers try to access sensitive data	B.3	Supply User Assertions Queries identified as part of Req #1 above.
8	A PCP retrieves clinical summaries for males patients over the past 5 years to analyze using a 3rd party analytics application	B.4	Get Document(s) for multiple patients
9	PCP extracts male patients with cardiovascular disease and diabetes	B.5	Identify Patient(s) based on discrete data elements

Req#	Requirement Text	User Story	Query Name
	within the past 5 years from the list of patients within a patient panel		Get Document(s) for multiple patients
10	PCP queries all lab results with HbA1c > 7% over the past 12 months for a single patient	B.6	Get clinical data for a patient based on discrete data elements
11	An application queries the EHR for patient demographics, and admitting diagnosis and any clinical documents to prepare patient instructions.	B.7	Find Patient Demographics based on Patient Identifiers. Queries identified as part of Req #1 above.
12	Physician identifies all patients with Hepatitis C diagnosis but have not had fasting glucose tests since the start of their Hepatitis C treatment	B.8	Identify Patient(s) based on discrete data elements
13	A nurse during preop queries for the patients data such as problems, meds and allergies	B.9	Get clinical data for a patient based on discrete data elements Queries identified as part of Req #1 above.

These requirements address several different aspects of query which are further described in the subsections below.

3.1.1 Support Multiple Levels of Enterprise Governance

There are three key levels of enterprise governance as shown in the table below.

Table 3.1.1-1: Enterprise Query Governance

Data Location	Governance
Within an Enterprise	Intra-Enterprise
Between specific Enterprises	Inter-Enterprise
Multiple external Enterprises	Federated

Enterprise complexity reflects the degree of governance necessary over query interactions. Enterprise in this sense often is mapped onto organizations; however, the size of the organization is what really matters. For example, a multi-site, multi-regional organization may require a federated governance model even though it is controlled by a single legal entity (organization).

3.1.1.1 Intra-Enterprise Governance

Policies within a single enterprise are relatively static, and under the control of a single organization. Intra-enterprise interchange can often be preconfigured and is fairly static with respect to governance being applied to the exchange. Security decisions with respect to authentication, authorization, audit, and access controls are readily controlled at a local level, and can be preconfigured. Endpoints are readily known and relatively static. Changes at this level can readily be managed because the organization making the change also has the authority and capability to change the configuration of the systems that need to be changed to account for it.

For example, to provide or remove access to an individual person, or an existing information system is something that is completely under local control within the enterprise, and does not require a great deal of coordination with respect to governance.

3.1.1.2 Inter-Enterprise Governance

While Inter-Enterprise complexity may still involve relatively static policies, those policies must be negotiated, and are not under the control of a single governance body. Trading partners are relatively stable. Some forms of exchange which are centrally managed (e.g., such as those negotiated through a regional health information exchange organization) still fit this model, even though trading partners are more dynamic. This is because the trading partner exchange agreements are negotiated between individual members of the exchange and the regional health information exchange organization, which limits the number of governing bodies that need to be negotiated to two for most participants.

Exchanges can still be preconfigured, but may need some intermediation to reflect the need to adapt to local deployment or policy decisions. For example, end-point addresses may need to be looked up in some sort of directory (usually DNS suffices), and information about local users or systems may need to be coordinated when new systems are deployed (e.g., user or node authentication or authorization policies).

3.1.1.3 Federated Governance

Federated governance often requires the ability to make dynamic, perhaps even human intermediated policy decisions to enable exchange across changing governance bodies. Trading partners enter and leave the information exchange much more dynamically, and the types of policies they support within the exchange may vary. At this level of governance, there is much more reliance on directories, and on asynchronous responses, since service level agreements cannot always be prenegotiated.

3.1.2 Query Targets

Access to data may be limited to data about a single patient, or populations of patients, or even aggregate data specific to a population of patients.

Gap: While IHE has some profiles supporting aggregated results (Quality Measure Execution – Early Hearing QME-EH), they are limited to very specific use cases, and a generalized model supporting a variety of aggregated measures are not available, and so will not be further considered in this white paper.

3.1.3 Query Granularity

Queries can be issued at either a coarse or find grained level. Queries can be for:

1. Specific documents captured during workflows
2. Documents containing specific data
3. Specific clinical data at the atomic level

4. Pre-computed aggregated data on a population (e.g., specific quality measure results)
5. Knowledge artifacts informing decision support, e.g., clinical guidelines, decision support interventions, or educational content.

Gap: As previously noted, queries for pre-computed aggregated data are not generalized in IHE profiles at this time, so item #4 above will not be supported.

Note: Item #5 is not based on clinical data stored about patients, and will not be addressed in this modular framework.

3.1.3.1 Query for Existing and Generated Documents using Encounter Documentation

DAF queries can retrieve existing documents which are already present in repositories. These documents get created during clinical workflows and document the events, actions, instructions relevant to a patient's encounter. DAF queries can also retrieve documents which are dynamically generated when the queries are executed by the responding system. In either case DAF queries would use the Encounter Documentation such as creation time, type of document etc. to query for documents.

3.1.3.2 Query for Existing and Computed Data using Detailed Clinical Information

DAF queries can retrieve granular data (such as problems, medications, allergies) which are already present in repositories. The relevant data gets captured during clinical workflows and documented as Detailed Clinical Information relevant to the patient. DAF queries can also retrieve data that is computed based on certain criteria about the Detailed Clinical Information present in the system.

3.1.3.3 Query for data within Enterprise (Intra-Enterprise)

Queries described in this white paper can readily retrieve existing metadata, documents and detailed clinical data present within the enterprise. This query framework can use simpler computational patterns without need for federation, and may also require the use for fewer protocols to support privacy and security, favoring instead the intra-enterprise security and privacy controls already in place within the enterprise. These controls may include manual processes for granting and revoking user access, patient consent processes that can be assumed to be adopted organizationally, and physical security used to ensure that connected computers can only be accessed through specific physical network connections, et cetera.

While manual and operational controls (e.g., physical security) may be sufficient in this environment to support the data access framework, we would recommend application of the IHE Audit Trail and Node Authentication profile as it mitigates risk against access attempts from systems attached to the network (either authorized or not), encrypts network communications so that other systems on the network cannot eavesdrop, and provides an audit log of user actions that is often required in many regulatory environments.

3.1.3.4 Query for data from Specific External Enterprise (Inter-Enterprise)

When requesting information from external enterprises, additional security is often required. This kind of data access will need to ensure sufficient security controls (Authentication, Authorization etc.) are in place to allow data access from specific external enterprises. Query Requestor and Query Responder will belong to two distinct enterprises in this case.

In this environment, we would recommend application of Cross Enterprise User Assertion (XUA) or Internet User Authorization (IUA) to support user authentication across enterprises, allowing for each enterprise to manage its separate user bases, and still provide the necessary authentication/authorization information.

3.1.3.5 Query for data from multiple external Enterprises (Federated)

DAF queries can retrieve existing or computed data from multiple external enterprises. The query framework used to access data from multiple external enterprises needs to support Federated access and needs to ensure sufficient security controls (Authentication, Authorization etc.) are in place to allow data access from a multiple external enterprises. Query Requestor and Query Responder will belong to multiple distinct enterprises in the case of Federated queries. In addition there will be dynamic behavior where Query Requestors can be sending requests to new Query Responders as enterprises as discovered and removed from the eco-system.

3.1.4 Query Response Granularity

The granularity of a query response can be at several levels. A response might include:

- Metadata associated with documents or encounters
- Clinical documents (aggregations of clinical data organized by encounter)
- Discrete data from single or multiple encounters

3.1.5 Security Aspects

Query interactions must also meet a variety of security requirements. A complete list of security requirements for the use cases and user stories described above would be another document the size of this white paper. The following is summary of security requirements for the purpose of this white paper.

- Protect message integrity and confidentiality
- Supporting appropriate audit logging associated with exchanges
- Support authentication of the end user or system performing the query
- Support access control checks before accessing data
- Support documentation of patient consent before allowing access to specific data elements

IHE has prepared two separate documents which discuss security planning for profiles and access controls.

3.1.6 Transport Requirements

There are a number of different ways that information exchanges can occur over a computer network, including the use of healthcare specific, and more general SOAP and RESTful transport protocols. Where feasible, the framework must identify how multiple transport protocols can be used to enable systems of varying capabilities to interoperate with each other.

3.2 Information Models

Most IHE profiles contain a Concepts or Overview section in Volume 1 which describe the key business concepts needed to understand the problem and solution space that are addressed by the profile. Many also describe the data elements necessary to support the interoperable solution at a high level.

This section identifies the various conceptual data models that can be used to meet the DAF data requirements. In IHE, there are three models that are commonly used to query for patient data. These models address the following three topic areas:

1. Patient Demographic Information
2. Encounter Documentation
3. Detailed Clinical Data

Note: IHE also has a conceptual data model supporting provider information. This is typically used for directory purposes, such as Personnel White Pages (PWP), Healthcare Provider Directory (HPD) or Care Services Directory (CSD) and is not further discussed in the context of this white paper.

3.2.1 Patient Demographic Information

Patient demographic information follows the conceptual model first developed in the IHE PIX and PDQ profiles, and subsequently mapped into the PIX/PDQ V3 profiles. The requirement to federate queries across multiple patient identity domains led to the development of the XCPD profile. This profile uses the same conceptual model, but supports federation requirements to access patient identifiers and demographics. The latest profile to make use of patient demographics fully develops a conceptual model in Volume 1 of the Patient Demographics Query for Mobile (PDQM) Profile Supplement.

That conceptual model applies to the entire family of PIX/PDQ profiles (including XCPD). Due to the way that IHE presently develops profiles, these are treated as separate profiles. However, they are nearly functionally identical, having the same conceptual models. The conceptual data model identified by these IHE profiles is presented in Table 3.2.1-1 Patient Demographics Conceptual Data Model below. This table is derived from the February 12, 2014 draft of the PDQM profile.

Table 3.2.1-1: Patient Demographics Conceptual Data Model

Field	Queryable
Identifier List	Y
Name(s)	Y
Date / Time of Birth	Y
Gender	Y
Address(es)	Y
Telecommunications Address(es)	N
Language(s) of communication	N
Marital Status	N
Non-Medical Identifiers	N
Death Date/Time	N
Mother's Maiden Name	N
Patient Home Telephone	N
Patient Multiple Birth Indicator	N
Patient Birth Order	N
Last Update Date/Time, Last Update Facility	N

3.2.2 Encounter Documentation

Encounter documentation has three levels of detail in information models relevant to data access. The first level is directly relevant to query, and describes the metadata used to describe the documentation associated with an episode of care or encounter. The second describes the general organizing structure of encounter documentation. The third level is essentially identical to the detailed clinical data level and will be described in that section.

3.2.2.1 Encounter Documentation Metadata

The encounter documentation metadata was first developed for the Cross Enterprise Document Sharing (XDS) profile, but has subsequently been used to support Point to Point communications of encounter documentation via both reliable messaging (XDR), and media or e-mail exchange[2] (XDM). The same metadata is also used to federate queries in XCA. From an exchange perspective, the XDR, XDM, XDS and XCA content exchanged uses nearly identical metadata.

There are a few cases where extra metadata elements were added to support the needs of a specific use case (e.g., routing media over e-mail, or discriminating between federation sources).

[2] The Direct specifications developed by the US Office of the National Coordinator supports the use of both XDR and XDM profiles. Email is the principle transport for health information exchange in the Direct specifications.

XCA and XDS provide nearly the same query and retrieve capabilities, and from a receiver perspective, it is hard to tell whether the responder to a query or a retrieve is an XDS Registry, or a federated gateway.

The IHE MHD profile uses a RESTful model that subsets the query functionality found in XDS/XCA, and can in fact be readily mapped to XDS/XCA capabilities. However, it is lighter weight, mostly because it is intended for devices with less capacity for orchestration. The MHD profile maps XDS metadata into HL7 FHIR resource queries. IHE collaborated with HL7 on the development of the FHIR infrastructure resources that support the MHD profiles. These resources are patterned after the XDS metadata model.

Note that the XDS metadata conceptual model is based on metadata found in common in various healthcare standards, including DICOM, CDA Release 1, CDA Release 2, CCR and CEN/ISO 13606. It is described in the IHE ITI Technical Framework Volume 3 and will be used as a reference for further discussions throughout the document.

3.2.2.2 Document Organization

The XDS family of profiles is neutral with respect to content. These profiles have been used to exchange images, PDF documents, and various forms of XML documents. In the context of the data access framework being developed in this white paper, we will principally focus on the use of these profiles to exchange CDA documents, which are used to document encounters or episodes of care[3].

