Integrating the Healthcare Enterprise

IHE IT Infrastructure (ITI) Technical Framework White Paper

Cross-Community Information Exchange
including Federation of XDS Affinity Domains
Version 3.3

October 10, 2008

Table of Contents

1.0 Introduction

As electronic medical records become more prevalent there is an increasing need to share medical data across organizations. The model developing within the United States suggests that medical data sharing will happen first at a local level, as part of Health Information Exchanges (HIEs) and then between HIE's. The Integrating the Healthcare Enterprise (IHE) organization has defined an integration profile called Cross-enterprise Document Sharing (XDS) which defines document sharing within a HIE or, using XDS terminology, within an XDS Affinity Domain. The XDS profile defines a coupling of facilities/enterprises for the purpose of patient-relevant document sharing. This document looks at the issues of how to achieve the sharing of patient-relevant healthcare information among multiple HIE (or HIE-like) environments. This can be seen as the problem of supporting a query which will 1) identify other HIEs which have clinical data about the patient and 2) identify the patient identifier used by the other HIEs for that patient and 3) request patient information from the HIE.

The original version of this paper was published in August, 2006. This revision, 3.0, has been updated to reflect new terminology and technology as well as the existence of the XCA profile.

1.1 Expected knowledge and references

It is assumed that the reader has a working knowledge of three key integration profiles defined within the IT Infrastructure Technical Framework which can be downloaded from:

http://www.ihe.net/Technical_Framework/index.cfm#IT

The key integration profiles and section number in the above document are:

- XDS – Section 10
- PIX – Section 5
- PDQ – Section 8

The reader is also referred to the Cross-Community Access supplement which defines one aspect of sharing healthcare information across communities. This supplement can also be downloaded from:

http://www.ihe.net/Technical_Framework/index.cfm#IT

A brief overview of these profiles is included in this paper as a reference.

2 Goals

This paper addresses the following goals:

- Show a vision for support of communication among XDS Affinity Domains.
- Show a vision for support of communication among HIE's no matter what their internal sharing infrastructure.

The concepts presented in this paper are evolving rapidly as interest and technology adoption grows. The goal is to summarize current activities and set a statement of direction with full expectation that over time this direction will evolve as appropriate. This paper defines common technological building blocks which allow for a variety of strategies and policies to be used. The building blocks are described on a conceptual level only. Specific technology for the building blocks, other than the existing XCA profile, has not yet been chosen.

There are many aspects to sharing data among communities that this paper does not directly address. Things like patient consent, security, privacy, auditability and many more will all need to be considered by implementors. As the building blocks described herein are fleshed out and defined concretely we will ensure that a variety of policies related to these issues can be supported. Some things, like auditing, have already been defined by IHE (see IT Infrastructure Technical Framework Volume 1 referenced above). In those cases the relationship between the existing work and the new work will be explained.

2.1 Request for Feedback

The IHE IT Infrastructure Technical Committee requests feedback on the concepts described in this White Paper. In particular, we would like your thoughts on whether this paper captures the problem as you see it and what do you think of the solution. Comments can be submitted via the web discussion forum at http://forums.rsna.org.

2.2 Open Issues and Questions

- How does an organization declare itself as a community interested in sharing records? Is there a bootstrapping mechanism needed which lists communities? How will credentialing of organizations be handled? Is there some bootstrapping process needed to build trust?

3 Overview

This paper has evolved during several months of discussion and many versions of the paper. The discussion began with an analysis of use cases which is presented in *A.1 Appendix - Use Cases*. Not all use cases have been addressed. One use case was used to develop deeper understanding of the issues and this led to the discussion in *A.2 Appendix - Strategies for solving the selected Use Case*. Further discussion of the strategies and a review of current industry activities bring us to the current version of this whitepaper. The document is organized as follows:

- Define community
- Generalize the many strategies of sharing patient health information into two common types: hierarchical and lateral. Define these concepts.
- Present an approach to cross-community document sharing using existing IHE profiles.
- Describe existing IHE ITI Integration profiles.
- Present common themes in Cross-Community
- Propose future IHE ITI Integration Profiles in support of hierarchical and lateral Cross-Community communication.
- Appendices

4 Community

The XCA profile defines community as follows:

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

Communities can be composed into hierarchical collections of communities we will call meta-communities.

