Integrating the Healthcare Enterprise

IHE IT Infrastructure White Paper

Health Information Exchange: Enabling Document Sharing Using IHE Profiles

Date: January 24, 2012

Author: Karen Witting, John Moehrke

Email: iti@ihe.net

CONTENTS

1 Introduction

The Integrating the Healthcare Enterprise (IHE) standards profiling organization has developed a collection of profiles which can be leveraged for use by healthcare communities for the purposes of document sharing. One of the most significant applications of healthcare information technology is the exchange of health information among disparate clinical information systems and otherwise unaffiliated care providers. Across the world, various communities have developed or are developing methods for exchanging health information among healthcare providers, patients, and other authorized parties. The purpose of this white paper is to provide an overview of the collection of IHE profiles which are intended to be used by communities for exchanging health information. The collection of profiles includes support for patient identification, health document location and retrieval, provider directories, and the protection of privacy and security. This white paper will show how various profiles work together to provide a standards based, interoperable approach to community and cross-community health information sharing.

This document will be annually reviewed for improvements and extensions. Please submit comments to be considered as part of the annual review. Comments can be submitted at http://www.ihe.net/iti/iticomments.cfm.

1.1 Scope

Effective health information exchange involves a diverse set of activities and a broad set of challenges, whether that exchange takes place among affiliated or unaffiliated care providers. The IT Infrastructure (ITI) domain of IHE has addressed many of these challenges by defining a series of integration profiles to address specific aspects of exchanging healthcare information. Each integration profiles addresses part of the broad set of challenges involved in health information exchange. The profiles, however, do not attempt to address governance and policy choices that significantly affect how the profile is adapted in a particular community. This white paper cannot address all such governance and policy issues but will provide some guidance on where governance and policy issues are applicable and offer some common approaches.

It is very important to note that IHE focuses only on interoperability and does not attempt to solve every issue involved in exchanging health information. These solutions are meant to be plugged into an architecture that is designed and executed by the exchange communities themselves. Thus, while each community will generate an architecture that meets its individual needs, the use of IHE profiles will lead to the creation of standards-based communities.

This white paper will focus on explaining the IHE profiles used to address interoperability aspects of document sharing and how they work together to solve common document sharing problems. It is does not assume or describe an implementation architecture. Another IHE White Paper, “Template for XDS Affinity Domain Deployment Planning”, provides support for policy and deployment planning. For application of Document Sharing for particular clinical use cases consider the work of the clinical IHE domains: Anatomic Pathology; Cardiology; Eye Care;

Laboratory; Patient Care Coordination; Patient Care Device; Pharmacy; Quality, Research and Public Health; Radiation Oncology; and Radiology.

1.2 Intended Audience

The assumed audience for this white paper includes those involved in a current or planned health information community of any size and scope that needs an overview of a framework for building a health information exchange model based on open standards. This paper does not cover the technical details as they are found in the IHE Profiles, White Papers, and Webinar material.

1.3 Overview of Health Document Sharing Communities

A health document sharing community (community) exists for the purpose of increasing the accessibility of patient health information across multiple organizations so that clinicians can make more informed decisions about the care that they provide. Today, there are many communities already in production and many more are being planned. The size, nature and scope of communities varies widely but can be characterized by a number of different aspects.

First, some communities are geographically focused while others are not. What often comes to mind when speaking of a community is a regional organization that facilitates information exchange across multiple organizations that are relatively close in proximity. Major metropolitan areas tend to be the focus of these communities, but often a regional community encompasses several rural locales. On the opposite extreme of the geographic aspect of communities is the network of United States Veterans Hospitals. The VA (Veterans Administration) hospitals are spread across the entire map of the US and beyond, yet significant efforts have been spent on being able to exchange data among these geographically separated care centers.

A second characteristic by which to categorize communities is the organizational structure of the community. In some cases, the community consists of a single hospital and several out-patient clinics that have a referral relationship with the hospital. In other cases, a network of competing hospitals, laboratories and private clinics may collaborate to form a community.

A third means by which to describe communities is the scope of the content shared. Some communities have very limited exchange functionality. For instance, a community may focus entirely on electronic lab result delivery or e-prescribing. Most communities define a moderate scope to their exchange activities that might include results delivery, electronic referrals, and perhaps some sharing of encounter-based information (e.g., dictations). More advanced communities leverage their network to include even larger scopes (perhaps including the sharing of documents with the patient's Personal Health Record, exchange of clinical summaries, regional patient centric workflows, etc.). No two communities are alike in terms of the set of exchange activities that they facilitate.

Finally, a fourth aspect of a community is the size, scope and political jurisdiction(s) that regulate it. The simplest community uses only an adhoc arrangement to push documents from

one organization to another. National and sub-national jurisdictions have significant effects on the organization and operations of a community.

Despite all the variance among communities, each has the same ultimate goal: to increase the authorized exchange of patient health information across organizations so that clinicians can make more informed decisions about the care that they provide. This ultimate goal provides the reason why the community exists, it is their affinity.

Once communities are formed there is a need to exchange health documents across the communities as well as within them. IHE uses the concept of cross-community to describe a federation of communities which use mostly peer-to-peer interactions for the purposes of health document sharing. A community may be a single organization, like the USA Veterans Administration, a complex community of many organizations, or a more simple organization like a single small hospital or facility. Cross-community describes an environment where multiple communities, be they simple, small, complex or large, interact without any understanding of or access to the internal structure of any of the other participants. A large federation of communities is exemplified by the multi-national exchange “European Patients – Smart Open Service” (epSOS).

This paper will describe a set of building blocks for health document sharing and each environment will use some set of those building blocks to enable the architecture desired by the community or communities participating.

2 Principles of IHE for Health Document Sharing

This section describes several principles which are foundational to IHE's approach to health document sharing.

2.1 General IHE principles

The following general IHE principles are applicable to the set of IHE profiles used for Document Sharing:

- IHE profiles describe the interactions between systems and not the implementation within systems. Interactions between systems are typically described by transactions which are technically specific and detailed enough to ensure interoperability among implementing systems. The internal implementation of the systems is not prescribed by IHE. For example, for patient demographic matching IHE specified the format of the query and response but not the algorithm or method used for the demographic matching. This allows freedom for implementations to address scalability, creative functionality, reliability, and other value-add.
- IHE profiles are designed to support a wide variety of governance and policies. Because IHE supports adoption of its profiles around the world it is rarely possible to define policies that are applicable in all countries. For this reason IHE profiles are designed with a variety of governance and policies in mind and are therefore applicable to a wide variety of environments. IHE profiles are designed to be policy neutral and support a broad set of governance; before they can be deployed there are many governance and policy issues that the communities must agree on. Examples of governance and policy issues are things like: roles and responsibilities, privacy, signature requirements, authorization, when to publish, what to publish, administrative roles, configuration, service level agreements, clinical pathways, long-term availability, etc.
- IHE assumes there is a general understanding of widely implemented Information Technology Standards. IHE profiles typically leverage underlying technology like XML, TCP/IP, DNS etc. without detailed explanations.

2.2 Document Sharing Governance

IHE enables interoperable sharing of documents but assumes this sharing occurs under a document sharing governance structure agreed to by all parties involved. The governance structure addresses all policy issues necessary to enable document sharing; content format and coding; and other operational characteristics. The IHE profiles are designed to be agnostic to governance and policy, while also being designed to support and enforce those governance and policy choices. The governance may apply only within a small group, such as a hospital and small physician's office, or may apply at a large level, like an entire nation. In fact, sometimes temporary or informal governance (e.g., via phone call) based on understanding of existing laws or customs is used for exchange among participants. Typically, in order to allow for effective and efficient interactions, the governance structure is formalized through some legal mechanism. Overlapping governance is common, where one set of agreements exist in the region and a

different set of agreements exist across the nation, yet most organizations will eventually want to exchange documents regionally, nationally and internationally.