Document exchanges often use HL7 CDA documents used to document care. At a conceptual level, these are simply clinical documents. Clinical documents are compositions of clinical data relevant to an encounter or episode of care, associated with a patient (whose demographics are modeled in much the same way as IHE PIX/PDQ profiles handle information). These clinical documents are organized in a common way based on the evolution of clinical judgment, from which a second conceptual data model can be derived. That model is organized into sections that generally fall into one of the categories answering the following questions about the encounter or episode of care:

- What is the context of care provided?
  - Who are you (see patient demographics)?
  - What encounter, organization, or providers are involved?
  - What kind of healthcare service is being provided?
- Why are you here?
  - Reason for Visit

[3] An encounter is usually defined as an interaction with a single healthcare provider or provider organization (e.g., a single ambulatory visit or an inpatient stay). An episode of care can cover multiple visits with multiple providers or organizations.

- Reason for Referral
- Chief Complaint
- Reason for Procedure
- Reason for Operation
- Admitting Diagnosis
- Preoperative/Preprocedure Diagnosis
- History of Present Illness
- Et cetera.

- What do we already know about you?
  - Problem List
  - Medication List
  - Allergy List
  - Family History
  - Social History
  - History of Hospitalizations
  - Surgical History
  - Pregnancy History
  - Immunization History
  - Et cetera.
- What was done while you were here?
  - Hospital Course,
  - Procedure Description
  - Operative Description
  - Review of Systems
  - Procedures Performed
  - Et cetera
- What did we find out?
  - Physical Examination Findings
  - Diagnostic Tests and Results

    - Findings
- What do we think is going on?
    - Diagnosis
    - Impression
    - Assessment
    - Post-Procedure Diagnosis
    - Post-Surgical Diagnosis
    - Discharge Diagnosis
- What should happen next?
    - Care Plan
    - Goals

The first item (context of care) is found in the CDA Header. The last six items are stored as section content within the CDA document. Each of the sections within these categories sections are classified using section codes from Logical Observation Identifiers, Names and Codes (LOINC®). The IHE PCC Technical Framework organizes the clinical content almost along these lines, but also combines that information with more the detailed clinical information classifications described below.

3.2.3 Detailed Clinical Information

Detailed Clinical Information doesn't necessarily capture the high level structuring and context surrounding it as described above, however it does capture information that fall into the following categories:

- Problems
    - Problems may include Findings, Chief Complaints, Diagnosis etc.
- Allergies
    - Note that allergy is a specialization of problem
- Medications
- Observations
    - Observations can be subdivided into multiple categories including diagnostic results (labs, imaging and other studies), vital signs, social history, and family history, et cetera.
- Immunizations
- Advance Directives

- Care Plan
- Goals
- Encounters
- Procedures
- Interventions

There are more detailed clinical models that can be applied to the above concepts, and many have been through CDA templates, including those found in the CCD and IHE templates which have made their way back into the HL7 template framework through the Consolidated CDA efforts.

These templates are applications of constraints on the HL7 Clinical Statement model. This model originated in CDA, and is the basis of many HL7 Version 3 standards. One standard in which the basic clinical statement model of CDA Release 2 is largely recognizable is the HL7 Version 3 Care Record standard. Much of the implementable XML representation uses the same elements and structure.

This standard is the basis for a number of IHE profiles supporting access to detailed clinical information. The most significant of these is the IHE Query for Existing Data (QED) profile, which uses both the Version 3 Care Record and Care Record Query Standards to provide access to detailed clinical data for the patient.[4]

No formally published conceptual model presently exists in HL7 or IHE which captures the kinds of detailed clinical information that is available in IHE profiles based on either the CDA or Care Record specification. However, a conceptual model can be readily derived from the C-CDA templates by abstracting upwards from the implementable information models in C-CDA. The following sections depict this conceptual information model and were developed based on this process.

3.2.3.1 Context Information

The contextual information model shown in figure 3.2.3.1-1 below describes the key data elements found in both the C-CDA and Care Record specifications. These are aligned with the conceptual model for Encounter Documentation Metadata described above in section 3.2.2.1 above.

[4] While QED is designed to support access to data for a single patient, there is a rather simple "cheat" to support a population level query using the same profile

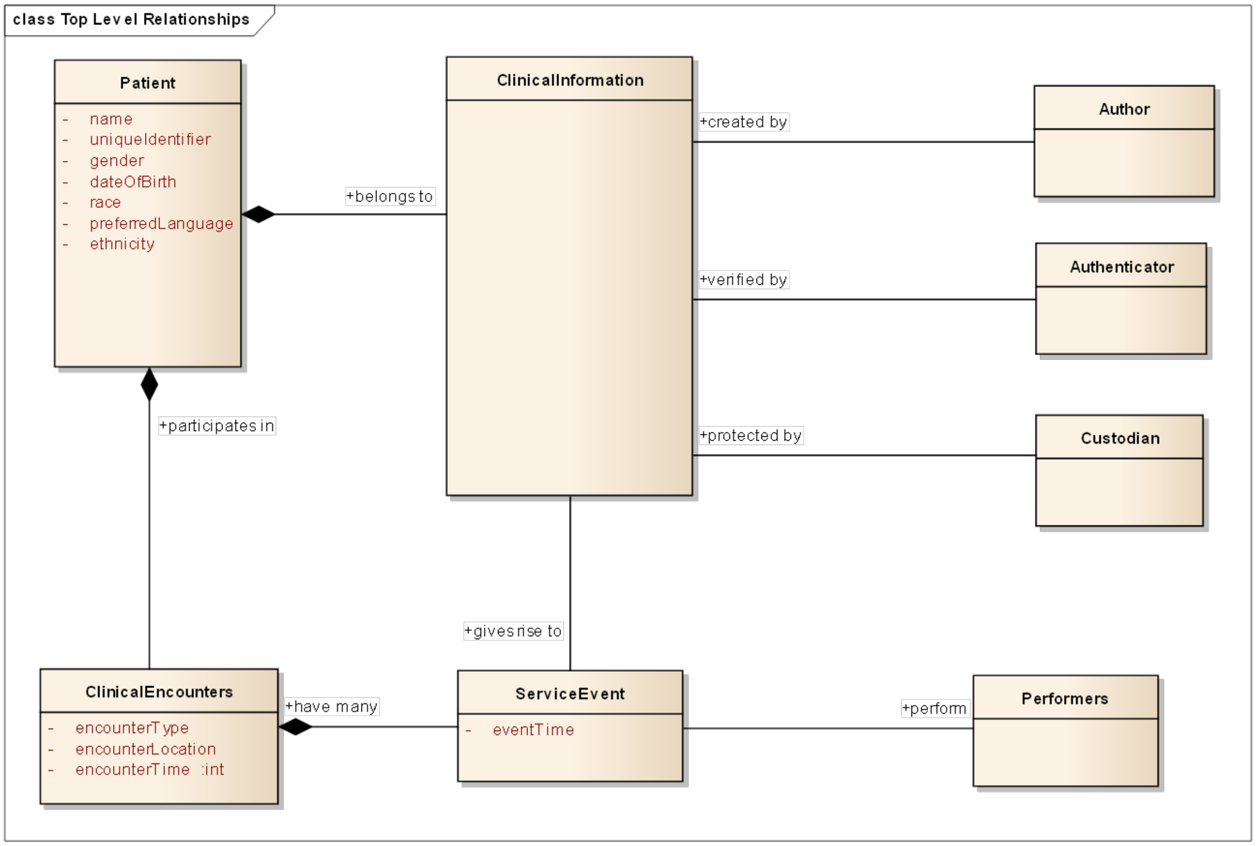

Figure 3.2.3.1-1: Contextual Information Model

560 ### 3.2.3.2 Detailed Clinical Data

The conceptual model for detailed clinical data is shown below in Figure 3.2.3.2-1 Detailed Clinical Data and Figure 3.2.3.2-2 More Detailed Clinical Data depicted below.