5 Hierarchical vs. Lateral

The many strategies of sharing patient health information can be generalized into two basic types: hierarchical and lateral. In a hierarchical approach some higher level authority is trusted to enable the sharing of data under a set of policies and procedures. Lateral cross-community communication is characterized chiefly by the lack of a hierarchically higher organizational entity to enable policy and technology decision making.

The higher level authority utilized in the hierarchical approach provides services and patient id cross referencing which enable sharing across the meta-community. Organizations and communities agree to participate by making a formal declaration and agreeing to follow the rules. We will call this a meta-community; although the only operational difference between a meta-community and a community would be that a meta-community would have a community within it with a different set of policies from the meta-community. A meta-community is any group of communities which has agreed to work together using a common set of policies for the purposes of sharing clinical information. In defining a meta-community each member would be a community of its own, although some may be so simple as to be a single enterprise or organization. In this sense a meta-community is a collection of communities where some of the communities may be just a single enterprise which can be expected to have well defined policies and data sharing mechanisms.

Lateral cross-community communication is characterized chiefly by the lack of a hierarchically higher organizational entity to enable policy and technology decision making. Thus we use the term lateral when two entities are communicating without any prearranged or declared agreements about how the interaction will be managed. Lateral communication is very common today for healthcare enterprises sharing data. Commonly a manual process is followed where phone or FAX numbers are shared, patient consent is acquired and finally data is transferred. Once two enterprises have collected phone numbers and consent the next transfer may move quicker, but the initial work is manual and sometimes administratively challenging. Thus the policy and technology choices are defined in an ad-hoc manner between every two communicating peers. Given a small set of partners this interchange can be built up as needed. As the group of partners grows the connections grow exponentially and eventually cannot be handled in an ad-hoc manner. Support for lateral communication using manual processes for policy and technology choices is available today from IHE (see XDR and XDM profiles). Automating lateral communication among communities is potentially many years out. This paper discusses the topic of automated support for lateral communication as a long term vision which will be supported by the short term plans.

We expect hierarchical and lateral models to coexist and overlap since they deal with different issues and solutions will be chosen based on the balance of the issues presented to the organization developing its communication model. The figure below shows a model of collecting three communities using a centralized index into a meta-community. That meta-community might then communicate laterally with other communities as show on the right side of the figure.

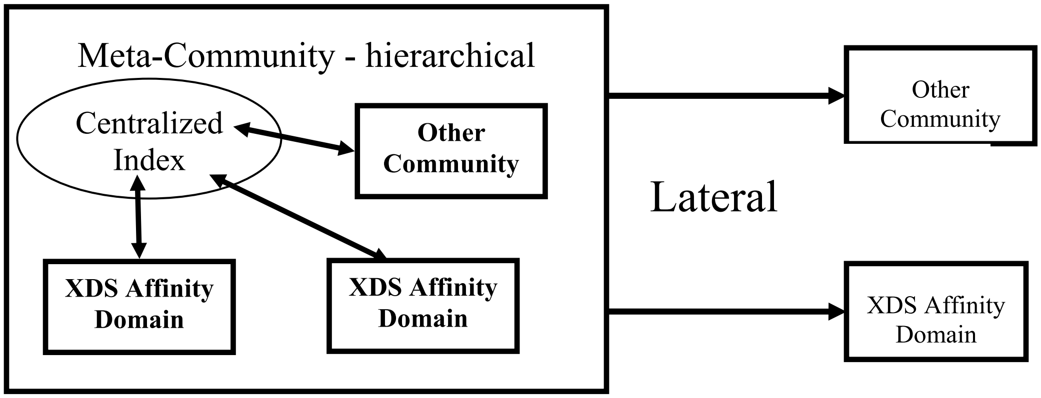

Figure 5-1: Hierarchical and Lateral coexistance

5.1 Meta-community

When a set of clearly defined communities is willing to agree to common policies and mechanisms, composing communities into higher level communities using a hierarchical structure is the most effective mechanism for cross-community communication. A community may participate in more than one meta-community but it would need to handle the resulting privacy considerations. We make no attempt to address the privacy policy issues introduced when a community participates in multiple hierarchical communities.