In addition to general governance agreements, a document sharing community should address the following issues:

- **Format of document content**: To enable interoperable transfer of documents the receiving side must understand the format and structure generated by the sending side. Typically there is an agreement on a set of document formats which must or may be supported. This could include unstructured content like PDF or text documents. Or a more structured format like CDA or a specific implementation guide applied to CDA for a particular purpose. The key is to ensure that whatever type of content is shared, the receiving system is able to interpret the content in an appropriate way, either through human review or machine processing.
- **Coding within documents**: Structured documents often include coded data derived from a given coding system. Agreeing on which coding systems to use for which data is often covered by an implementation guide for the structured document. Agreeing to an implementation guide, or a general guideline for coding systems to use, is necessary to enable semantic understanding of the document received.
- **Coding of metadata**: Metadata are data that provide information about one or more aspects of the document. In the case of IHE-defined document exchange, specific metadata are coded within the structure of the content being exchanged. See section 2.6 where the metadata defined by IHE are introduced. Some of that metadata have values chosen from a coding system defined by the governance of the sharing community. Because IHE profiles can be applied in many parts of the world where coding systems are different, IHE has not specified which code sets to use and this decision must be made among the systems exchanging documents.

The purpose of this aspect of governance is to enable semantic interoperability among participating partners. When the Cross-enterprise Document Sharing (XDS) profile is used the governance is provided through the XDS Affinity Domain, see section 3.2.

2.3 Distinction between Documents and Messages

The HL7 standard for Structured Documents Section 1.2 describes the document vs. message distinction as follows "A document is designed to be persistent for long periods of time, whereas messages are more often expected to be transient. There is a place for both of these constructs in healthcare." HL7 characterizes a document by the following properties:

- *Persistence* – Documents are persistent over time. The content of the document does not change from one moment to another. A document represents information stored at a single instance in time.
- *Wholeness* - A document is a whole unit of information. Parts of the document may be created or edited separately, or may also be authenticated or legally authenticated, but the entire document is still to be treated as a whole unit.

- *Stewardship* –A document is maintained over its lifetime by a custodian, either an organization or a person entrusted with its care.
- *Context* - A clinical document establishes the default context for its contents
- *Potential for authentication* - A clinical document is an assemblage of information that is intended to be legally authenticated.

Health messages, on the other hand, are not expected to be persistent, but represent a unit of information at a moment in time where the context is often implied by the transaction partners. The content is not always whole, where context may exist in the messaging environment rather than inside the message itself. The distinction between message and documents can get blurry at times, as messages sometimes can be persisted and can contain all necessary context. In fact, messages can be converted to documents and can carry documents within their content. But documents are expected to be persistent, relevant over time and having the same meaning regardless of environment. And messages need not be any of those things.

2.4 Longitudinal Patient Record

Building on the document concepts described above in section 2.3 of persistence, wholeness, stewardship and context we can identify the principle of the longitudinal patient record which is foundational and central to health document sharing. Document Sharing Communities are patient centric; and the patient identity is associated with every document shared

Care providers, which may support a broad variety of healthcare facilities: private practice, nursing home, ambulatory clinic, acute care in-patient facility, etc., are typically the sources or creators of health documents. Typically a patient will go through a sequence of encounters in different care settings over the course of their lifetime. With each encounter there is the potential that a provider will produce a health document that can be shared with the community. Documents shared by the provider and tracked by a centralized registry (see section 3.2) or federation of communities (see section 3.3) form a longitudinal record for the patients that received care among those providers within the community. Longitudinal records therefore are expected to last over the span of many decades, just as the documents that comprise them are expected to have persistence, wholeness, stewardship, context, and potential for authentication. As a health information exchange is adopted it is a common practice to use an historical bulk data load, or comprehensive patient summary to initialize the electronic patient record with data for historical purposes.

Within a care setting Clinical Data Repositories (CDR) or Clinical Information Model Infrastructure databases might be used to enhance Clinical Decision Support as a complement to document discovery. These databases would not be nationwide, but, like EHRs themselves, be local to the patient's care facility. Document Sharing supports interoperability amongst local systems and supports a longitudinal patient record that spans across many local systems potentially using multiple different database systems.

2.5 Use of Documents

IHE Document Sharing profiles are content neutral, meaning that any type of clinical information without regard to content and representation is supported. A document is any

collection of bytes, including proprietary and textual formats. It is expected that a deployment of Document Sharing will restrict the format and content of documents exchanged to those agreed to by the partners in the exchange, as stated in Section 2.2. While the format and content of a document is not restrictively defined, it is expected to be a coherent set of healthcare data that includes enough context to be useful to a practitioner. A document should have the characteristics as described in Section 2.3 namely, persistence, wholeness, stewardship, context and potential for authentication.

IHE Document Sharing profiles assume that a patient identity is associated with every document shared (See section 2.4).

The most common document content standard that is profiled by IHE is HL7 Clinical Document Architecture (CDA). This standard supports the coding of the clinical content which allows for use of the content both for display purposes as well as machine processing. Although IHE encourages the use of CDA as the document content type of choice, it does not restrict the content of a document in any way. Many times a document will be encoded in PDF or simple text (e.g., U.S. Department of Veterans Affairs "Blue Button" program). Images and manifest documents may also be exchanged using the same infrastructure. By defining a document so liberally, IHE enables a common health record sharing infrastructure that is flexible enough to handle the content types agreed to by the partners in the exchange.

IHE and other organizations have integration profiles which define document content for specific, commonly occurring cases. For example, the IHE Laboratory domain has defined an XD-Lab content profile to support sharing laboratory reports. Likewise, the IHE Patient Care Coordination (PCC) domain has defined various content profiles including a Medical Summary (XDS-MS) content profile and an Emergency Department Referral (EDR) content profile. XDS-MS supports a patient's transfer of care from one care setting to another, and EDR supports the situation where a physician determines that a patient should proceed directly to an emergency department for care. In each of these cases it is useful for IHE to profile (define) both the transport and the content of the documents so that true interoperability can more easily be achieved throughout the healthcare continuum.

2.6 Value of Metadata

Another key principle leveraged by IHE Document Sharing is the use of metadata. As defined in section 2.2, metadata are data that provides information about one or more aspects of the document. While a document may be any collection of bytes, IHE defines a collection of metadata about the document that aid its identity, discovery, routing, security, provenance, privacy, authenticity and electronic pre-processing. The set of metadata are defined to facilitate interoperability, so that receiving systems can manage, route and administer documents even if they are unable to interpret the contents of the document. The metadata are defined in such a way that additional metadata, defined outside of IHE, can be sent. Of course, systems not enabled to understand the additional metadata will ignore them, but this capability allows the set of metadata defined by IHE, which is already extensive and robust, to be extended when local needs arise.

Metadata serve multiple purposes. They allow systems to perform:

- automated management of the documents – like assigning priorities or work tasks
- automated patient identification – adding the new information to the correct patient's local record
- support for provenance management – making decisions based on authority of creator of content
- support for episodic searches – by type, date of service
- support relationships between documents
- support privacy/authorization controls – enabling access to content only where appropriate
- support security and integrity controls

Any metadata element may support overlapping purposes but the combination of metadata elements provides a robust understanding of the document and enables automated and manual management of the document without the requirement access to the detailed clinical information contained within the document.

2.7 Document Relationships

The metadata defined in the IHE Document Sharing model encompasses more than just characteristics of documents. In fact, the metadata model is very rich, encompassing the relationships between documents through use of folders, submission sets and associations.

Documents: Each document shared using IHE-defined constructs comes with a collection of metadata which describes the document. The metadata describing the document includes things like: document identifier, patient identifier and demographics, document author, class of document, confidentiality of document, creation time, and events causing creation of document, document format and several more. For a complete list of document metadata refer to ITI TF-3: Section 4.1.7 [see endnote [i]].

Folders: Metadata shared using IHE-defined constructs can also describe folders and document's membership in folders. A folder may be used to collect documents for many purposes, like ease of access or describing a functional purpose.

Submission Set: When documents are published or pushed using IHE transactions they are collected into submission sets to reflect the collection of documents sent at a given moment. Since a submission set reflects a collection of documents it shares some of the same metadata as a document, like patient identifier and author, and adds metadata reflecting the collection like identifier of the source, intended recipient and submission time.

Document Associations: The document sharing metadata supports the description of associations between documents. The associations supported are: append, replace, transform, transform with replace, and signs (i.e., digital signature). The append, replace, and transform associations support representation of document lifecycle events, where a document is associated with documents which are created as part of lifecycle events related to the original document.

2.8 Document Sharing Models

IHE has enabled three distinct Document Sharing Models that share the principles in this section. Because the principles are the same it is relatively simple to implement more than one model to accomplish multiple objectives. The three models are:

- Direct Push – in this model, clinical content in the form of documents and metadata is sent directly to a known recipient, or published on media for delivery
- Centralized Discovery and Retrieve – in this model, a centralized locator is used to discover the location of documents which enables a retrieval of the document from a custodian who has registered existence of the document with the centralized locator
- Federated Discovery and Retrieve – in this model, a collection of peer entities are enabled to query each other to locate documents of interest, followed by retrieval of specific documents.