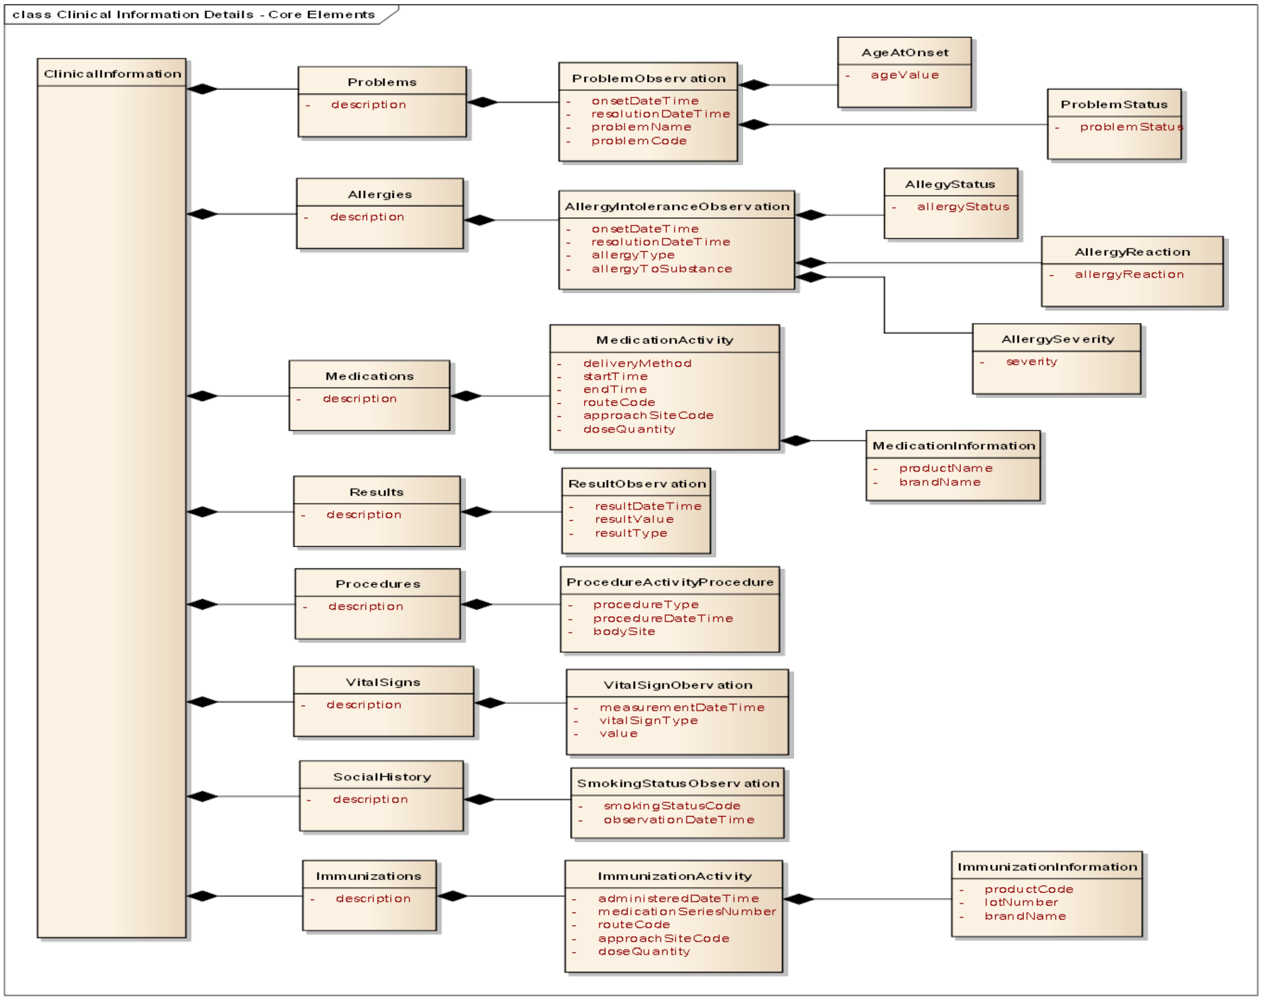

Figure 3.2.3.2-1: Detailed Clinical Data

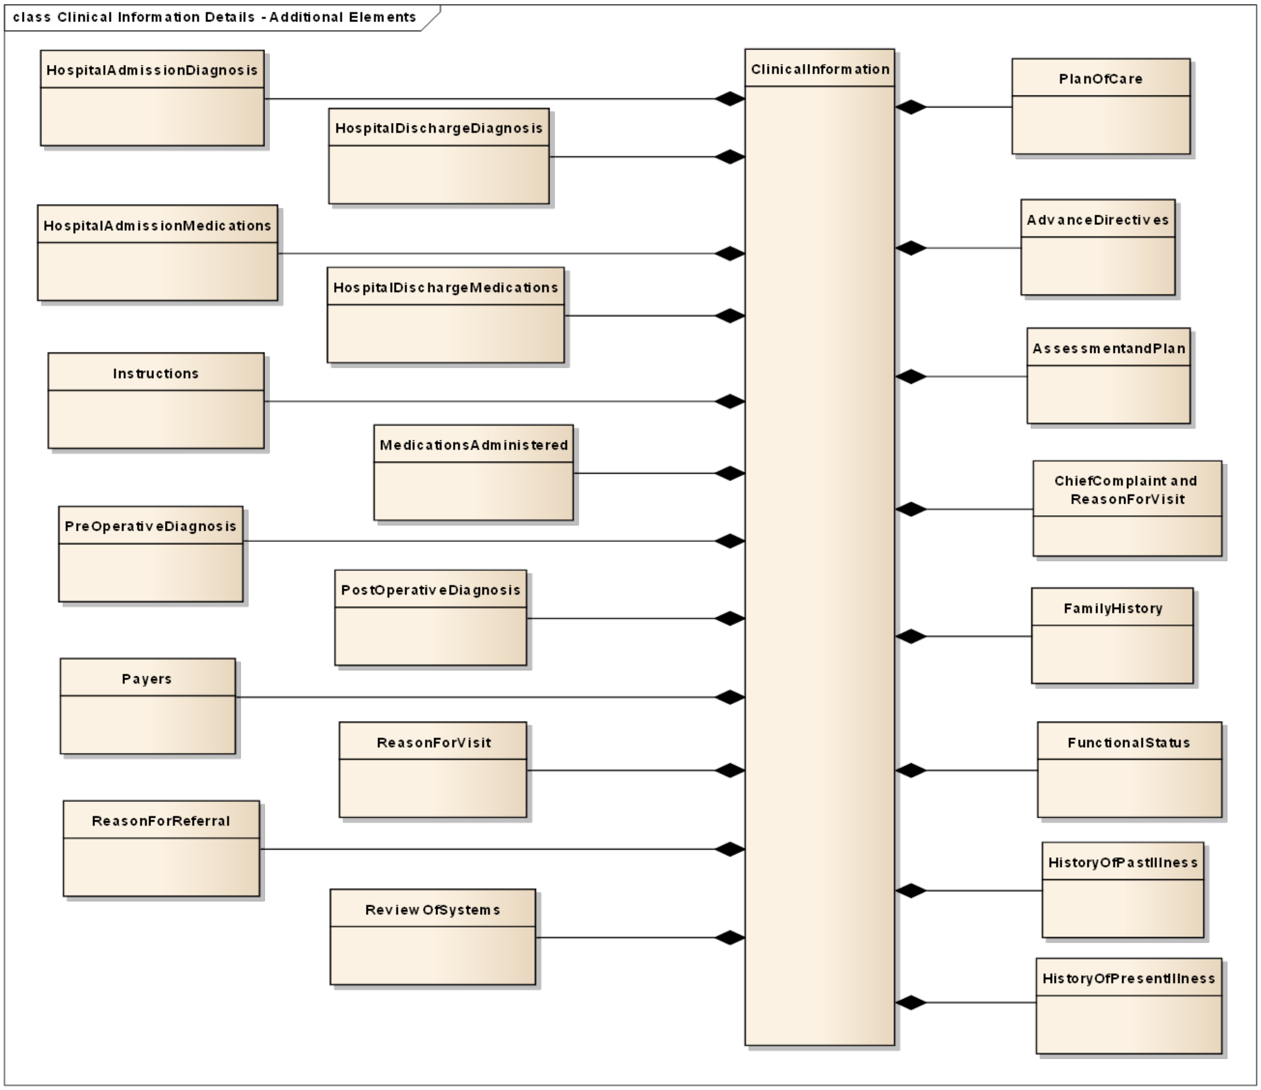

Figure 3.2.3.2-2: More Detailed Clinical Data

3.3 Behavioral (a.k.a. Computational) Models

The behavioral perspective captures the behavioral aspects and requirements of the system. This includes the various actors, roles, types of interactions and transactions that need to be supported along with the events that trigger the various interactions.

Behavioral models describing the conceptual flow of activities can be found in Volume 1 under Process Flows. System components providing the profile solution are identified as actors in IHE profiles, and the communications between those components are called transactions. The Actor / Transaction section of Volume 1 describes the conceptual arrangement of information flow, and links those flows to the implementable transactions.

Transactions describing the implementable system communications are found in Volume 2. Each transaction identifies the actors that use the transaction, and the roles that they play. The transactions also describe the events that trigger the execution of the transaction.

When the implementation platform is HTTP/SOAP, the structure of the WSDL is defined as described in Appendix V: Web Services for IHE Transactions found in Volume 2 of the IHE ITI Technical Framework.

3.3.1 Basic Behavioral Design Patterns

The subsections below describe a number of different behavioral (computational) design patterns related to query.

3.3.1.1 Request/Response

The request/response design pattern has two actors, a requestor which makes a request (1), and a responder (2) which replies to the request, providing the results in a single interaction.

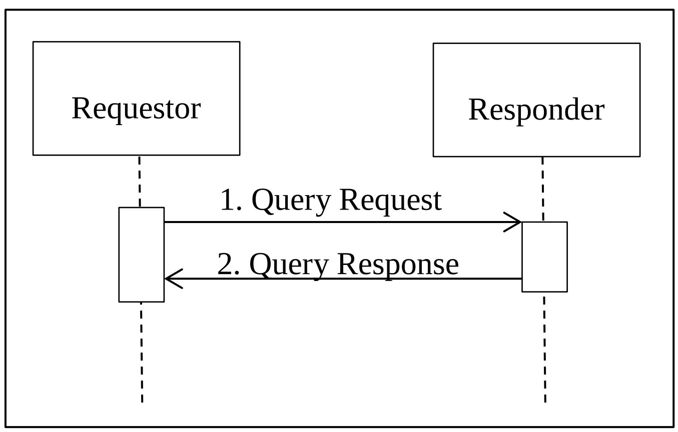

Figure 3.3.1.1-1: Request/Response Behavior Pattern

3.3.1.2 Request/Batched Responses

The basic request/response design pattern can be extended by responding to the request (1) by batching the first N results (2) that are returned by a query, and supports capture of complete results by requesting additional batches (3 and 4). Many variations of this design pattern exist, including those that allow only traversal in the forward direction, traversal forwards and backwards, or direct access to a set of N results starting at a given position. When finished, the query is canceled (5) to release server resources used to maintain state.

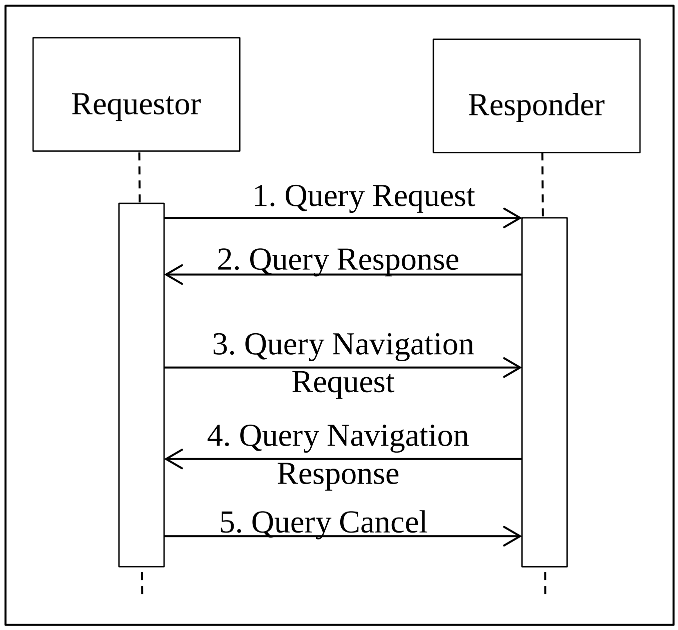

Figure 3.3.1.2-1: Statefull Query Request with Batched Responses

This behavioral pattern described above can require the storage of state information in a request. This depends entirely on upon whether the query request (1) is itself repeated in each subsequent interaction (3) or whether only a query identifier (a handle to state information) is passed back and forth. When the entire query is repeated in subsequent interactions, the stateless interaction pattern shown below can be used. To facilitate the stateless behavior, the query parameters are often duplicated in the query responses (2 and 4).

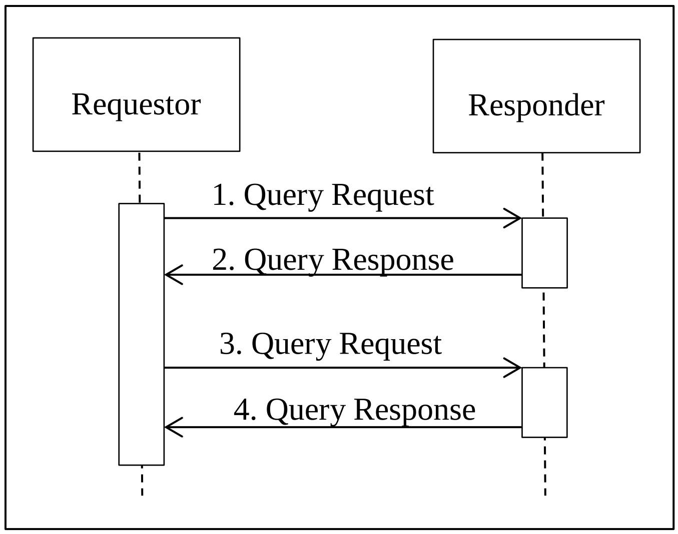

Figure 3.3.1.2-2: Stateless Request with Batched Responses

3.3.1.3 Subscribe/Publish

In the subscribe/publish behavioral model, the subscriber requests information based on a query to a publisher (1). The publisher then responds to the subscriber with individual responses (2,3) that match that query as they become available. This continues until the subscriber cancels the subscription (4). Individual responses are, of necessity, asynchronous.

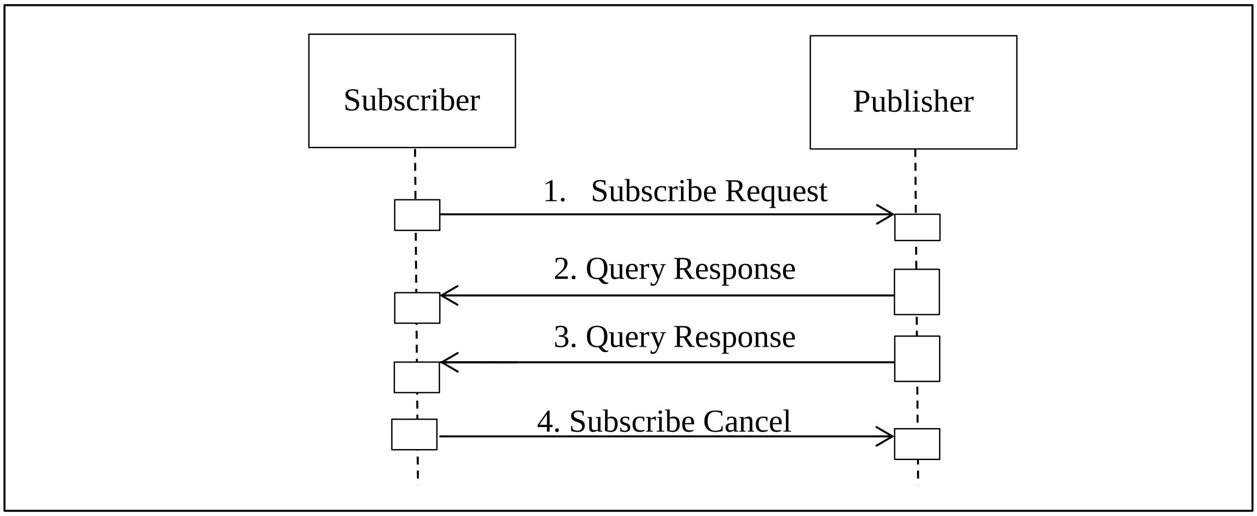

Figure 3.3.1.3-1: Subscribe/Publish

3.3.2 Federation Patterns

The federation pattern generalizes interactions between two systems. Requests are made to a gateway (a Server grouped with a Client) (1) that forwards them on to other servers (2). Responses (3) are either brokered by the gateway as shown in the diagram below (4), or returned directly. Responses can be handled synchronously (within the same activation line), or asynchronously with separate activations. Subsequent patterns refine this generalization.