It is clear that defining the policies under which a meta-community operates is a challenging process. We make no attempt to define those policies in this document, but do ensure that the design supports a broad range of policies. The design must support security and privacy policies defined by governmental agencies as well as a wide variety of organizational preferences.

A significant technical challenge in building a hierarchical grouping of communities is dealing with patient identifiers and patient record locations. Each community will have its own mechanism for handling patient identifiers and a mapping across communities is needed. The communities we are aware of are using some variation of a cross-referencing method for handling patient identifiers. Another alternative would be to assign a global (or global across the communities within the meta-community) identifier for each patient. This “meta-community identifier” is quite far from a national patient identifier, but may trigger a similar debate for large meta-communities. Currently our analysis and experience has been that a patient (or consumer) identifier cross-referencing scheme is used.

All mechanisms for support of the patient cross referencing require cooperation from all communities in the meta-community to feed and update the service or services which perform the cross referencing. This is the main distinction between a meta-community and lateral cross-community communication. A meta-community requires regular patient identification updates by all participating communities to enable the cross-community communication.

Early work on this paper proposed a Patient-Data-Existence Locator which would tie communities together for the purposes of communication. This approach has evolved into the definition of a meta-community which uses the equivalent of a Patient-Data-Existence Locator within a hierarchical organizational structure. It is believed that patient demographic and record location information must be controlled by organizations which have a relationship with the patient. Accumulating this information at the community and meta-community level is acceptable but sharing it beyond that brings up significant concerns regarding privacy and security as well as organizational questions like what entity would maintain and run it, how would that entity be controlled and how would the data be kept up-to-date. The meta-community concept solves all these problems by operating under a set of policies and constraints agreed to by all participants.

5.2 Lateral

Lateral communication requires a process for identifying communities of interest and, for each one found, agreeing on a set of policies and a communication mechanism. Today's medical environment does this manually with significant patient involvement. The patient might supply a FAX number or address for documents being sent.

Our vision for automation of lateral communication among communities (or meta-communities) includes sharing of attributes of a community, searching those attributes and using capabilities to agree on a communication mechanism. Please refer to *9.2 Cross-Community Discovery* for details.

5.3 Summary

The two mechanism under which cross-community communication operates show the variety of policies and practices that must be supported by IHE profiles. The work to analyze and understand these techniques has been used in selecting the first building block and future building blocks defined in Section 9 *Proposed Future IHE profiles.*

6 Existing IHE Profiles useful in Cross-Community Environments

This section lists a few existing IHE profiles which could help support Cross-Community interactions. These profiles are introduced at a very high level. Please refer to the profile for details.

6.1 Cross-Enterprise Document Sharing (XDS)

The ***Cross-Enterprise Document Sharing*** IHE Integration Profile facilitates the registration, distribution and access across health enterprises of patient electronic health records. Cross-Enterprise Document Sharing (XDS) is focused on providing a standards-based specification for managing the sharing of documents between any healthcare enterprise, ranging from a private physician office to a clinic to an acute care in-patient facility. There are two version of XDS, XDS.a and XDS.b. We present XDS.b here, although XDS.a is functionally identical.