These models share the common definition of a document and metadata describing documents, folders, submission sets and document associations. Each requires some level of governance structure in order to operate, although there is some difference in the governance needs. For instance, the centralized model requires knowledge only of the centralized locator which can then provide connections with distributed document repositories. For Direct Push and Federated approaches a detailed directory of participating entitles is typically used to ensure that the push or query transactions are sent to the proper place. All include strong support for authenticity and encryption on transport. Privacy requirements vary especially between the Direct Push, where privacy policy is generally determined prior to initiation of the action, and Discovery mechanisms where privacy policy is most often determined prior to responding to the request. So, while the issues that need to be resolved through governance are largely the same, the resolutions will sometimes vary depending on the model chosen.

It is expected that most communities of exchange will start with one of the three forms of document exchange and, if needed, adopt the others later. The addition of a new model to an existing deployment is relatively simple because the IHE profiles are based on common principles.

2.9 Patient Identity Management

The Document Sharing mechanisms enabled through IHE assume that a patient is associated with every documents shared. That patient is described within the metadata describing the document.

In the case of a Direct Push, it is up to the receiving entity to resolve the patient by using the metadata containing identifiers and demographics of the patient. It is preferable to resolve the patient prior to sending documents and using the patient identifier metadata element to unambiguously communicate the patient's identity.

In the Discovery models the document query requires the specification of a patient identifier as known by the query recipient. So, in these models it is necessary to resolve the patient prior to searching for documents. In fact, the query does not carry any patient demographic data beyond the patient identifier.

Resolving the patient is a complex subject made more complex through historic norms, regulations, and business factors. Some regions have a universal identifier, but most regions don't. IHE provides several profiles that aid the resolution of the patient identifier. The profiles are described in Section 4.

2.10 Locating sharing partners

One of the challenges of Document Sharing that is not directly addressed by IHE is the identification of Document Sharing partners. Each Document Sharing model has a different type of need: where a centralized discovery approach requires the identification of the central locater, the peer based push and discovery mode requires identification of each of the peers. This ability to discover sharing partners can be accomplished in many different ways and a clear preference is not yet apparent. The approaches can be broadly characterized as a) locating electronic services which can provide information and b) locating patient specific source of information.

For locating electronic services which can provide information, some approaches currently used in various parts of the world are:

- Local configuration files – many organizations keep a local configuration file or address book which is managed manually whenever a new sharing partner is identified or updated.
- Service Registry – a services registry is sometimes used as a centralized service available to all participants.
- Healthcare Provider Directory (HPD) profile – the HPD profile enables a directory of individual and organizational entities along with electronic services provided by those entities. See section 5 for more information about HPD.

For locating source of information about a particular patient, some approaches are:

- Patient Specific Health Data Locator – the Cross-Community Patient Discovery (XCPD) profile enables a special type of locator which can be used to find entities holding data about specific patients. See Section 4 for more information about XCPD.
- Patient Identity Cross-Reference (PIX) – may be used to find an assigning authority of an organization which has registered patient demographics for the patient.
- Patient Demographic Query (PDQ) – may be used to find an assigning authority of an organization which has registered patient demographics for the patient.
- Cross-Enterprise Document Sharing (XDS) – used to locate documents related to a specific patient, see Section 3.2.

2.11 Security/Privacy

IHE addresses Privacy and Security through the use of Risk Assessment and Management. Each profile is assessed for various types of risks and the profile includes mitigations identified through that assessment in the privacy and security considerations.

IHE includes profiles specific to interoperability of security and privacy. Interoperability profiles are not enough to fully address privacy or security. Privacy and security are enabled and

enforced at many levels of depth including policy, physical environment, procedures, organizational, departmental, functional, and information technology.

IHE provides profiles that support privacy and security audit logging, user and system identification and authentication, access control, encryption, data integrity, digital signatures, and privacy consent management. Security and Privacy and the profiles IHE offers are discussed in section 6.

3 Document sharing profiles

The key actors in health information exchange are the document source actors – those applications or modules that create the document to be shared, and the document consumer actors – those applications or modules that retrieve the document to act on it (i.e., present it to the user, import it into the receiving system, etc.). The strength of the Document Sharing profiles is that they enable effective sharing of data among multiple, disparate systems in a way that minimizes the burden that data sharing imposes on those systems. These profiles may be categorized according to three different data sharing models:

- Direct Push – supports point-to-point push of documents where content is sent directly to the intended recipient found through manual means or infrastructure enabled directory
- Centralized Discovery and Retrieve (XDS Affinity Domain) – a community of sharing partners agrees to use a common infrastructure to enable Health Document Sharing. A document source will publish the existence of documents to a location that is accessible to other systems. Then, document consumers can discover document locations that have been previously published and pull a copy of the document.
- Federated Discovery and Retrieve – content is pulled directly from the content holder who is found through manual means or a directory

The three models are designed to support different use cases. The Direct Push model can be relatively simple but it cannot satisfy all use cases because it relies on the source of documents to know where those documents will be needed. The Discovery models can also handle use cases like:

- Treatment of a new condition where prior conditions may be relevant
- Open Referral, where the patient is allowed to choose the specialist
- Highly mobile patient
- Emergency
- Patient with many medical conditions
- Patient with complex condition

The IHE profiles addressing these models are:

- Direct Push – Cross-Enterprise Document Reliable Interchange (XDR) and Cross-Enterprise Document Media Interchange (XDM)
- Centralized Discovery and Retrieve (XDS Affinity Domain) – Cross-Enterprise Document Sharing (XDS)
- Federated Discovery and Retrieve – Cross-Community Access (XCA)

Figure 3-1 shows the flow of data for each of these models.

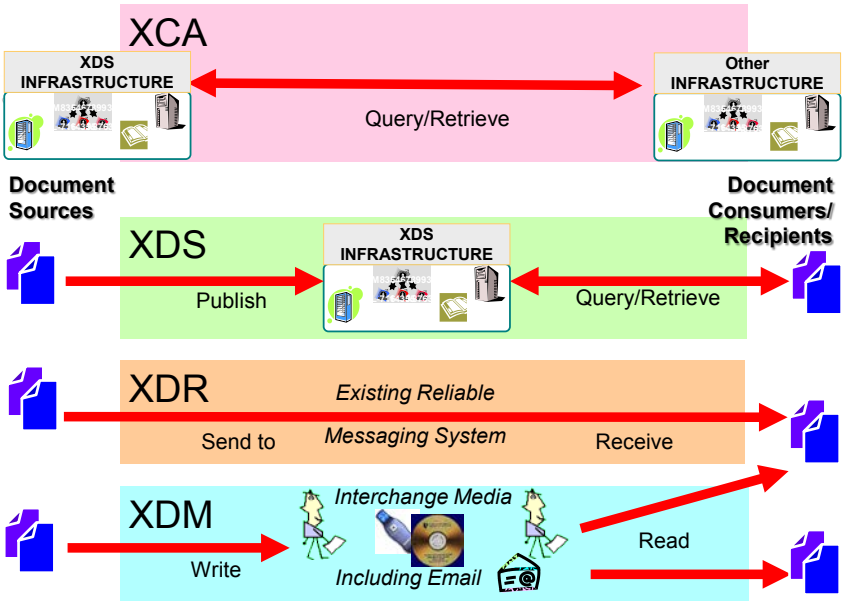

Figure 3-1: Document Sharing Models

Figure 3-2 shows this as a continuum from a simple point-to-point push model on the bottom left to a highly scaled multi-community federated discovery on the top right. Across the bottom are the use-cases we have been discussing and coming from the left are the IHE profiles that address these use-cases.