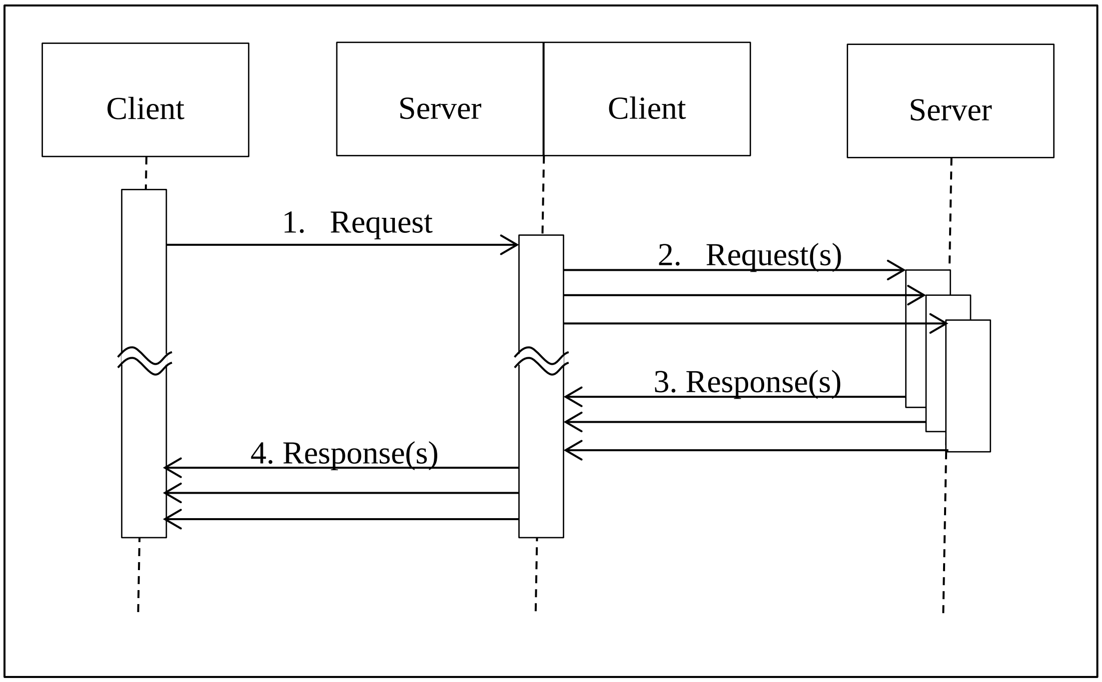

Figure 3.3.2-1: General Federation Pattern

3.3.2.1 Federated Request/Response Pattern

In the federated request/response pattern, the basic request (1) it passed to a gateway system (a responder grouped with a requester), which then forwards it to other systems (2) on the requestor's behalf. The results are returned to the gateway (3) which then aggregates all results, deduplicating as necessary (4), before passing them back to the requestor (5).

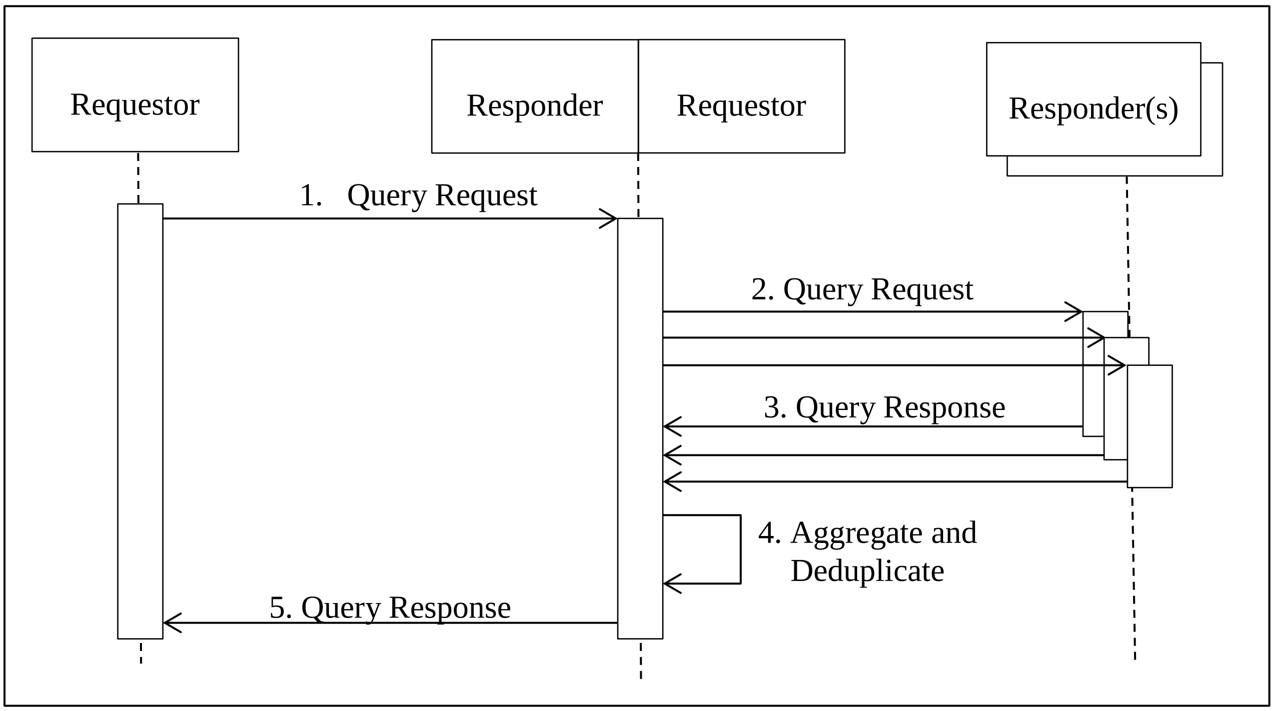

Figure 3.3.2.1-1: Federated Request/Response Pattern

3.3.2.2 Asynchronous Federated Request/Response Pattern

In the asynchronous federated request/response pattern, a request is made to a gateway (a responder grouped with a requestor) (1), which forwards them to other responders (2). The responses (3) are returned directly to the requestor without intermediation by the gateway. The requestor listens to all responses and after some time has passed (4) is responsible for the aggregation and deduplication of results (5). In this behavioral pattern, the requestor must make a (possibly arbitrary) determination when all results have been sent.

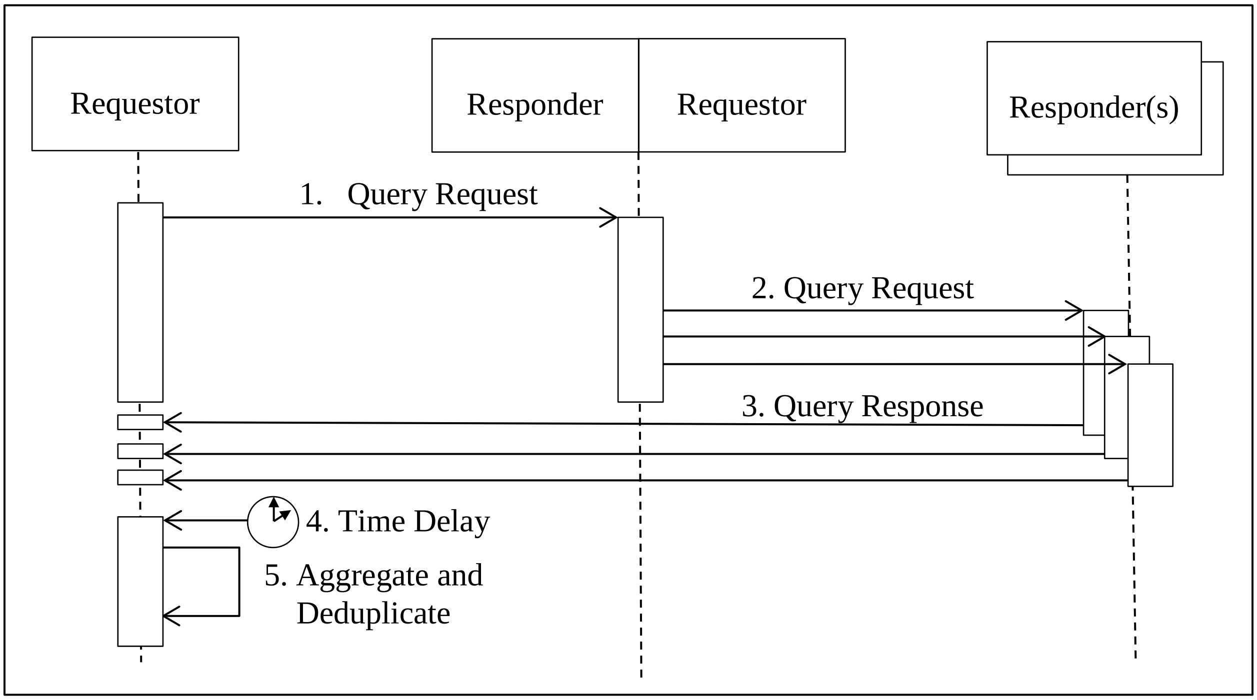

Figure 3.3.2.2-1: Asynchronous Response

3.3.2.3 Federated Subscription

The federated subscription model is one in which the subscriber passes its request to a gateway (a publisher grouped with a subscriber) (1) which forwards it to other information sources as a new subscriptions (2). Responses can be returned asynchronously by the gateway (3) and after it receives them, it just forwards them to the subscriber (4), or returned directly (5) as in the Asynchronous federated request/response pattern. A cancelation sent to the gateway (6) is also fanned out to all publishers (7) who received a copy of the original subscription.

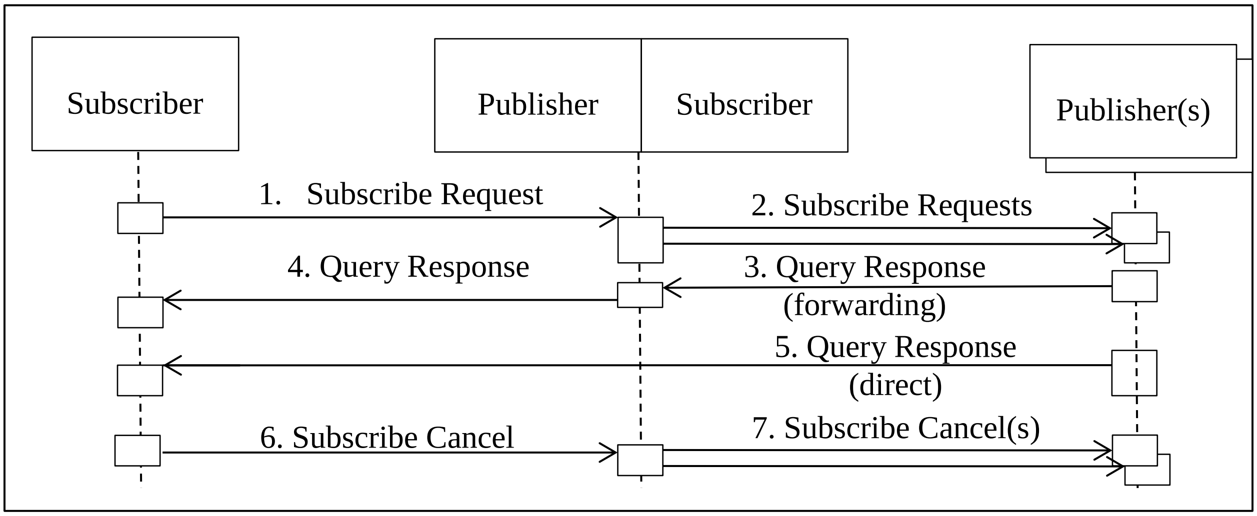

Figure 3.3.2.3-1: Federated Subscription

3.3.3 Security and Privacy Overlays

Many of the behavioral patterns described above must be combined with other capabilities to ensure that security and privacy requirements are supported. These requirements can often be supported by the application of general behavioral patterns described in the subsections below.

3.3.3.1 Node Authentication

Node authentication entails grouping two interacting systems with an actor that exchanges node certificates securely (1). Note that the activation line of the Secure Node continues throughout the entire transaction to show that the exchange(s) continue with assurance that the systems that are communicating have authenticated each other. This pattern is implemented in the IHE Audit Trail and Note Authentication (ATNA) Profile.

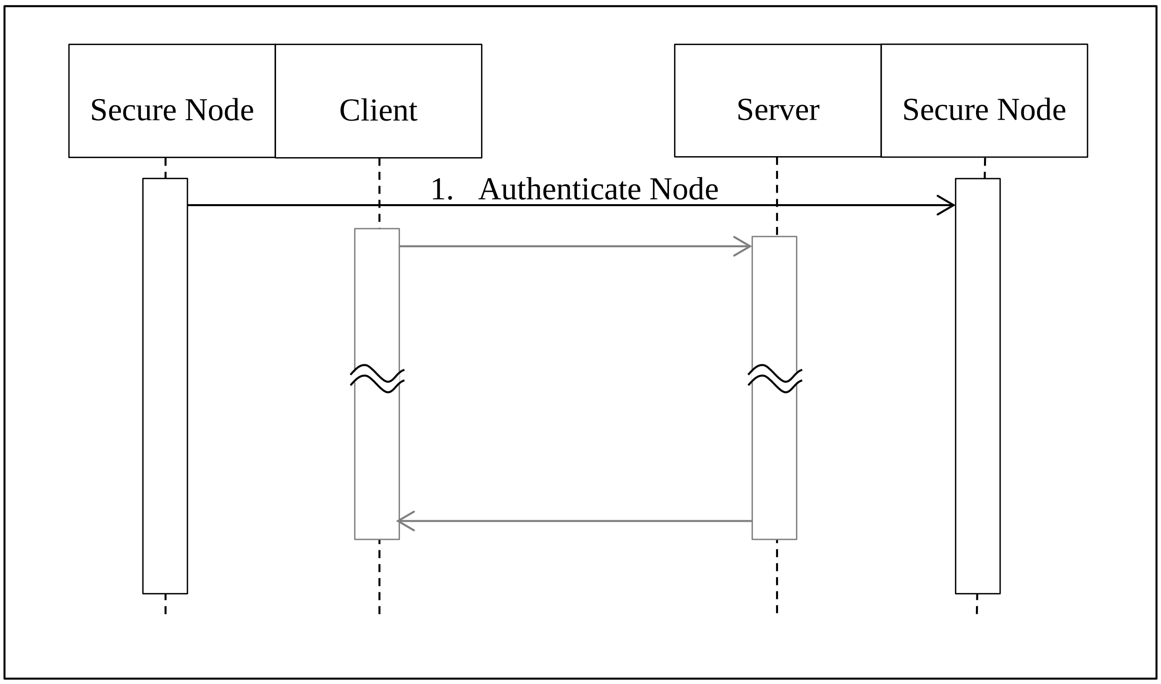

Figure 3.3.3.1-1: Node Authentication

3.3.3.2 Encryption

Encryption makes use of node authentication (1), and establishes a secure channel which is then encrypted (2). Again, this capability is enabled by the IHE ATNA Profile.