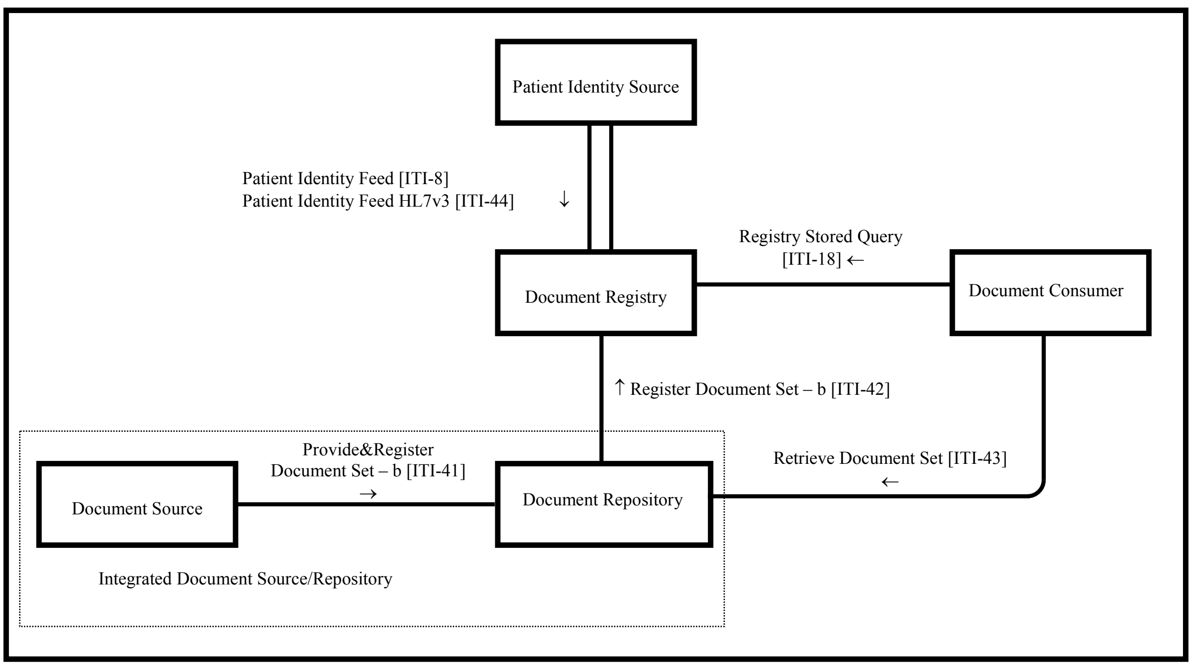

Figure 6.1-1 Cross-Enterprise Document Sharing – b (XDS.b) Diagram

XDS supports:

- The submission of meta-data and documents from a Document Source to a Document Repository
- The submission of meta-data about documents from a Document Repository to a Document Registry
- The query of document meta-data by a Document Consumer to a Document Registry
- The retrieval of documents by a Document Consumer from a Document Repository

6.2 Patient Identifier Cross-Referencing (PIX)

The ***Patient Identifier Cross-referencing Integration Profile (PIX)*** supports the cross-referencing of patient identifiers from multiple Patient Identifier Domains via the following interactions:

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

The following diagram shows the scope of this profile.