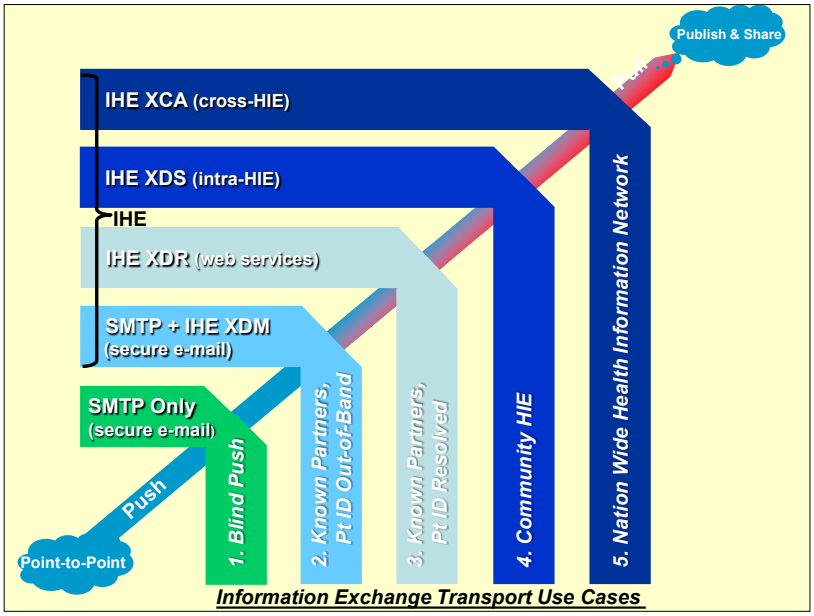

Figure 3-2: Document Sharing Use Case Continuum

The following sections will describe in more detail the three models.

3.1 Direct Push

3.1.1 Cross-Enterprise Document Reliable Interchange (XDR)

The XDR profile provides a point-to-point method of sending documents to a specific recipient. It leverages common principles as described in Section 2. It sends documents and metadata using the same Web-Services reliable transport that is used to publish documents in the Centralized Discovery and Retrieve Document model.

The typical use case for XDR is the patient referral. Dr. Suwati may wish to refer her patient Mary to a specialist, Dr. Lima, who works across town. XDR may be used to send the referral document (and possibly other clinical documents) from Dr. Suwati's Apollo EMR to Dr. Lima's Great Charts EMR.

The Point-to-Point Transmission of Documents slide deck and webinar, [see endnote [ii]], provide more detail about XDR.

3.1.2 Cross-Enterprise Document Media Interchange (XDM)

The Cross-Enterprise Document Media Interchange (XDM) profile addresses situations where the electronic exchange of clinical information does not rely on networked connections between the parties exchanging the information. In these cases, electronic media (such as CDs and USB drives) or email may be employed to transport the data from one system to another.

The XDM e-Mail option is a logical advance for directed e-mail exchange that provides content packaging and metadata to enable accurate processing. The XDM profile has been adopted in national exchange specifications such as the USA defined “Direct Project”.

The XDM CD-ROM and USB-Memory options are logical methods for physically handing the patient records to the patient them-selves; or for delivery via secure courier.

Whether the data is transferred via electronic media or e-mail, the same format is used to express the documents, metadata, and encapsulation structure. The XDM profile can be especially useful when there is no established infrastructure in place between those who have documents and those that need them. Thus the XDM profile can be used in environments where much of the Governance is managed manually, out-of-band. The receiver of an XDM exchange does need to be robust to high variability of the content due to this lack of automated Governance.

Again, the patient referral use case is a typical one that may employ XDM. Dr. Suwati may wish to refer her patient Mary to an orthopedist who does not have an electronic endpoint to receive the referral documentation or to an orthopedist of Mary’s choice in which case the point of service is unknown at the time of the referral. Dr. Suwati decides to write a referral letter and to create an image file of the X-ray of Mary's leg. Dr. Suwati employs her EMR to write this letter and the image file to a USB key using the XDM profile. She then gives the USB key to Mary so that she may take the files with her to the orthopedist.

The Point-to-Point Transmission of Documents slide deck and webinar, [see endnote ii], provide more detail about XDM.

3.2 Centralized Discovery and Retrieve

The Cross-Enterprise Document Sharing (XDS) profile enables centralized discovery of health documents and retrieval of those documents from distributed document repositories.

The following scenario describes a typical exchange of clinical information using XDS. Dr. Suwati works for New Hope Medical Partners which provides her with an EMR system. Her patient, Mary Gomez, just explained to the doctor that she was recently hospitalized at Norwalk General Hospital. Dr. Suwati would like to review the medical records that documented Mary's hospital stay. Using her EMR, Dr. Suwati searches for recent documents for Mary Gomez created by Norwalk General Hospital's EHR. Having found several documents (lab results, radiology reports, a discharge summary, etc.), Dr. Suwati chooses first to view Mary's radiology reports. Having read the reports, she discards them. However, Dr. Suwati reads the discharge summary and then saves it to Mary's record in the local EMR.

In this scenario of health information exchange, the primary player (Dr. Suwati) has three principal objectives: find patient records available from external systems, view a selection of those records, and incorporate a select number of those records to her local system. This sequence of actions is repeated continually in the healthcare setting. To address this very common scenario, IHE has created the XDS profile, a method to coordinate the authorized discovery and sharing of medical documents among disparate information systems.

XDS minimizes the burden imposed on the document sources and consumers when sharing documents by establishing the use of two infrastructure components (the document registry and document repositories), which handle most of the effort involved in exchanging clinical data. This separation allows for minimal yet rich metadata to be centrally managed in a document registry while the full clinical details stay protected within distributed document repositories. The IHE profiles enable the automation of discovery and retrieve by more advanced health information systems.

3.2.1 Document Publishing

The document registry and document repositories always work hand-in-hand, the one being useless without the other. It may be convenient to think of the document registry and document repositories like a public library. The document repositories are a library's set of shelves, an organized resting place for books (i.e., medical documents) that are available to library patrons. The document registry is the library's card-catalog, a tool for locating specific books that lie on those rows and rows of shelves. Unlike a library, the bookshelves are potentially deployed within each participating organization; thus the books are controlled by the original organization until the moment that another organization requests a copy.

It is the responsibility of the publisher to put the books (documents) on the shelf and provide the information for the card-catalog (metadata). The library will step in to update the card-catalog with the data needed to find the new book. In XDS jargon the publisher is called the document

source, whereas the act of putting the book on the shelf and then cataloging it is referred to as "provide and register." Thus, the document source sends a copy of medical documents and associated metadata to the document repository, and the document repository subsequently sends the metadata to the document registry (see figure 3.2-1).

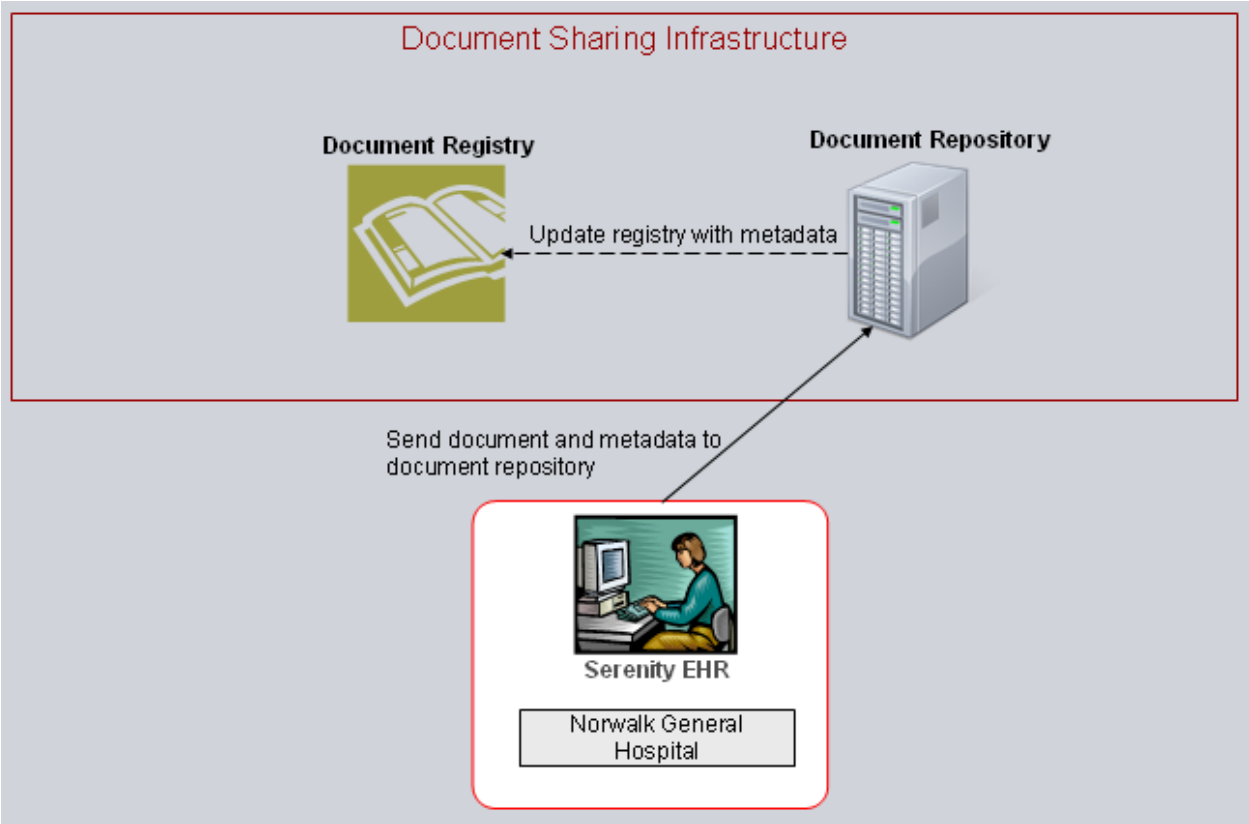

Figure 3.2-1: Provide and register document

The actual location of the document repository will depend on the local deployment. IHE provides flexibility to enable many different deployment approaches.