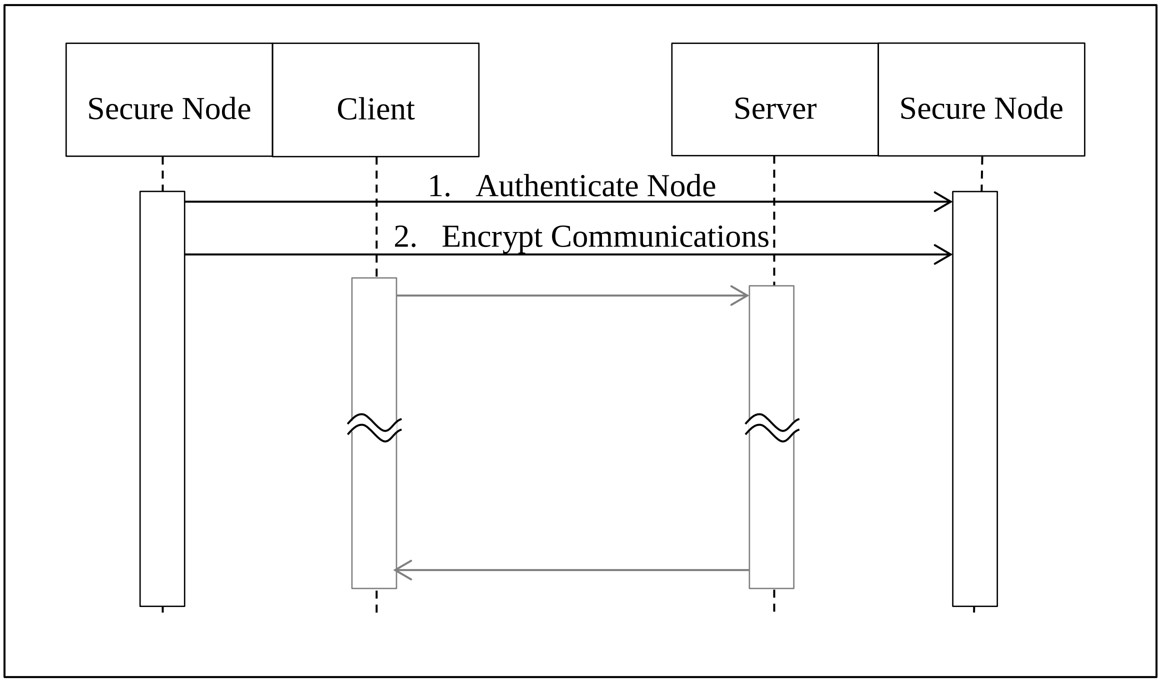

Figure 3.3.3.2-1: Encryption

3.3.3.3 Audit Logging

Audit logging is handled by grouping an actor with the secure node. Upon execution of a communication (1 and 4) with another actor, the sender requests generation of an audit event (2) by the Secure Node. The Secure Node sends the audit event to the Audit Repository (3). Upon receipt of the communication, the receiver also requests generation of an audit event (4) by its Secure Node. That audit event is also sent to an audit repository (5). Note that the audit repository used by the Sender and the Receiver may be different. The mechanism by which a request is made of the Secure Node to generate an audit event is left to the implementer. Note that each communication should result in the appearance of two correlated audit events in the Audit Repositories, one initiated by the sender, and the other initiated by the receiver. The order of receipt of the two audit events is immaterial; however, what is material is that the time stamps of the events be based on a synchronized clock (not shown).

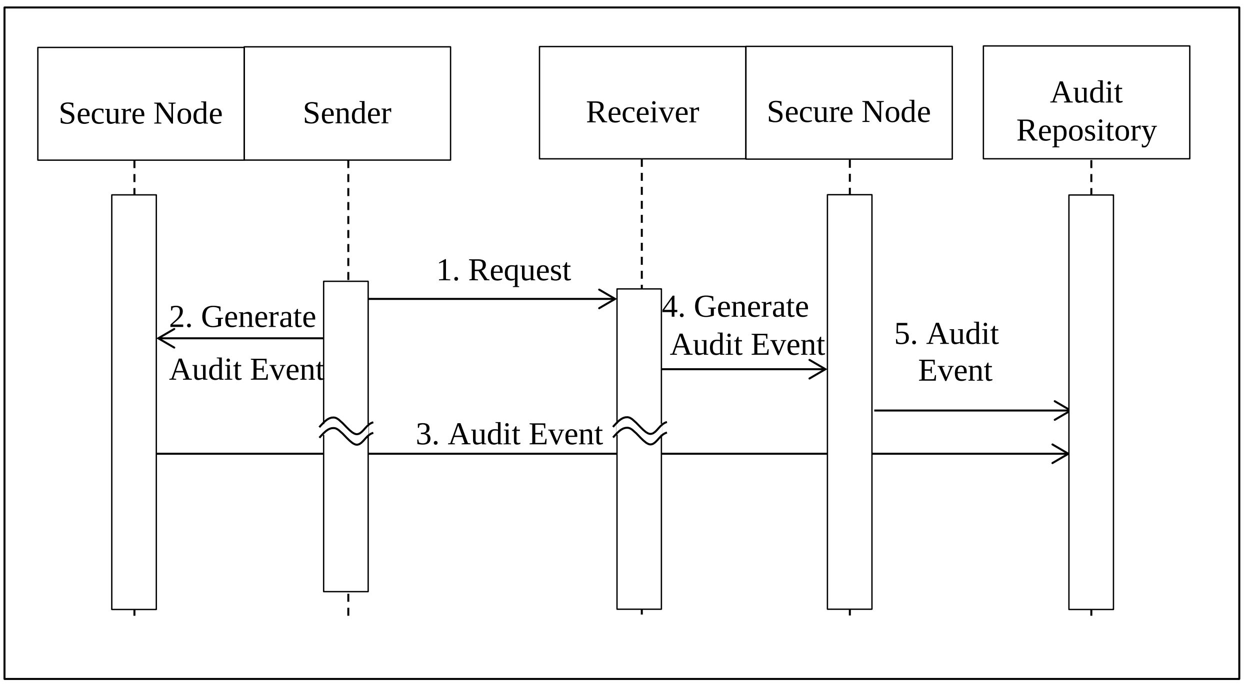

Figure 3.3.3.3-1: Audit Logging

3.3.3.4 User Authentication / Authorization

User Authentication involves requesting proof of authentication (which is called a token here) (1) with a separate token issuer that is responsible for ensuring that the user is authenticated and/or has authorized the sending application (the Sender) through an unspecified mechanism (2). This results in the generation of a token (a.k.a. ticket or assertion) that can be passed to a Receiver Actor along with a request (3). The Receiver validates the token using the Validator Actor (4). That identity validation may be done by internally (e.g., using cryptographic methods), or by requesting verification that the token or assertion is valid from the original authenticator (5). The sender may indicate that the user session is terminated (6).

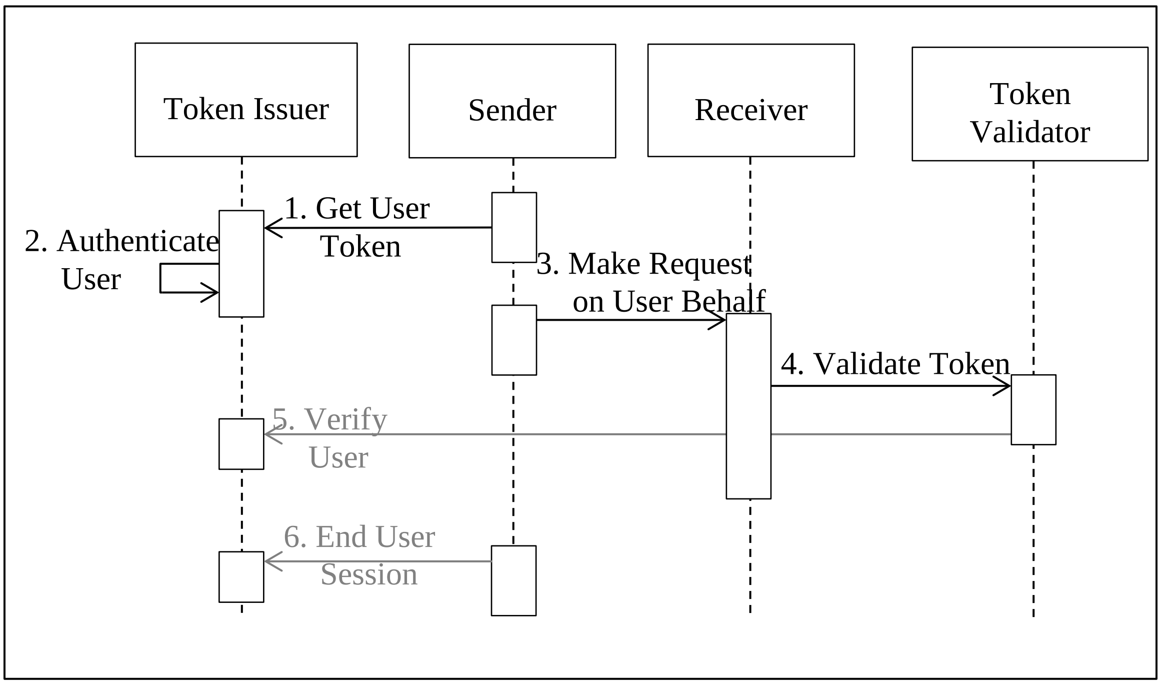

Figure 3.3.3.4-1: User Authentication and/or Authorization

This general pattern is followed by the IHE Enterprise User Authentication (EUA) profile, the Cross-Enterprise User Assertion (XUA) profile, and the Internet User Authorization (IUA) profile.

In EUA, the Authenticator is the combination of the User Authentication Provider and the X-Assertion Provider. Steps 1 and 2 don't necessarily occur in that order. The communication of the token (called a ticket) is handled through the Kerberos protocol.

In XUA, the token is a SAML assertion, and the Token Issuer and Validator become the Assertion Provider Actor of that profile. In the XUA profile, the Assertion Provider and User Authentication Provider Actors are grouped. These are shown together above as the Token Issuer. Note that again Steps 1 and 2 don't necessary occur in that order.

In IUA, the Authorization Server Actor is represented as the Token Issuer above. IUA does not specify the mechanism by which the token is validated; however, it uses OAuth 2.0, which follows the same patterns described above. SAML assertions may be included in, or retrieved from OAuth tokens.

3.3.3.5 Access Control

Access control decisions made within the data access framework involve the aggregation of an Access Controller with a Service Provider. The Service User makes its request (1) to the Service Provider. It checks with the grouped Access Controller, which somehow is provided with access to information about user authorizations, identity assertions and other metadata to verify that the data in the response can be requested and returned (2) before it is in fact returned (3).

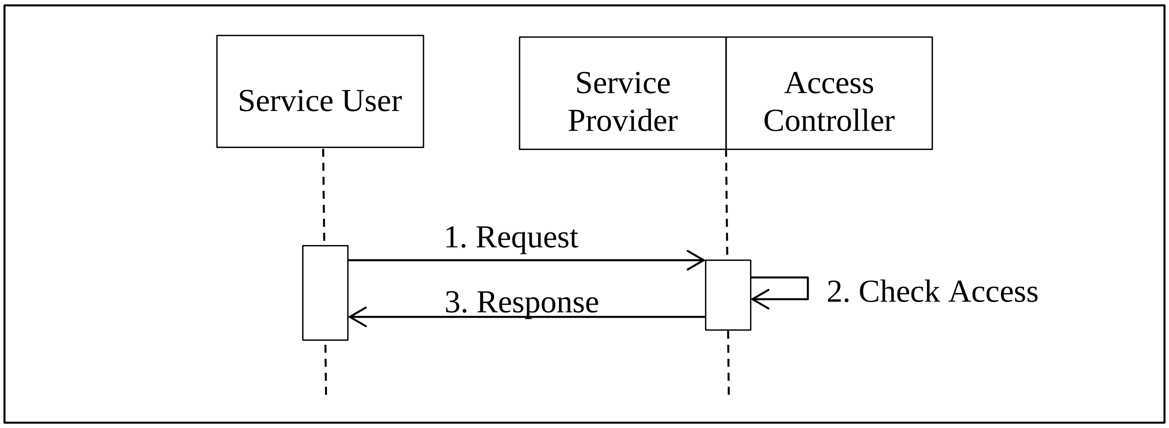

Figure 3.3.3.5-1: Access Control

There is no single pattern of access checks that work for all service users, as various combinations of requests and results may be subject to access control decisions at different access control points. Some responses may not be authorized (e.g., through patient consent) to be returned, while others that are allowed may be returned. This may result in an incomplete set of possible responses to a request.