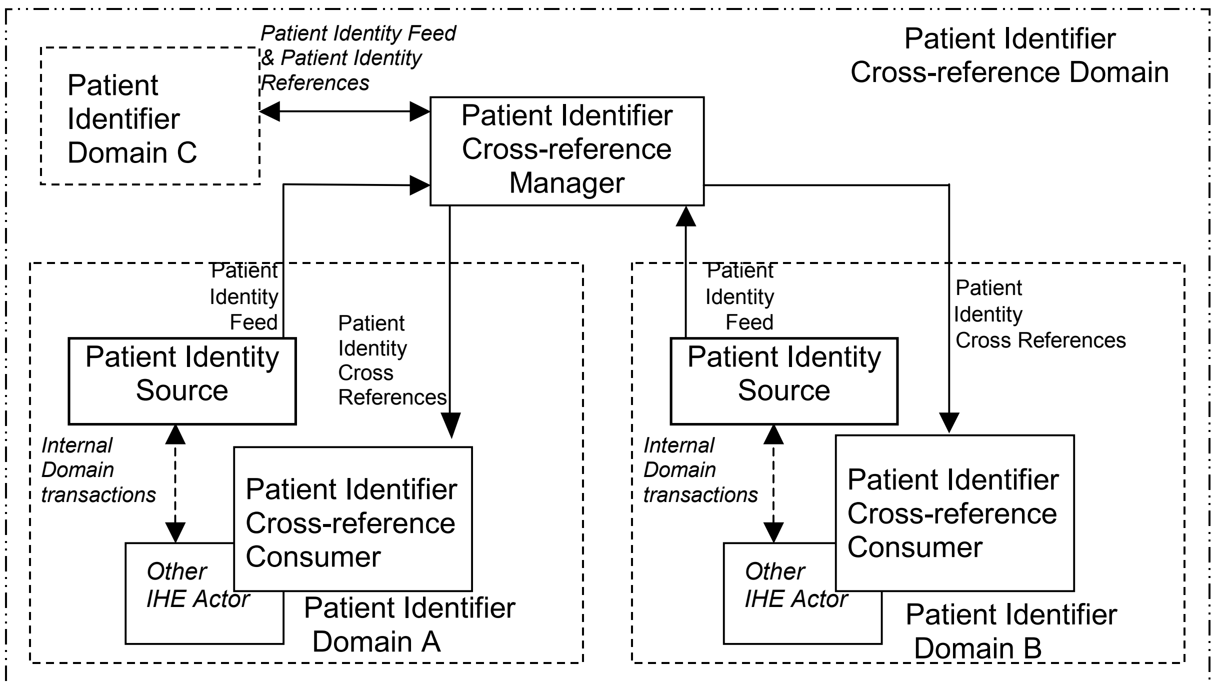

Figure 6.2-1 Process Flow with Patient Identifier Cross-referencing

Figure 6.2-2 shows the actors directly involved in the Patient Identifier Cross-referencing Integration Profile and the relevant transactions between them.

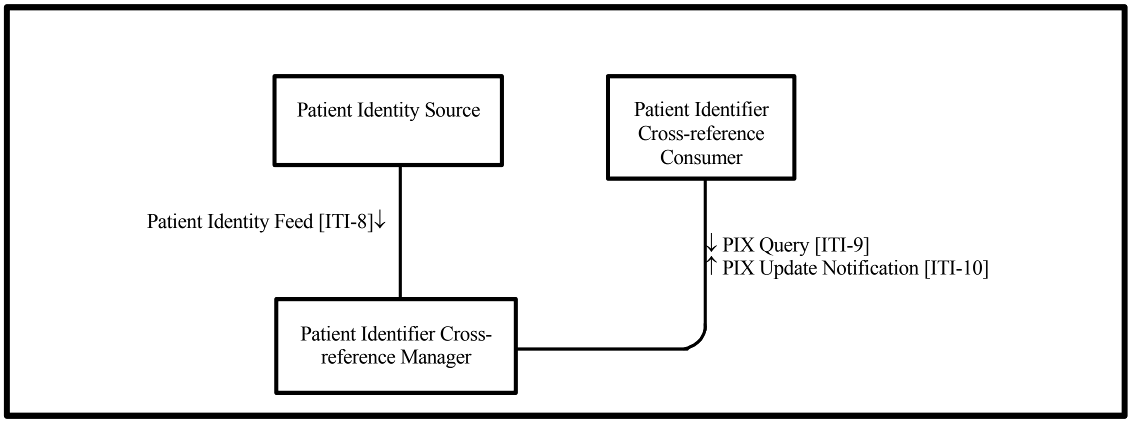

Figure 6.2-2 Patient Identifier Cross-referencing Actor Diagram

6.3 Patient Demographics Query (PDQ)

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

Figure 6.3-1 shows the actors directly involved in the Patient Demographics Query Integration Profile and the relevant transactions between them.

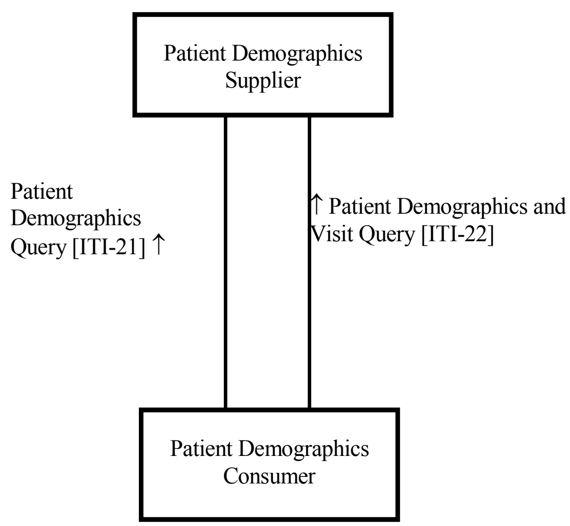

Figure 6.3-1. Patient Demographics Query Profile Actor Diagram

6.4 Cross-Community Access (XCA)

The Cross Community Access profile supports the means to query and retrieve patient relevant medical data held by other communities. The actors and transactions defined by XCA are useful within a meta-community and in lateral interactions.