- The document repository may be combined with the document registry, allowing for an integrated environment where no external “update registry” transaction is needed.
- The document repository could be combined with a document source allowing a large hospital system to enable its local EMR system to also act as a document repository. In this case there is no externally recognized “provide and register” transaction, but simply the “update registry” transaction from the hospital system to the central document registry.
- There is no restriction on how many document repositories can be associated with a single document registry.
- There are no constraints on where a document repository is hosted, the decision is based on many implementation considerations. For instance, a hospital may want to keep its clinical

content local in which case it supplies a repository hosted locally. Or a small physician office may have no ability to support a repository and will prefer to use a repository provided by an external organization, like a hospital system of an infrastructure only partner.

3.2.2 Document Discovery

To complete our analogy, we must consider the library patron (Dr. Suwati in our case), whose goal is to find specific books. The patron interacts with the catalog; sometimes searching for specific books, other times browsing what is available. Once the locations of interesting books are discovered, the patron fetches them from the shelves. In our XDS drama, the document consumer (our library patron) interacts with the document registry to find medical records of interest. This process is known as the "query registry" transaction. The act of fetching the medical record from a document repository is known as the "retrieve document" transaction. Of course with the structured and coded metadata, this step of discovery can be highly automated.

3.2.3 Governance

As described in Section 2: *Principles of IHE for Health Document Sharing* section, the XDS profile is document content neutral; uses document metadata that are represented in a structured, standard format; and supports longevity of document storage.

XDS requires a governance structure as described in Section 2.2 and defines the XDS Affinity Domain as the agent for that governance. An XDS Affinity domain is a group of healthcare enterprises that have agreed to work together using a common set of policies and XDS-based infrastructures for sharing patient clinical documents. Some examples are:

- Regional community of care
- Nationwide EHR
- Specialist (cardiology, oncology) or disease-oriented (diabetes) care networks
- Government-sponsored or federation of enterprises
- Insurance provider supported communities

The XDS profile is patient centric thus requires that an XDS Affinity Domain use a single Patient Identification Domain called an XDS Affinity Domain Patient Identification Domain (XAD-PID). A XAD-PID is a patient identifier assigning authority which provides a single unique identifier for each patient for which documents are registered in the document registry. This ensures that, for example, when submitting documents for Mary Gomez the same unique patient identifier is associated with each document for Mary Gomez, and thus a search can reliably find all of Mary's documents by using this single unique identifier. XDS does not specify how the document source and consumer systems find the unique patient identifier assigned by the XAD-PID, but does require that they find and use it. Use of the PIX and/or PDQ profiles are often applied in support of this task, see Section 4.

Further detail regarding deployment of an XDS Affinity Domain may be found in the "Template for XDS Affinity Domain Deployment Planning" IHE ITI White Paper [see endnote i].

XDS is commonly deployed in an integrated environment which makes use of many IHE profiles working together to enable all aspects of a Document Sharing environment. The profiles most frequently deployed in an XDS Affinity Domain environment are:

- Patient Identification Profiles like PIX and/or PDQ. See Section 4.
- Subscription and Notification Profiles like NAV and DSUB. See Section 3.2.4.
- Security and Privacy Profiles like ATNA, CT, XUA, BPPC and others. See Section 6.

3.2.4 Notifications

The XDS profile supports registration of content and the ability to query and retrieve from a centralized service, but there are many use cases where a system (i.e., a clinician) may wish to be proactively notified when a new document is made available. The simplest type of notification is a personal email requesting that the receiver look for new content. Beyond this approach, IHE specifies profiles that further refine the ability to notify or subscribe for notifications.

- Notification of Document Availability (NAV) – supports out-of-band notifications of documents between systems via an email containing a machine process able message.
- Document Metadata Subscription (DSUB) – uses subscription for new documents fitting specified metadata and notification of available documents.

The Publication and Discovery slide deck and webinar, [see endnote ii], provide more detail about XDS and DSUB.

3.3 Federated Discovery and Retrieve

A community, such as an XDS Affinity Domain, is a means for a specific set of related organizations/facilities to exchange clinical information. But care-givers need access to a patient's entire longitudinal health history, regardless of where that historic information was created. The population is mobile, sometimes due to a temporary situation like a vacation, reoccurring like seasonal housing, work related, or some other choice. Therefor many patients receive care outside of their home community and sometimes the care received externally can be very significant clinically. A very specific healthcare related use-case is when a patient is seeking the care of a specialist. So, there is a need to share health information between two communities. The Cross-Community Access (XCA) profile was developed to address this need.

To implement XCA, a community builds two services called gateways through which all inter-community transactions will flow. An Initiating Gateway is used to send queries to other communities, while a Responding Gateway is employed to receive queries and respond to them. Behind the gateways may be a single organization, like the USA Veterans Administration, a complex community of many organizations, or a more simple organization like a single small hospital or facility. The gateways hide the internal structure of each participant in the cross-community exchange.

XCA is based on the discovery and retrieve pattern defined in XDS, but does not require that either community use XDS. Rather than the centralized model of XDS, which aids locating

documents of interest, the XCA model enables a federated approach, where discovery of documents of interest requires a query to each community that might hold such documents. The XCA gateways are the conduit through which these transactions flow.

The extension of XCA beyond XDS based communities is an important characteristic. A community that is not XDS based, such as the aforementioned USA Veterans Administration, can develop services that implement the interface characteristics of XCA Initiating and Responding Gateways. Since XCA defines only the interface characteristics, proprietary networks are able to support Document Sharing without changing their internal architecture.

The Federation aspects of XCA allow for easy expansion of a network to add new gateway participants in a way that has minimal impact on the creators of content, which still publish information locally, and those that need the data, which use the same query mechanism no matter how broadly the query is federated.

The Federation of many communities does create a larger patient identity problem, and thus a federation approach to patient identity is needed as well. This patient identity federation is profiled in the Cross-Community Patient Discovery (XCPD) profile. For more information on this profile see Section 4.3.

The Cross-Community slide deck and webinar, [see endnote ii], provide more detail about XCA and XCPD.

4 Patient identity management

The Document Sharing defined in this white paper is patient centric, meaning that a patient is associated with each document shared. When data related to an individual patient is exchanged among healthcare information systems it is critical to ensure that the participating systems are referring to the same patient. This requirement can be accomplished in several different ways.

One possible way would have each transaction carry enough demographic data to ensure that the partner is able to match the patient through demographic matching with locally held characteristics. The challenges of "enough" demographic data is a difficult problem. It includes issues around demographics changing over time (name changes) and other aspects of demographics matching rules. There is also concern around privacy when unnecessarily transporting patient demographics.

Thus IHE recommends that the identification of the patient be done through patient identifiers in a common or accepted patient identification domain. Thus, prior to the exchange of healthcare information the partners agreed on a commonly known patient identifier to refer to the patient. Essentially any identifier that a patient provides can be used to correlate identities, with a Voluntary Health Identifier (VHID) being a specific example of an identifier assigned outside of treatment. This requirement, however, is often non-trivial and the patient identity management profiles serve the purpose of enabling this aspect of Document Sharing. Some regions and nations have enabled the use of a unique patient identifier that is widely available but many places still need profiles which aid in patient identifier discovery.

Systems participating in Document Sharing frequently use locally assigned patient identifier domains. A patient identifier domain is defined as a single system or a set of interconnected systems that all share a common patient identification scheme (an identifier and an assignment process to a patient) and issuing authority for patient identifiers.