Communication of access control decisions may or may not be desirable, as communication that “access is not permitted” to a request for sensitive data may provide a covert channel that would communicate that sensitive data does in fact exist for a patient.

In other cases, the responses themselves may be insufficiently anonymized (e.g., as in population queries where a small number of results could lead to patient identification). Thus, access control is expected to be integrated with a responder. Figure 3.3.3.5-2 Loosely Coupled Access Control Implementation below is taken from the IHE Access Controls white paper. It describes a mechanism by which access controls can be loosely coupled with Service Users and Services providers providing the recommended access control architecture. However, it does not specify the mechanisms or standards by which access control policies, decisions, or metadata required to make decisions are communicated between these systems.

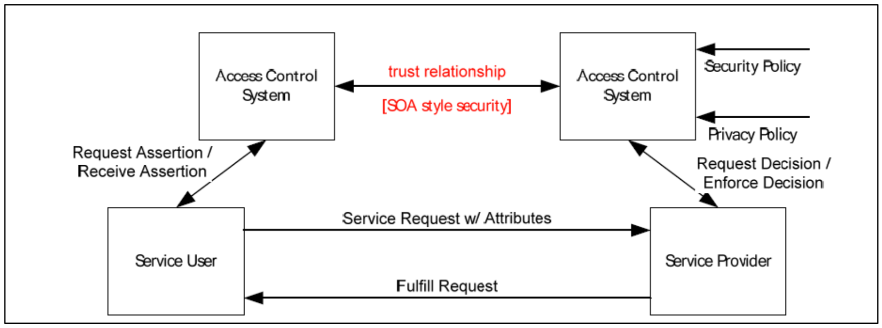

Figure 3.3.3.5-2: Loosely Coupled Access Control Implementation

3.3.3.6 Patient Consent

Patient Consent is a specific kind of access control decision that allows or denies a request to be fulfilled based on the existence of a positive patient consent (opt-in), or lack of existence of a specific negative consent (opt-out) associated with specific information and access policies.

The layering of consents within a query request/response pattern can be performed in a number of ways. The information about the consent can be provided in an assertion (e.g., a SAML assertion) provided when a request is made.

The IHE Basic Patient Privacy Consent (BPPC) profile provides a content profile for structuring a consent document (defining the policy), describes the metadata that can be shared to identify key features of a consent for making access control decisions, and specifies how access checks (see step #2 in Access Control above) can be resolved.

4 Implementable Specification

Most of the effort in making an IHE profile implementable is through the development of content found in Volume 2 and 3 of an IHE Domains technical framework, with some additional information found in volume 4 to support Regional requirements at the implementation level.

4.1 Substitutability

One of the non-functional requirements of DAF is that it supports data access through multiple protocol stacks in order to maximize interoperability across different systems. These requirements are described in more detail below:

1. Allow for substitutability of transport stacks
   Vendor systems implement different transport stacks. Examples of transport stacks include (HTTP + SOAP, HTTP + RESTful resources, MLLP + ER7 etc.). Interoperability across the transport stacks is a challenge, however allowing for defined ways to substitute transport stacks will provide greater levels of interoperability.
2. Design for a modular Query stack
   In order to enable flexibility in the usage of multiple transport, content and security standards it is important for DAF to modularize the query stack with defined APIs that allow for substitution.
3. Allow for substitutability of content standards
   The content that is exchanged using the various transport stacks vary widely and requires specific bindings based on the transport stack to enable interoperable exchange. The Data Access Framework supports multiple content standards and profiles, and is expected to evolve over time.

The following four protocol stacks were considered during the development of this white paper.

Queries using HTTP and SOAP (SOAP)

Many IHE profiles use HTTP as the transport, along with appropriate SOAP protocols to perform queries. SOAP protocols use XML structure to package content. Within the XML package other structural standards and formats are leveraged to further define data structures. Some of these additional standards that are leveraged include ebRIM, SAML, DSML etc.

Queries using HTTP and RESTful resources (REST)

Some IHE profiles use HTTP as the transport, along with RESTful resources to perform queries. The RESTful resources are using XML or JSON structures to package content. Within the package other structural standards and formats are leveraged to further define data structures. Some of these additional standards that are leveraged include OAuth2 and in the future HL7 FHIR.

Queries using MLLP and HL7 V2 (MLLP)

The Minimal Lower Layer Protocol (MLLP) is a standard for transmitting HL7 messages via TCP/IP. Since TCP/IP is a continuous stream of bytes, a wrapping protocol is required for communications code to be able to recognize the start and the end of each message. The Minimal Lower Layer Protocol is the most common mechanism for sending unencrypted HL7 via TCP/IP over a local area network, such as those found in a hospital. The HL7 messages further use specific delimiters and structures to encode the data within the message.

Queries using SMTP and S/MIME (SMTP)

Currently there is limited use of SMTP (IHE XDM profile has an SMTP option) stack for data access; however the SMTP stack may play a role in asynchronous query implementation. A derivative of the SMTP stack (i.e., Direct) is required for Meaningful Use stage 2 in the US due to which many EMR systems are supporting the STMP stack as one of the protocols to enable push based messaging. The S/MIME standard is used to structure the package in the SMTP stack. The package itself leverages other standards and formats to further define the data structures.

4.2 Regional Requirements

Implementable variances introduced by regionally specific requirements are identified in Volume 4 of an IHE technical framework. Content within this volume is developed by IHE Regional Deployment organizations into National or Region Extensions and is integrated into a domain's technical framework by that domain's technical committee.

Gap: At this point in time, few profiles identified in the Data Access Framework have specified any national or regional extensions.

5 A Data Access Framework

Table 5-1 Implementable Information Models below illustrates the implementable specifications specified in standards and implementation guides which are used by IHE profiles for each of the combinations of data accessed, behavior model for access and network transport. It also identifies gaps where profiles do not yet exist to support these capabilities.

These models are used regardless of the kind of governance applied to the systems exchanging information.

Table 5-1: Implementable Specifications

Data Access	Behavior Model	Network Transport	Implementable Information Model	IHE Profile	
				Patient	Population
Patient Demographics	Request/Response	MLLP	HL7 V2 ADT	PIX/PDQ	Gap[1]
		SOAP	HL7 V3 Patient Administration	PIX/PDQ V3 XCPD (Federated)	
		REST	FHIR Patient	PDQM	
	Publish/Subscribe	SOAP	HL7 V3 Patient Administration	Gap[2]	
		REST	FHIR Patient		
Encounter Documents and Metadata	Request/Response	SOAP	ebXML RIM+CDA	XDS/XCA	MPQ
		REST	FHIR+CDA	MHD	Gap[3]
		SMTP	Gap[10]		
	Publish/Subscribe	SOAP	ebXML RIM+CDA	DSUB	Gap[4]
		REST	FHIR+CDA	Gap[5]	
		SMTP	Gap[10]		
Detailed Clinical Data	Request/Response	SOAP	HL7 V3 Care Record	QED[8]	Gap[6]
		REST	FHIR	Gap[7]	
	Publish/Subscribe	SOAP	HL7 V3 Care Record	CM[9]	
		REST	FHIR	Gap[7]	

Implementable information models are represented in the Message Semantics section describing transactions in Volume 2 or in Content Modules found in Volume 3.

Table 5-2: Data Access Framework Gaps and Proposed Resolution

Gap	Proposed Resolution
1	IHE ITI has not developed profiles to support unrestricted queries for patients matching demographics criteria. Most information systems in a healthcare environment have access to an organizational MPI. This could be a profile submission to the Quality, Research and Public Health Domain. We note that most user stories for queries based on patient demographics also inquire about patients with a particular disease, or presence or absence of diagnostic test.
2	IHE ITI has not published a profile supporting a publish/subscribe model for accessing patient demographics data. This could be a profile submission to the ITI Technical Committee.

Gap	Proposed Resolution
3	The IHE Mobile Access to Health Documents (MHD) Profile does not support population level queries; however, the base standards do support this. This could be submitted as a new profile to Quality, Research and Public Health Domain.
4	This requirement could be met by extending the IHE Document Subscription (DSUB) profile to support the queries specified in the IHE Multipatient Query (MPQ) profile. This might be added as a change proposal or a new profile submission to IHE IT Infrastructure.
5	As written, the IHE Mobile Access to Health Documents profile does not presently support a publish/subscribe model, however this capability is supported by the base standards and could be incorporated as a change proposal or new profile submission to the Quality, Research and Public Health Domain.
6	The IHE QED profile does not support population level queries. A simple extension to this profile could be made to allow for matching without a patient identifier being specified, or a new profile submission could be made.
7	The base standard (FHIR) supports RESTful queries, but has not yet been profiled in IHE PCC. This is on PCC's roadmap for future years. It was not considered for submission in the 2014/2015 development cycle due to the fact that FHIR had not yet reached DSTU status.
8	The IHE QED profile needs to be updated to support the HL7 Version 3 Care Record Standard that is now normative. The current content relies on the HL7 Version 3 DSTU.
9	The IHE CM profile should be updated to support the HL7 Version 3 Care Record Standard, and should also take advantage of HQMF Release 2, which supports identification of data elements that are needed in support of patient care (the HQMF data element model was informed by this profile).
10	The IHE XDM profile could be used with an SMTP protocol as an alternative mechanism to respond to publish/subscribe requests reusing existing IHE transactions. No query model has been developed for SMTP submission.

5.1 Security Considerations

5.1.1 Intra-enterprise

For intra-enterprise queries, the enterprise controlling the Query Requestor and Query Responder Health IT systems will prescribe appropriate security controls based on local policies. These systems can use the IHE ATNA Profile to encrypt and secure information exchanged between systems, BPPC to support patient consent, and IHE EUA, XUA or IUA as appropriate to ensure user authorization and authentication.

The IHE EUA profile can work with any HTTP-based protocol stack. The XUA and IUA profiles can support communication of a SAML assertion within an exchange, either through HTTP or other protocol that supports communication of an assertion or token.

Gap: Communications using the MLLP protocol and HL7 Version 2 standards, such as PIX/PDQ can also support EUA by pre-adoption of the HL7 Version 2.7 UAC segment, which includes the ability to communicate Kerberos ticket information. HL7 Version 2 standards allow segments to be pre-adopted in communications. While this solution has been available for several years, there has been no request to update the profile to support this capability; it seems to be little needed within an enterprise.

5.1.2 Inter-enterprise

Inter-enterprise queries are executed between two specific enterprises belonging to two different security domains. In order for these queries to be executed, appropriate security information (Authentication, Authorization etc.) needs to be included as part of the query request and could be pre-negotiated between the enterprises. These systems can use the IHE ATNA Profile to encrypt and secure information exchanged between systems, BPPC to support patient consent, and IHE XUA or IUA as appropriate to ensure user authorization and authentication. We do not recommend the use of EUA as it is designed for use within a single enterprise.

5.1.3 Federated

Federated queries are executed across multiple enterprises belonging to multiple security domains. In order for these queries to be executed, appropriate security information (Authentication, Authorization etc.) needs to be included as part of the query request and may need to be dynamically negotiated. The nature of the dynamic negotiation may depend on local or regional policies and is also dependent on the trading partners who get added or removed into the eco-system. Many implementations may use policy engines to deal with the above variability. These systems can use the IHE ATNA Profile to encrypt and secure information exchanged between systems, BPPC to support patient consent, and IHE XUA or IUA as appropriate to ensure user authorization and authentication.

6 Conclusions

The Data Access Framework presented in this white paper illustrates how four common conceptual data models (Patient Demographics, Encounter Document Metadata, Clinical Documents and Detailed Clinical Data) are used within IHE profiles. It further demonstrates how a common set of base standards including: ebXML RIM, HL7 Version 2 ADT, HL7 Version 3 Patient Administration, HL7 Version 3 Care Record, and HL7 FHIR can be used to develop a highly consistent data access framework.

Industry experience has shown that IHE profiles using a common conceptual model, such as in the case for PIX/PDQ, PIX/PDQ V3, and XCPD are readily adopted by products implementing the service provider (server) capabilities. For example, many MPI products used for Health Information Exchange support all of the above IHE profiles. The IHE Mobile Access to Health Documents profile has already been prototyped making use of the IHE XDS and XCA profiles using the NIST XDS Reference implementation as the back end information source, with a gateway service providing a Façade implementing the MHD profile. It was designed with this implementation pattern in mind.

7 Recommendations

Originally, IHE profiles using different implementation stacks were not specified, because a single implementation stack provided the most interoperable solutions for systems. However, over the years, various regions had developed requirements to implement different stacks, so an IHE profile that resolved a problem with one stack was revised to support additional stacks. These updates were structured as new profiles, often borrowing much of the same Volume 1 content.