XCA introduces the concept of a Gateway, which encapsulates all incoming and outgoing cross-community communication. This keeps the management of cross-community policies and practices in one place. The outgoing part of the Gateway is the Initiating Gateway Actor. The incoming part of the Gateway is called the Responding Gateway Actor. Components within the community interact with the Initiating Gateway Actor to carry out transactions with other communities. Other communities interact with the Responding Gateway Actor as the point of contact for all requests to the community.

Internally a community may, or may not, be an XDS Affinity Domain. Implementation of a Gateway depends on its need to support interaction with XDS Affinity Domain actors. The case where the Initiating and Responding Gateway's use non-IHE specified methods for communicating within the community is presented in Figure 6.4-1. Please note that this figure shows only one direction of travel although in most cases a community would participate as both an Initiating and Responding Community and would, therefore, implement both the Initiating and Responding Gateway Actors.

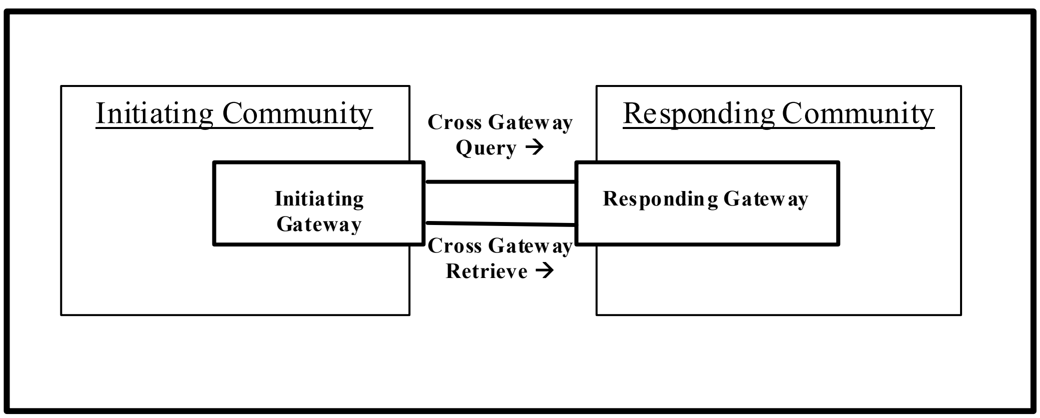

Figure 6.4-1: Base XCA

If both sides support XDS Affinity Domains the IHE specified interactions look like:

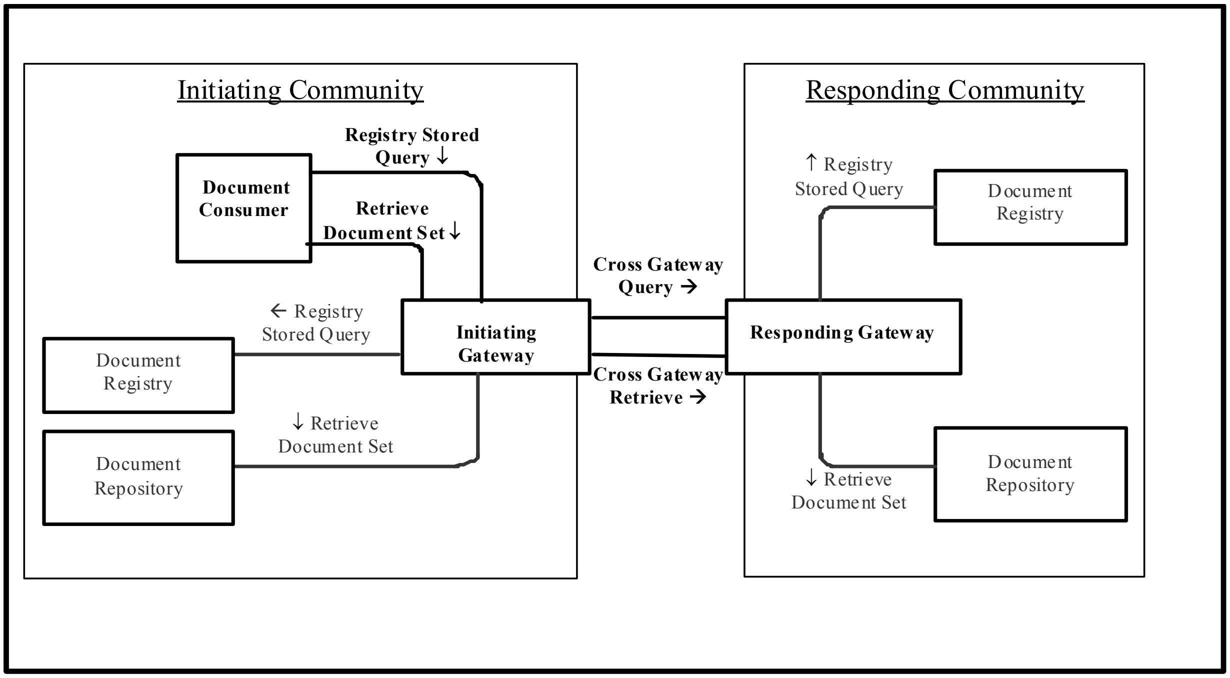

Figure 6.4-2: XCA supporting XDS Affinity Domains

The XCA profile contains a gap in the communication of patient identities. It requires the initiator of a query to determine the patient identifier of the correct patient under the assigning authority of the receiving community. But XCA does not contain detailed, profiled, specification for doing so. For more details see Section *8.6 Identifying patients of interest*.

6.5 Security and Privacy Profiles

The "*HIE Security and Privacy through IHE"* White Paper published by IHE presents an overview of IHE Security and Privacy profiles. You may access this White Paper at:

http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Whitepaper_Security_and_Privacy_of_HIE_2008-08-22.pdf

Please review section 3.4 for a high level understanding of the relevant profiles.

7 Creating a meta-community using existing IHE profiles

A meta-community can be created using existing IHE profiles by using a PIX Manager, called topmost PIX Manager, which supports patient identifier mapping across the patient identity domains of the communities in the meta-community. This is a meta-community because all communities have agreed to push to the topmost PIX Manager all patient demographics for patients which consent to sharing across the meta-community. The following requirements apply.

- All communities participating in the meta-community have agreements in place which address privacy and security policies which support sharing of patient data with other communities in the meta-community as well as sharing patient demographics with the topmost PIX Manager. These policies include, but are not limited to: security, privacy, access control, auditing, The transactions and behavior described here are compatible with a reasonable set of policies in support of privacy and security
- The topmost PIX Manager will receive Patient Identity Feeds for all patients who have agreed to share data within the meta-community. The feeds are delivered in a reliable way and none are lost. The PIX Manager will use a demographics matching algorithm to match patients across the communities. It is assumed that the population is small enough or discrete enough to make this automatic matching effective enough to satisfy patient safety and privacy concerns. For some populations, automatic matching is unlikely to work and some form of manual intervention may be required. In cases where the likelihood of an inappropriate match is high the use of a topmost PIX Manager may not be appropriate unless it includes a mechanism to support manual matching when appropriate.
- A mapping between the XDS Affinity Domain patient identification domains (assigning authority) and Responding Gateways must be defined. This environment uses the patient identification domain to determine the set of Responding Gateways to direct a patient query to. Preferably this mapping would be one-to-one, but it does not need to be.
- Agreements regarding the coding systems in the metadata and documents must be defined. Preferably a common coding system is agreed to and the gateway actors do mapping when necessary.

7.1 Detailed Description

The following figure shows the transactions and behaviors of the actors involved. All transactions shown are existing IHE transactions without modification. For the purposes of the PIX Query transaction, the Cross-Community Bridge acts as a PIX Consumer when it queries the topmost PIX Server.