The Patient Identifier Cross-Referencing (PIX) profile supports the linking of patient identifiers from multiple patient identifier domains. The Patient Demographics Query (PDQ) profile supports the ability to query by a set of demographics and get in response a complete set of demographics, usually including patient identifiers in domains of interest.

The Patient Identity Management deck and webinar, [see endnote ii], provide more detail about PIX and PDQ.

4.1 Patient Identity Cross-Reference (PIX)

Most health information systems assign to each patient an identifier (usually a string of letters and/or numbers) that is unique to the patient within only that information system. Thus, Gary Collins may be identified as 3562A at the office of his Primary Care Physician (PCP) and 0320 at his specialist's clinic.

IHE utilizes the concept of Patient Identifier Domains which defines a domain of patient identifiers, like identifiers assigned within a PCP office, assigned by a single authority and an identifier for each assigning authority. For example, the PCP office identifier is unique within the

assigning authority for the PCP. If the PCP's system wants to communicate with the specialist's system about Gary Collins, both systems must be able to know that 3562A assigned by the PCP offices is equivalent to 0320 assigned by the specialist's office, and that neither of those identifiers is equivalent to Garry Collin with an ID of 333 at a local Hospital. This is known as a cross-reference that links the two patient identifiers for Gary Collins.

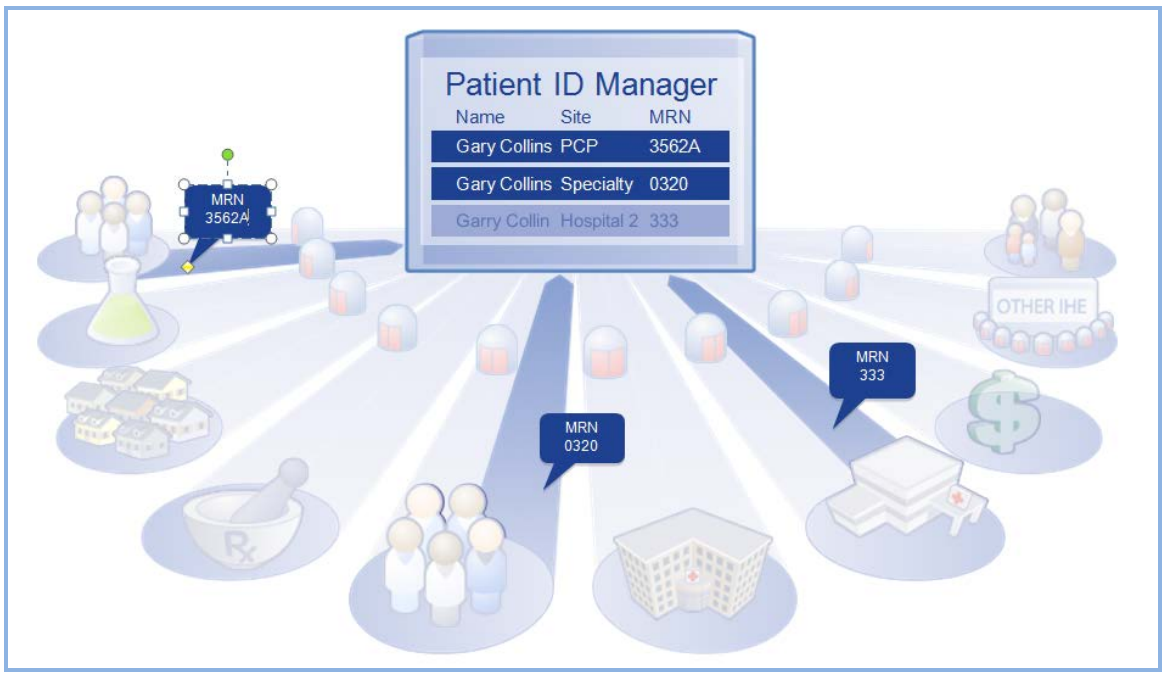

Figure 4.1-1: Patient identifier cross-referencing

The PIX profile is IHE's answer to the difficulty of managing an individual patient's multiple Identifiers. A PIX Manager system receives feeds from multiple patient identity domains, such as the PCP and specialist offices, and uses the demographics in those feeds to create a cross-referencing table which associates identities with matching demographics and does not associate identities found not to match. It should be noted that the PIX profile does not specify how patient matching occurs. Each vendor is welcome to use their own matching algorithms to determine which IDs should be cross-referenced. The IHE profile focuses only on the interfacing characteristics that would be consistent regardless of how the PIX Manager matches the identifiers.

A consumer system may query the PIX Manager to receive a list of identifiers which are cross-referenced with the identifier specified in the query. In this way the PCP office can discover the identifier used by the specialist's system and thus can communicate with that system using a known patient identifier.

A primary use of the PIX profile is to enable document consumers and document sources using the XDS profile to find the patient's identifier in that XDS Affinity Domain Patient Identifier Domain (XAD-PID). See Section 3.2. The PIX profile can be used for Cross-Community if the communities are willing to have a centralized patient cross-reference (See Section 4.3 for federation of patient identity.)

4.2 Patient Demographics Query (PDQ)

Demographics (information describing the patient in general) are used to help identify the patient. With information on dates of birth and sex, information about Leslie Ramsi, a male born on May-2-1968, can be distinguished from that of Leslie Ramsi, a female born on July-23-1987. To help information systems improve their management of patient demographic information, IHE defines a profile called patient demographics query (PDQ). The premise of this profile is that some information systems will have more comprehensive and more accurate demographic information about a patient than other systems. The following paragraph describes a typical use of the PDQ profile.

A typical use of PDQ is to discover the patient's XDS Affinity Domain Patient ID. Imagine that Justin McCarthy heads to the local public health department for a vaccination. The public health department's clinical system does not assign local patient identifiers and thus cannot use the PIX profile to discover Justin's XDS Affinity Domain Patient ID (a required element for the XDS transactions described above). The public health department can use PDQ to find matches for Justin and will receive Justin's XDS Affinity Domain Patient ID as part of the demographics returned. With the knowledge of Justin's XDS Affinity Domain Patient ID, the public health department can now publish his vaccination record to the community via the XDS profile.

4.3 Cross-Community Patient Discovery (XCPD)

The Cross-Community Patient Discovery (XCPD) profile offers a means to discover mutually known patients and a method to correlate the patient's identifiers across those communities.

XCPD uses the same transaction standard as PDQ but adjusts the profiling of that standard in such a way that it is suitable for environments where there is no centralized source of patient demographics or identifiers. For this reason XCPD is most likely to be used with the Federated Discovery and Retrieve (XCA) model of Document Sharing, which adjusts a subset of the transactions of XDS for use in environments where there is no centralized patient record. Thus, XCPD and XCA are designed for environments where the implementation of a centralized source of patient demographics, identifiers or record locations - like is required when using PIX/PDQ/XDS - is not considered acceptable. In environments where communities are willing (and allowed) to feed all patient demographics to a single, central server, or even a small number of duplicated central servers, the use of PDQ or PIX is a more efficient technology to resolve patient identifiers. But as numbers of systems grow, multiple centralized authorities are needed to accommodate the scale and XCPD and XCA are designed to enable communication across multiple such central authorities. Thus XCPD supports a hierarchical approach which bridges communities that might use PIX or PDQ internally.

To illustrate the use of XCPD, imagine that Dr. Holsen has an encounter with his patient, Trudy Levitz. At the moment, Trudy lives in Indianapolis; however, she recently moved there from Chicago. Thus, the majority of Trudy's past medical history is stored in the clinical systems of provider institutions in Chicago. Fortunately, Dr. Holsen's EMR has the ability to discover patient data that exists outside of the local, Indianapolis-based community. Dr. Holsen queries to the Chicago community and finds the relevant patient identifiers from the Chicago community

that represent Trudy. With this information, Dr. Holsen can subsequently use XCA to look for documents containing Trudy's past medical history held within the Chicago community.

The Cross-Community slide deck and webinar [see endnote ii], provide more detail about XCPD.

5 Common Provider Directory

As with patient identity management, the management of data related to healthcare providers (both individual providers and provider organizations) is a fundamental challenge for communities. IHE has defined the Healthcare Provider Directory (HPD) profile to addresses this challenge. There are two principal benefits of using the HPD profile. First, HPD provides a means to disambiguate the identity of providers (i.e., allow one to distinguish between the 58 year-old male pediatric nurse named Lindsay Smith and the 32 year-old female orthopedic surgeon Lindsay Smith). Second, HPD offers a method to discover a provider's contact information (e.g., phone numbers, street address, etc., as well as an electronic endpoint and digital certificate that may be used for trusted communication).