The use of multiple protocols will likely continue as new protocols and standards such as FHIR are developed. Ideally, Volume 1 content in IHE profiles would change little, with the principle exception that linkage to other implementable transactions would be allowed. IHE has not yet systematically addressed this issue across domains. While such a refactoring effort could take place, the challenge to address is how that refactoring would be done to have as little impact as possible on existing profiles, and to support continued development of new profile work in IHE.

In addition to refactoring and organizing existing IHE profiles, the gaps and resolutions identified earlier to satisfy Data Access Framework requirements can be scheduled as part of future IHE work.

Appendix A: Sample Integration Statements

Integration Statements are used by vendors to declare their implementation of IHE profiles. Table A-1 below links each service (query or other function) specified in section 3.1 Business Requirements to the IHE profiles and actors necessary to realize that capability on the client or server side for queries using the request/response pattern.

Table A-1: Data Access Framework Queries for Request/Response pattern using existing IHE profiles

ID	Query	Profile	Client Actor	Server Actor
1	Find Document(s) based on Patient Demographics *Note: To perform this query, the demographics must first be resolved into a patient identifier using PDQ or substitutable profiles.*	PDQ PDQ V3 XCPD* PDQM	Demographics Consumer Demographics Consumer Initiating Gateway***	Demographics Supplier Demographics Supplier Responding Gateway***
		XDS MHD	Document Consumer Document Consumer	Document Registry Document Responder
2	Find Document(s) based on Patient Identifier	XDS XCA*	Document Consumer Initiating Gateway	Document Registry Responding Gateway
3	Find Patient Demographics based on Patient Identifiers.	PDQ PDQ V3 XCPD* PDQM	Demographics Consumer Demographics Consumer Initiating Gateway***	Demographics Supplier Demographics Consumer Initiating Gateway***
4	Find Patient Identifier for Patient Demographics	PDQ PDQ V3 XCPD* PDQM	Demographics Consumer Demographics Consumer Initiating Gateway***	Demographics Supplier Demographics Consumer Initiating Gateway***
5	Get clinical data for a patient based on discrete data elements	QED	Clinical Data Consumer	Clinical Data Source
6	Get Document(s) based on Document Identifier	XDS XCA* MHD	Document Consumer Initiating Gateway Document Consumer	Document Registry Responding Gateway Document Responder
7	Get Document(s) for multiple patients	MPQ	Document Consumer	Document Registry
8	Identify Patient(s) based on discrete data elements	QED**	Clinical Data Consumer	Clinical Data Source

*Federated

**In QED, the patientId must be specified. The profile says nothing about the use the wildcard identifiers, but these could be used to support this capability.

*** This is still being defined by the IHE Technical committee.

Table A-2 below links each service (query or other function) specified in section 3.1 Business Requirements to the IHE profiles and actors necessary to realize that capability on the client or server side for queries using the publish/subscribe pattern.

Table A-2: Data Access Framework Queries for publish/subscribe pattern using existing IHE profiles

ID	Query	Profile	Client Actor	Server Actor
1	Find Document(s) based on Patient Demographics *Note: To perform this query, the demographics must first be resolved into a patient identifier using PDQ or substitutable profiles.*	PDQ PDQ V3 XCPD* PDQM	Demographics Consumer Demographics Consumer Initiating Gateway***	Demographics Supplier Demographics Supplier Responding Gateway***
		DSUB	Document Metadata Subscriber Document Metadata Notification Recipient	Document Metadata Notification Broker
2	Find Document(s) based on Patient Identifier	DSUB	Document Metadata Subscriber Document Metadata Notification Recipient	Document Metadata Notification Broker
3	Find Patient Demographics based on Patient Identifiers.	Gap[2]		
4	Find Patient Identifier for Patient Demographics	Gap[2]		
5	Get clinical data for a patient based on discrete data elements	CM	Care Manager	Clinical Data Repository
6	Get Document(s) based on Document Identifier	N/A as there is no need to use Pub/Sub to subscribe to a single document		
7	Get Document(s) for multiple patients	MPQ	Document Consumer	Document Registry
8	Identify Patient(s) based on discrete data elements	CM*	Care Manager	Clinical Data Repository

* In CM, the patientId must be specified. The profile says nothing about the use the wildcard identifiers, but these could be used to support this capability.

Table A-3 below links each security and consent capabilities specified in section 3.1 Business Requirements to the IHE profiles and actors necessary to realize that capability on the client or server side

Table A-3: Data Access Framework Security and Consent capabilities using existing IHE profiles

ID	Query	Profile	Client Actor	Server Actor
1	Supply and Consume User Assertions.	EUA XUA IUA	Client Authentication Agent X-Service User Client Authorization Agent	Kerberos Authentication Server X-Service Provider Resource Server
2	Capture Patient Consent	BPPC XDS	Content Creator Document Source	Content Consumer Document Repository
3	Authenticate Node	ATNA	Secure Node	Secure Node
4	Authenticate User	ATNA	Secure Node	Secure Node

5	Encrypt Communication	ATNA	Secure Node	Secure Node
6	Audit Logging	ATNA	Secure Node	Secure Node

To build an Integration Statement for DAF, decide upon which of the above capabilities your application wants to support.

Table A-4: Sample Integration Statement for DAF

IHE Integration Statement	**Date**	**12 Oct 2005**
Vendor	**Product Name**	**Version**
Any Medical Systems Co.	DAF	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**

Indicate the query capabilities supported as a client. The numbers to the left of each boxed row identify which business requirement the collection of profiles and actors support.

Indicate the profiles supported as the client.

1	Patient Demographics Query V3	Demographics Consumer	none
	Cross Enterprise Document Sharing	Document Consumer	none

Indicate the profiles supported as a server.

1	Patient Demographics Query V3	Demographics Supplier	none
	Cross Enterprise Document Sharing	Document Registry	none
	Document Metadata subscription	Document Metadata Notification Broker	none

Add rows to indicate the profiles supporting security requirements

1	Cross Enterprise User Authorization	X-Service User	None
	Audit Trail and Node Authentication	Secure Node	

Add the links to general information.

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		
In North America: www.ihe.net	In Europe: www.ihe-europe.org	In Japan: www.jira-net.or.jp/ihe-j

Appendix B: Sample User Stories

The User Stories represent real world examples of the data access framework. This section contains example user stories to illustrate the specific instances of the Data Access Framework use cases. By design the Data Access Framework is expected to support multiple user stories, many now unforeseen, and therefore the Appendix does not attempt to enumerate all possible uses.

B.1 Document metadata based access - Patient Level Query

A Provider accesses clinical summary documents on an ad hoc basis for a new diabetic patient with documented poor glucose control

A new patient arrives to a small family practice in Boston, MA. The PCP sees a 48 year-old male, with Diabetes Mellitus Type I (DM I) diagnosis since age 12. The patient has a history of myocardial infarction (MI) at age 37 and a stroke at age 43. The patient admits that he often forgets to take his medication as prescribed and often forgets to check his blood sugar levels throughout the day. The patient travels for work and has been admitted to different ER's numerous times for acute complications due to elevated blood sugar levels. All healthcare facilities where the patient was admitted generated clinical summaries and sent the information to patient's new physician at the patient's request. The clinical summaries have been stored in the local document repository database within the organization. ***For today's visit, the physician's practice generates an ad-hoc query in preparation for the patient's arrival within the EHR to access all clinical summary documents produced locally and those received from other healthcare facilities, so that he can check if the patient's HbA1c levels were greater than 7% and if the glucose levels were greater than 100mg/dL over the past 5 years the EHR system queries the document repository database to retrieve the requested information and sends back multiple clinical summary documents to the physician for additional review.*** This information provides the physician required context to understand the severity of circumstances that led to the patient's ER admission, the severity of the patient's non-adherence to medications and formulate a plan to improve the patient's lifestyle and adherence to medications to mitigate future ER visits and reduce or prevent the progression of established comorbidities.

B.2 Document metadata based access - Patient Level Query

A provider needs to access information for one of his patients' who recently moved to a new state and that has a new care team.

A patient is moving from Michigan to Florida for retirement. The patient has diabetes and has also undergone multiple open heart surgeries to correct irregular heartbeats and other ailments related to the heart. His new care team in Florida is preparing for an initial visit and has requested the patient to retrieve his medical history from as many sources as possible. The patient approaches the Michigan hospital, the PCP and the cardiologist office who are part of the current care team and where he has received treatment before. ***He requests each one to provide his medical records (clinical documents) to date. The providers query each of their local EHR***

systems to obtain the clinical documents, requested by the patient. Now that the patient has all necessary records, he can carry them with him on his initial visit to a new care team in Florida.

B.3 Data Element based access - Patient Level Query

Physician referral to Endocrinologist within the same organization using different EHRs with system alerts for patient protected information

In accordance with best practice, the Gastroenterologist orders fasting glucose lab tests for new or current Hepatitis C patients. ***The Gastroenterologist's EHR receives results from source systems based on queries*** which are set up to run automatically, and alerts him when a patient's fasting glucose lab results are between 100 mg/dL and 125 mg/dL. During an initial encounter with a VA patient for Hep-C, the Gastroenterologist is alerted that the patient's glucose intolerance lab results are very high. The Gastroenterologist wants to refer the patient to an Endocrinologist in his practice. ***In preparation for the referral, the Gastroenterologist queries the repository for all of the patient's records including sensitive records disclosed to him by the VA per the patient's consent. The Gastroenterologist receives a response to this query*** and is alerted that information related to the patient's Hep-C, which was diagnosed during substance abuse treatment, is protected under Title 38, and may not be disclosed without patient consent. Before making the referral, the Gastroenterologist asks the patient whether she consents to disclose protected information to the Endocrinologist. The patient agrees, and signs an electronic consent directive. The Gastroenterologist's EHR updates the security labels on this patient's protected information authorizing the Endocrinologist to query for her records. ***When the Endocrinologist's EHR system queries Gastroenterologist's EHR, it is authorized to receive the patient's records including information protected by regional policies. When researchers within the Endocrinologist's practice query for Hepatitis C patients, they will not receive the results for patients who have not consented to disclosure for research, because they are not authorized.***

B.4 Document metadata based access - Population level Query

PCP searches for office visit summaries in local EHR system to further analyze them using 3rd party software system (external to EHR) to understand severity of illness in patient population

A primary care physician's patient panel has a significant number of male patients who have cardiovascular disease and diabetes over the past 5 years. She wants to further analyze the clinical summaries of her male patient population over the past 5 years using a 3^{rd} party analytical application external to the EHR System. She queries her EHR system to retrieve clinical office summary visit documentation for patients over the past 5 years. The results of the query are returned to her in a structured document format for each of the patients fitting those criteria. Once she receives the results, she further analyzes the summaries by using an external 3^{rd} party analytical application to break down cohorts of those patients with mild, moderate, and severe disease to determine who are missing recommended preventive and disease management services such as lab checks and diabetic foot exams.

B.5 Data Element based access - Patient Level Query

PCP searches for office visit summaries in local EHR system to further analyze them using 3rd party software system (external to EHR) to understand severity of illness in patient population

A primary care physician’s patient panel has a significant number of male patients who have cardiovascular disease and diabetes over the past 5 years. She already has a list of male patients and their clinical office visit summary documents that she was able to retrieve through a previous query search in her EHR. She wants to use that list of patients now to drill down within each of these documents to identify patients with cardiovascular disease and diabetes over the past 5 years. The PCP sends one query to her EHR system for all identified patients to retrieve patients with diagnoses of cardiovascular disease and diabetes over the past 5 years. The query returns a list with associated documents that match the query request. Once she receives the results, she further analyzes the summaries by using an external 3rd party application to break down cohorts of those patients with mild, moderate, and severe disease to determine who is missing recommended preventive and disease management services such as lab checks and diabetic foot exams.

B.6 Element based access – Patient Level Query

PCP querying lab data results over past 12 months for a patient whose HbA1c is >7%

A Primary Care Provider (PCP) at Virginia Family Medicine Center (VFMC) recently ordered an HbA1c test for a new patient with established Diabetes Type 1 diagnosis. The patient had been to VFMC several times before, but just recently switched her PCP internally at VFMC. The PCP received the test results for a specimen drawn on 7/5/2013 in her EHR system indicating that the patient’s HbA1c was 8.3%. Her PCP would like to determine her patient’s glucose level trend over the past 12 months. The PCP formulates a query in her EHR system to retrieve all HbA1c results where the patient’s levels were above 7% at VMFC. The PCP receives a single response of available results from one or more responding application(s) where this data was documented. The PCP is able to obtain all of the results requested from the responding application(s). Upon receiving the results, the PCP confirms that the patient’s glucose levels have been progressively increasing based on available results for each visit since 7/5/2012. The PCP then schedules a set of diagnostic tests to aid her in developing an effective rehabilitation plan to proactively manage her patient’s health condition.