The referral process (one provider referring a patient to the care of another provider) is one of the most common uses of the HPD profile. When Dr. Palov wishes to send his patient Mary Blythe to a female endocrinologist who speaks Spanish, he may query the Directory to find contact information for providers that match those criteria. Similarly, Dr. Palov may wish to refer another patient, Thomas Reed, to the local Mercy Hospital. Dr. Palov could query the Directory to discover the hospital's electronic endpoint (e.g., a secure email address or an XDS repository endpoint) so that he may forward some of Mr. Reed's medical records to the hospital in advance of his visit.

The Healthcare Provider Directory profile describes both how to store data regarding healthcare providers and also how to subsequently access that information. Within the directory, one may also store relationships between providers. For example, Nurse Joe may be an individual provider who belongs to the organizational provider General Hospital.

HPD does not support attributes intended directly for Access Control.

The Healthcare Provider Directory slide deck and webinar, [see endnote ii], provide more detail about HPD.

6 Security and Privacy

This section will discuss how a community that leverages IHE Profiles for document sharing can protect patient privacy and information security. The topic of Security and Privacy is covered in two slide decks and webinars, see endnote [ii].

A very important aspect that is beyond the scope of IHE is the definition of the overall Policies of the community. There is guidance in the IHE Technical Framework, but there is no single policy that must be put in place by an IHE based community to ensure privacy and security. In this section we will discuss potential policy decisions and positions with regard to the profiles. It is very important for the reader to understand that the scope of an IHE profile is only the technical details necessary to ensure interoperability. It is up to any organization building a community to understand and carefully implement the policies of that community and to perform the appropriate risk analysis. Although this section is not going to define the policies that a community should have, it will explore some of the policy building activities to demonstrate how such policies can be supported.

The Policy Environment is made up of many layers of policies. These policies work together in an interlocking hierarchy. We will introduce some of these layers in this section and show how they influence the technology. At the highest layer are international policies, like the International Data Protection Principles. Countries or regions will have specific policies. Some examples are USA HIPAA Security and Privacy Rules, with further refinement by the states. There are horizontal policies that are common among a specific industry, such as those from medical professional societies. Then within the enterprise will be specific information technology policies. As shown in this section, the IHE Profiles offer not only the means to exchange information, but to do so in a way that is supportive of many of the policies mentioned.

The policy landscape that the community is built on needs to be defined well before the community is built.

6.1 Policies and Risk Management

IHE solves Interoperability problems via the implementation of technology standards. It does not ***define*** Privacy or Security Policies, Risk Management, Healthcare Application Functionality, Operating System Functionality, Physical Controls, or even general Network Controls.

While community Policies and Risk Management are outside its scope, IHE does recognize that these elements are a necessary piece of a system implementation. IHE IT Infrastructure technical white paper, “Template for XDS Affinity Domain Deployment Planning” [see endnote i] outlines some of the issues that should be evaluated for inclusion in the local Policy creation and Risk Management decisions. It is therefore the duty of system implementers to take this guidance into account as part of their Risk Management practices.

Implementers need to be aware of different kinds of policies that need to be harmonized with those policies of the local health enterprises connected to the community. The following is a list of sample policy fragments to stimulate discussion:

- Policies for who has access to what type of documents in the community
- Policies for who is allowed to publish documents into the community
- Policies on the acceptable types of documents that can be published into the community
- Policies that indicate acceptable levels of risk within community
- Policies that indicate what sanctions will be imposed on individuals that violate the community policies
- Policies on training and awareness
- Policies on user provisioning and de-provisioning within the community and local operation
- Policies on emergency mode operations
- Policies on acceptable network use (browser, decency, external-email access, etc.)
- Policies on user authentication methods that are acceptable
- Policies on backup and recovery planning
- Policies on acceptable third party access
- Policies on secondary use of the information in the community
- Policies on the availability of the community systems (are the community systems considered life critical, normal, or low priority)
- Policies for maintenance downtime
- Policies for length of time that information will be maintained in the community

These policies are not a flat set, but often interlock and at other times cascade. An important set of policies are those around emergency modes. There are wide definitions of cases that are referred to as emergency mode. These emergency modes need to be recognized for the risks they present. When these use cases are factored in up-front, the mitigations are reasonable.

- Natural or manmade catastrophic disaster (e.g., Hurricane, Earth Quake) – often times additional workforce migrates into the area from other places to help out. These individuals need to quickly be screened and provisioned with appropriate access.
- Utility failure (e.g., electric failure) – this situation is common and easily handled through uninterruptible power supplies and backup generation
- IT infrastructure failure (e.g., hard drive crash) – this situation is also common and handled through common infrastructural redundancy
- Need to elevate privileges due to a patient emergency, often called break-glass (e.g., nurse needs to prescribe)
- Need to override a patient specified privacy block due to eminent danger to that patient – this override is not a breaking of the policy but would need to be an explicit condition within the policy.

Often times being in the emergency department is considered as an emergency mode, but the emergency department is really a normal mode for those scheduled to work there. When looked at as normal mode, the proper privileges and workflow flexibility can be specified.

Policy development often is frustrated by apparent conflicts in policies. These conflicts are often only on the surface and can be addressed upfront once the details of the policy are understood. For example in Europe there are policies that forbid the recording of race, yet this is an important clinical attribute. This superficial conflict might be addressed by recording genetic markers instead of race. Another good example of a policy conflict is in records retention requirements at the national level vs. at the Medical Records level. Medical Records regulatory retention is typically fixed at a short period after death, yet if the patient has black lung then the records must be preserved well beyond.

6.1.1 Technical Security and Privacy controls

In 1980, the Organization for Economic Cooperation and Development ("OECD") developed Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. These guidelines were intended to harmonize national privacy laws, uphold human rights, and promote the free flow of information among its 30 member countries. The OECD guidelines have served as a basis for data protection laws in the United States, Europe, Canada, Japan, Australia, and elsewhere. Together, these principles and laws provide a useful framework for developing general data protection requirements for health information systems. For more information see http://oecdprivacy.org.

Based on the experience of the IHE participants in implementing community environments there is a common set of Security and Privacy controls that have been identified. These controls are informed by a combination of the OECD data protection principles, experience with explicit policies at community implementations, and Security Risk Management.

These security and privacy controls are:

1. Audit Log Controls – The controls that can prove the system is protecting the resources in accordance to the policies. This set of controls includes security audit logging, reporting, alerting and alarming.
2. Identification and Authentication Controls – The controls that prove that a system or person is who they say that they are. For example: personal interactions, Digital Certificates, security assertions, Kerberos, and LDAP.
3. Data Access Controls – The controls that limit access by an authenticated entity to the information and functions that they are authorized to have access to. These controls are often implemented using Role Based Access Controls.
4. Secrecy Controls– As sensitive information is created, stored, communicated, and modified; this control protects the information from being exposed. For example: encryption or access controls.
5. Data Integrity Controls – The controls that prove that the data has not changed in an unauthorized way. For example: digital signatures, secure hash algorithms, CRC, and checksum.

6. Non-Repudiation Controls – The controls that ensure that an entity cannot later refute that they participated in an act. For example author of a document, order of a test, prescribe of medications.
7. Patient Privacy Controls – The controls that enforce patient specific handling instructions.
8. Availability Controls – The controls that ensure that information is available when needed. For example: backup, replication, fault tolerance, RAID, trusted recovery, uninterruptible power supplies, etc. (not an area where Interoperability applies)

6.2 Applying Security and Privacy to Document Sharing

IHE does not set policies but is policy sensitive. Therefore we now discuss the policy enabling technologies and not the policies themselves.

This section shows how the existing security controls in the local health IT system are leveraged and extended when they become interconnected through document sharing.

6.2.1 Basic Security

IHE recognizes that in healthcare, with patient lives at stake, audit control is the primary method of accountability enforcement. The profile that provides this basic security principle is Audit Trail and Node Authentication (ATNA). This profile requires three things of each system:

1. User authentication and Access Controls are enforced accordingly,
2. Security Audit Logs are recorded, and
3. Strong network authentication and encryption for all communications of sensitive patient data

The Security Audit Logging includes a set of security relevant events that must be audited. When one of these events happens the record of the event must be described a specific way. The systems are expected to support the recording of all of the security relevant events that might happen in the system. The ATNA profile offloads the recording, filtering, alerting, and reporting to an audit service. The more centralized this audit log analysis can be, the more easily it is to prove accountability across the whole Document Sharing exchange.

Once it is known that the system will enforce Access Controls and Audit Controls then it can be connected to other systems that have also been assessed positively. In this way these systems only talk to other systems that also agree to enforce the common policies. This creates a basis for a chain of trust through accountability among all of the systems participating in the Document Sharing exchange. The communications between these trusted systems is also encrypted.

For more information on the use of IHE ATNA to enable basic security see the security and privacy slide decks and webinars, [see endnote ii].

6.2.2 Protecting different types of documents

The IHE Document Sharing profiles, like XDS, allow for many different types of documents to be shared. These documents are likely to have different levels of confidential information in them. For instance, one document might contain the very basic health information that the patient considers widely distributable. Another document might be made up totally of information necessary for proper billing such as insurance carrier and billing address. Yet another document might carry the results of a very private procedure that the patient wishes to be available only to direct care providers. This differentiation of the types of data can be represented using a diagram like found in Table 6.2.2-1: Sample Access Control Policies

Sensitivity Functional Role	Billing Information	Administrative Information	General Clinical Information	Sensitive Clinical Information	Research Information	Mediated by Direct Care Provider
HL7 confidentialityCode (2.16.840.1.113883.5.25)	**L**	**N**	**D**	**R**	**V**	**T**
Administrative Staff	X	X				
Dietary Staff		X				
General Care Provider		X	X			
Direct Care Provider		X	X	X		X
Emergency Care Provider (e.g., EMT)			X			
Researcher					X	
Patient or Legal Representative	X	X	X	X		

Table 6.2.2-1 Sample Access Control Policies

Then documents can be labeled with one or more of the codes on the columns, and results in the specified Functional Roles to be given access to that type of document. In this way, the document sharing metadata informs the Role-Based Access Control (RBAC) decisions through self-describing sensitivity, known as confidentialityCode.

In the same way that the XDS metadata 'doctype' defines what the document is in terms of the clinical/administrative content, the confidentialityCode defines what the document is in terms of privacy/security content, sometimes referred to as sensitivity. The confidentialityCodes should be looked at as a relatively static assessment of the document content privacy/security characteristics. Some documents are so sensitive in nature that they simply should not be shared or published.

The rows are showing a set of functional roles. These roles would be conveyed from the requesting organization through the use of the Cross-Enterprise User Assertion (XUA) profile. This profile defines how a user and the security/privacy context of the request is defined. Additional information can be carried such as the purposeOfUse, what the user intends to use the data for. Note that Privacy Policies and Access Control rules can leverage any of the user context, patient identity, or document metadata discussed above.

For more details on enabling Role-Based-Access-Control and federation of identities see the security and privacy slide decks and webinars, [see endnote ii].

6.2.3 Patient Privacy Consent to participate in Document Sharing

The topic of Patient Privacy Consent (Authorization) to collect, use, and disclose is a complex topic. This complexity does not always need to be exposed in full detail across a Document Sharing exchange. That is, a request for information does need to consider the current status of any Patient Privacy Consent that the patient has given, but most of the time explaining the detail of this Privacy Consent to the requesting system/individual adds no value. Most often the requesting system/individual is either fully empowered to receive and use the content, or not authorized at all. In these cases the use of user identity context, as discussed above around the XUA profile, is sufficient to make the Access Control decision. The trust relationship of the Document Sharing exchange includes background governance on appropriate use, as discussed above around the ATNA profile.

Privacy Consents may need to be expressed in a way that all parties in a Document Exchange can understand. IHE has published the Basic Patient Privacy Consents (BPPC) Profile that can be used to enable basic privacy consent controls. At this time further standards are under development by organizations such as OASIS, HL7, ISO, and others. When these standards are complete, patient privacy consents will be more comprehensive and allow the patient to exert far more complex controls than are possible with BPPC. That said, BPPC still provides a rather extensive but coarse-grained level of controls, which may be sufficient in many cases. Some examples of the type of policy that can be enabled by BPPC are:

- Explicit Opt-In (patient elects to have some information shared) is required which enables document sharing
- Explicit Opt-Out (patient elects to not have information shared) stops all document sharing
- Implicit Opt-In allows for document sharing
- Explicit Opt-Out of any document sharing
- Explicit Opt-Out of sharing outside of use in local care events, but does allow emergency override
- Explicit Opt-Out of sharing outside of use in local care events, but without emergency override
- Explicit authorization captured that allows specific research project
- Change the consent policy (change from opt-in to opt-out)

The BPPC profile can be used as a gate-keeper to the document sharing community. BPPC does not define the policies, but does allow for a community that has defined its set of policies to capture that a patient has chosen one or more of those policies.

For example: Let's say that the above set of sample policy fragments was available to a patient sharing in a community. The patient could agree to Opt-In, and also agree to a specific research project. This set of acknowledgments would be captured as one or more BPPC documents. These documents would indicate the policy that is being acknowledged, the date it is being acknowledged, an expiration date if applicable, etc. Then the systems involved in the document sharing can know that the patient has acknowledged these policies and thus the patient's choices can be enforced. A system that is doing research can see that this patient has acknowledged participation in the research project, while other patients have not.

Let's further examine what happens when the patient changes their decision. For example, the patient is moving to a totally different region that is not served by this community. The patient can acknowledge the Opt-Out policy. This policy would then be registered as a replacement for the previous Opt-In policies including the research policy. Thus now if that research application tries to access the patient's data, it will be blocked as the patient does not have a current acknowledgement of the research policy.

6.2.4 Security and Privacy in a Patient Safety Environment

The IHE security and privacy model supports both centralized and distributed control. The entities that are allowed to participate in community based document sharing need to be evaluated to assure that they have the capability to enforce the policies they are expected to enforce. This may mean that access control is enforced at the edge systems, at the center, or more likely in both places.

In healthcare, beyond the basic security principles, we must additionally be sensitive to patient care and safety. The applications closest to the patient are best informed for determining the context of the current situation. It is primarily at this level that emergency mode can be handled in a robust way (often called break-glass).

The IHE security and privacy model is very careful to include security while allowing for flexible and safe provision of healthcare by individual participants.

6.3 IHE Security and Privacy Controls

The following is a breakdown of the security and privacy controls and in what way the IHE profiles can help. The following table shows the set of identified Controls (identified in above) as columns and the supportive IHE Profiles as rows. In this table a '√' indicates a direct relationship. A direct relationship means that the Profile addresses the security and/or privacy principle. An '.'' indicates an indirect relationship, meaning that the Profile assists with the principle. Further details on the '√' direct and '.' Indirect relationships can be found in the profile text or through other webinars.

Table 6.3-1: Profiles relationship to Controls

Security & Privacy Controls / IHE Profile	Profile Issued	Audit Log	Identification and Authentication	Data Access Control	Secrecy	Data Integrity	Non-Repudiation	Patient Privacy
Audit Trails and Node Authentication	2004	√	√	√	√	√	√	√
Consistent Time	2003	√	·				√	
Enterprise User Authentication	2003		√	·			·	·
Cross-Enterprise User Assertion	2006		√	·			·	·
Basic Patient Privacy Consents	2006			·				√
Personnel White Pages	2004		√	√			·	
Healthcare Provider Directory	2010		√	·			·	
Document Digital Signature	2005		√			√	√	
Document Encryption (in development)	2011			√	√	·		

Note: The topic of Security and Privacy is covered in two slide decks and webinars, see endnote [ii].

7 Further Reading

In the paragraphs above, the core set of IHE IT Infrastructure profiles are described in an introductory manner. Specific technical details were purposely omitted from those descriptions since the intent of this white paper is to offer a primer on how to apply IHE ITI profiles to document sharing. For those readers who wish to learn more details, please refer:

1. ITI Educational slides and webinars available at endnote [ii].
2. ITI Technical Framework and supplements available at endnote [i]
3. A white paper that covers deployment planning for an exchange "Template for XDS Affinity Domain Deployment Planning" available at endnote [i]

[i] See http://www.ihe.net/Technical_Framework/index.cfm#IT for published IHE ITI documents.

[ii] See http://wiki.ihe.net/index.php?title=Current_Published_ITI_Educational_Materials for a list of Educational presentations and Webinars related to the topic of this paper.