B.7 Document metadata based access - Patient Level Query

Two applications share data during a hospital visit to coordinate information about diagnoses, medications and treatments and queuing of appropriate patient education and instruction material. (Debbie Foss Submitted on Wednesday September 5th, 2013)

A patient enters the hospital for pneumonia. During his visit, he is diagnosed with CHF. Patient instruction located in Application X queries the information from Application Y and receives patient demographics and admitting diagnosis, triggering a preliminary list of education topics for introduction to pneumonia and medications for in-hospital teaching. Application X then receives (either via query or as and alert) for the CHF diagnosis, and begins to queue topics for daily teaching on a new diagnosis, new medications and diet. Prior to discharge, Application X queries Application Y -- perhaps seeking a C-CDA in whatever state of completion it's available -- and topics for discharge instructions are triggered for compilation by providers.

B.8 Data Element based access - Population level Query

Physician conducts ad hoc query to determine percent of Hepatitis C patients for research at an organization under treatment with no fasting glucose lab tests (EHR to CDR)

A new physician starts working at a health center where many patients with Hepatitis C are treated. The physician is aware of clinical practice guideline that specifies that patients with Hepatitis C diagnosis on active treatment must have fasting glucose test performed at the beginning of treatment and at predefined intervals during the treatment. The physician wants to conduct research on the quality assessment of patients being treated. The physician sets up a query to first identify all patients with a diagnosis of Hepatitis C and currently receiving Hepatitis C treatment that have not had a fasting glucose test since beginning of the therapy. The query is sent from the local EHR system to an identified application(s) (i.e., Clinical Data Repository) to retrieve a list of patient names fitting these criteria. Upon receiving this information back in his EHR system the physician learns that 3% of his Hepatitis C patients currently under treatment have not had their fasting glucose test. The physician then retrieves the list of individual patients who have consented to share their information for purposes of research.

B.9 Data Element based access-Patient Level Query

User Story Revised and Submitted by Nicole Antonson September 12th, 2013 Ancillary to EHR Query and Update (Pull and push)

Dr. Jones admits patient J to the hospital for pneumonia. During patient J’s visit, he is diagnosed with angina. While in the hospital, he is scheduled for angiogram. During preop, the cardiology nurse begins the data entry process into the cardiology system for the patient (e.g., completes assessment form.) The nurse selects the patients name and the cardiology system initiates a query to the EHR for demographic and patient profile data (e.g., problems, meds and allergies.) The EHR returns the information, the cardiology system uses this information to populate the assessment form, and the nurse completes any missing information through a patient interview. (During the assessment process the same information returned is used for decision support and reminders.) During the angiogram, patient J requires angioplasty. Medications are administered during the procedure and new ongoing orders are created. After the procedure is closed, the Cardiology system pushes the administered medications and ongoing medications to the EHR.

Appendix C: DAF Data Requirements

The dataset requirements section identifies the data elements based on the use cases and are described at a conceptual level. The descriptions of the data elements are independent of any particular standard and will serve as the starting point for detailed profile development activities. As a starting point these data elements have been derived from IHE XDS Metadata definitions.

Table C-1: Time related Encounter Documentation

Data Set Selection	Generic Data Element[5]	Generic Data Element Description
Time	Document Creation Time	Date and Time stamp for document creation.
	Service Start Time	The start time the service being documented took place.
	Service End Time	The stop time the service being documented took place.

Table C-2: Patient related Encounter Documentation

Data Set Selection	Generic Data Element[6]	Generic Data Element Description
Patient Data	Patient ID	The identifier assigned by a provider or healthcare organization to a patient (example: MRN)
	Patient Demographics	A set of demographic information about the patient. This information typically includes patient's first and last name, sex, birth date, race, and ethnicity.
	Patient Identifiers	ID assigned to a patient where the care was provided within the local organization, if different from Patient ID.

[5] **Note:** examples of data elements for document metadata based access can be found in the following types of profiles: XDS, CDAR2

[6] **Note:** examples of data elements for document metadata based access can be found in the following types of profiles: XDS, CDAR2

Table C-3: Organization related Encounter Documentation

Data Set Selection	Generic Data Element[7]	Generic Data Element Description
Organization Data	Author institution	Represents a specific healthcare facility where a document was authored.
	Health Facility Information	Information about the organizational setting in which the clinical encounter was documented and where clinical act occurred. This includes (Name of facility, Type of facility, code of facility, ID of facility)
	Source Organization Information	Information about the origin of the document (Name of the Organization, Type of organization Code of organization, ID of the organization)
	Practice Setting Information	Practice setting is the location where clinical care was provided and the document was created. (Name of the practice, Code associated with the type of practice, identifier associated with the type of practice) e.g., Family Practice, Laboratory Department, Radiology Department, Pulmonary Unit, Intensive Care Unit, etc.
	Document Custodian	Organization legally responsible for the document

Table C-4: Document related Encounter Documentation

Data Set Selection	Generic Data Element[8]	Generic Data Element Description
Document Level Data	Document Information	Information about the document where patient information has been recorded (code associated with document type, Name associated with the document type, IDs associated with the document)
	Comments	Comments associated with the Document, free form text.

[7] **Note:** examples of data elements for document metadata based access can be found in the following types of profiles: XDS, CDAR2

[8] **Note:** examples of data elements for document metadata based access can be found in the following types of profiles: XDS, CDAR2

1090

Table C-5: Document Author related Encounter Documentation

Data Set Selection	Generic Data Element[9]	Generic Data Element Description
Document Author	Author Specialty	Represents a specific specialty of the author who created the document. For example, Primary Care Physician, Nurse Practitioner, Anesthesiologist, Cardiologist etc.
	Author Contact Information	Represents the telecommunications address (e.g., phone, email, etc.) of the document author, intended to assist with automated routing of other messages intended for the document author.

[9]**Note:** examples of data elements for document metadata based access can be found in the following types of profiles: XDS, CDAR2

Table C-6: Detailed Clinical Information for DAF

1095 Note: This list includes an initial list of data elements from Meaningful Use summary documents exchanged among providers and/or patients.

Legend

X	DAF Query requests will be created using one or more of these data elements identified in the request parameter column
☐	DAF Query responses will include some or all of the data elements identified in the response parameter column

Data Elements	Request Parameter	Parameter Response Value
Patient (s) Identification	X (Null for population Queries)	☐ (Null for population Queries /De-ID/LDS)
Provider Identification	X	☐
Facility / Source	X	☐
Encounter Type	X	☐
Date (Date Range)	X	☐
Confidentiality Information		☐
Patient name[10]	X (Null for population Queries)	☐ (Null for population Queries /De-ID/LDS)
Sex	X	☐
Date of birth	X	☐
Race*	X	☐
Ethnicity*	X	☐
Preferred language*	X	☐
Smoking status*	X	☐
Problems*	X	☐
Medications*	X	☐
Medication allergies*	X	☐

[10] **Note:** The initial list of data elements are derived from MU2 data elements whose definitions can be accessed here. Data elements numbered 1-18 are from the MU2 data elements.

Data Elements	Request Parameter	Parameter Response Value
Laboratory test(s)*	X	☐
Laboratory value(s)/result(s)*	X	☐
Vital signs (height, weight, BP, BMI)	X	☐
Care plan field(s), including goals and instructions	X	☐
Procedures*	X	☐
Care team members	X	☐
Immunizations*	X	☐
Confidentiality Information	No Confidentiality Code Query Parameter	☐
Clinical Instructions[11]	X	☐
Cognitive Status	X	☐
Date and Location of Visit	X	☐
Dates and Location of Admission and Discharge- Inpatient Only	X	☐
Diagnostic Tests Pending	X	☐
Discharge Instructions- Inpatient Only	X	☐
Functional Status	X	☐
Future Appointments	X	☐
Future Scheduled Tests	X	☐
Immunizations Administered during the Visit*	X	☐
Medication List *	X	☐
Medications Administered during the Visit*	X	☐
Provider Name and Office Contact Information	X	☐
Reason for Hospitalization- Inpatient Only	X	☐
Reason for Referral- Ambulatory Only	X	☐

[11] **Note:** Data Elements in Blue Text have been cited from EHR Certification Criteria and can be found here. Data elements numbered 19-36 are from the EHR Certification Criteria.

Data Elements	Request Parameter	Parameter Response Value
Reason for Visit	X	☐
Recommended Patient Decision Aids	X	☐
Referrals to other Providers	X	☐

Appendix D: IHE Profiles Supporting the Data Access Framework

This appendix identifies the various existing IHE integration profiles that were considered for implementing the DAF requirements.

Table D-1: IHE Profiles under Consideration

IHE Profile	Summary	Applicability to			Notes
		Intra-enterprise	Cross-Enterprise	Federated	
ATNA	Audit Trail and Node Authentication Basic security through (a) functional node access controls (b) defined security audit logging and (c) secure network communications.	Y	Y	Y	
BPPC	Basic Patient Privacy Consents method for recording a patient's privacy consent acknowledgement to be used for enforcing basic privacy appropriate to the use.	Y	Y	Y	
DEX	Data Element Exchange leverages the concept of a metadata registry to add mapping metadata to an annotated data capture form at the point of form design instead of the exchange of data instances.	N	N	N	Provides ability for model correspondence similar to USHIK, so it is useful for dynamic discovery of data elements, so it would not be applicable for DAF, but maybe useful before DAF comes into play.
EUA	Enterprise User Authentication enables single sign-on inside an enterprise by facilitating one name per user for participating devices and software.	Y	N	N	

IHE Profile	Summary	Applicability to			Notes
		Intra-enterprise	Cross-Enterprise	Federated	
HPD	Healthcare Provider Directory supports discovery and management of healthcare provider information, both individual and organizational, in a directory structure.	N	Y	Y	HPD may be useful to discover electronic addresses supporting queries. Becomes relevant in the targeted and federated case to perform discovery because in the first cases the execution environment is static and well known.
IUA	Internet User Authorization provides user authorization for RESTful interface.	Y	Y	Y	
MHD	Mobile access to Health Documents provides a RESTful interface to Document Sharing including XDS.	Y	Y	Y	
MPQ	Multi-Patient Queries aggregates queries to a Document Registry for data analysis such as provider accreditation, clinical research trial data collection or population health monitoring.	Y	Y	Y	
PDQv3 and PDQv2	Patient Demographics Query lets applications query by patient demographics for patient identity from a central patient information server leveraging HL7 v3.	Y	Y	N	PIX/PDQ may be used across organizations based on legal and policy agreements to share patient identities.
PIXv3 and PIXv2	Patient Identifier Cross Referencing lets applications query for patient identity cross-references between hospitals, sites, health information exchange networks, etc. leveraging HL7 v3.	Y	Y	N	PIX/PDQ may be used across organizations based on legal and policy agreements to share patient identities.
PWP	Personnel White Pages provides basic directory information on human workforce members within an organization.	N	N	N	

IHE Profile	Summary	Applicability to			Notes
		Intra-enterprise	Cross-Enterprise	Federated	
QED	Query for Existing Data queries data repositories for clinical information on vital signs, problems, medications, immunizations, and diagnostic results.	Y	Y	Y	
QME-EH	Quality Measure Execution-Early Hearing describes the content needed to communicate patient-level data to electronically monitor the performance of early hearing-loss detection and intervention (EHDI) initiatives for newborns and young children.	N	N	N	
RID	Retrieve Information for Display provides simple (browser-based) read-only access to clinical information (e.g., allergies or lab results).	Y	Y	Y	This is similar to MHD and could be useful to access limited data
XCA	Cross-Community Access allows querying and retrieving patient electronic health records held by other communities.	N	Y	Y	XCA deals with query federation across multiple sources, Between two organizations XDS can provide the same capability; however one could use XCA between two organizations.
XCPD	Cross-Community Patient Discovery supports locating communities with patient electronic health records and the translation of patient identifiers across communities.	N	Y	Y	Between two organizations PDQv2 or PDQv3 provides the same capability as XCPD.
XDM	Cross-enterprise Document Media Interchange transfers documents and metadata using CDs, USB memory, or email attachments.	Y	Y	Y	
XDR	Cross-enterprise Document Reliable Interchange exchanges health documents between health enterprises using a web-service based point-to-point push network communication.	N	Y	Y	This could be useful for asynchronous queries such as Disability determination; this is unlikely to be used locally.

IHE Profile	Summary	Applicability to			Notes
		Intra-enterprise	Cross-Enterprise	Federated	
XDS	Cross Enterprise Document Sharing share and discover electronic health record documents between healthcare enterprises, physician offices, clinics, acute care in-patient facilities and personal health records.	Y	Y	N	
XDS-SD	Cross-enterprise Sharing of Scanned Documents enables electronic records to be made from legacy paper, film, and other unstructured electronic documents.	Y	Y	Y	
XDW	Cross Enterprise Workflow coordinates human and applications mediated workflows across multiple organizations.	N	N	N	Workflow requirements are not in-scope for DAF
XUA	Cross-Enterprise User Assertion communicates claims about the identity of an authenticated principal (user, application, system...) across enterprise boundaries - Federated Identity.	N	Y	Y	

1105 **Glossary**

No new glossary terms.