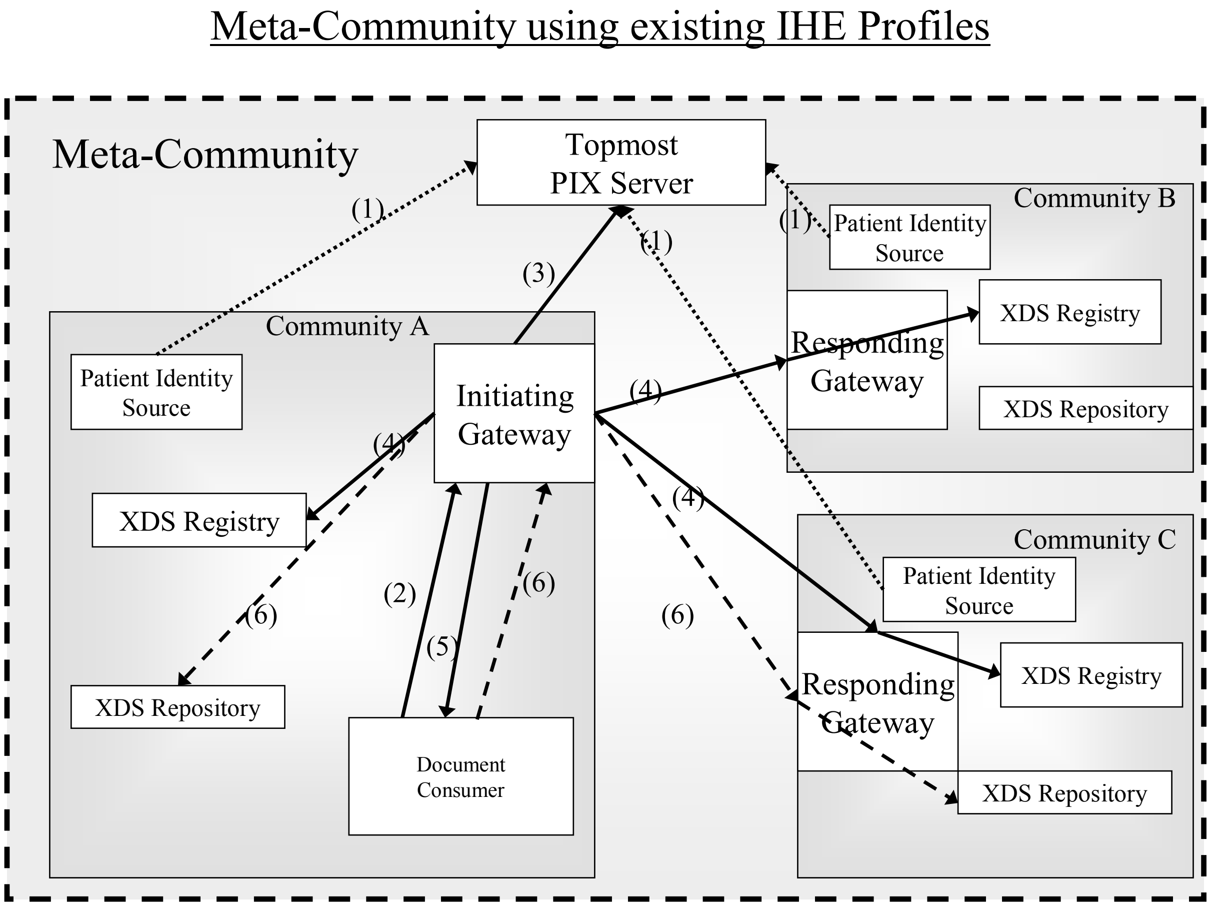

Figure 7.1-1: Meta-Community using existing IHE Profiles

1. The Patient Identity Source for each XDS Affinity Domain in the meta-community must send a Patient Identify Feed for every patient whose records will be shared across the meta-community to the topmost PIX server. This will allow the topmost PIX Server to cross reference the patient identifiers among the XDS Affinity Domains. The publishing of the patient id may be a consequence of a patient consent process reflecting the patient's consent to "opt in" to participation within the meta-community. There are other approaches to handling of patient consent which would not use this mechanism.
2. A Document Consumer issues an XDS Registry Stored Query transaction to the Initiating Gateway for Community A. The Initiating Gateway supplies the extra work of cross referencing the patient identifiers, querying multiple communities and combining the query results – see steps 3, 4, and 5.
3. The Initiating Gateway issues a PIX Query using the identifier specified in the XDS Registry Stored Query. The PIX Server returns a list of matching patient identifiers and their patient identification domains.

4. The Initiating Gateway uses the list returned by the PIX Server and the mapping from patient identification domain to community to initiate a Cross Gateway Query to each Responding Gateway and its local XDS registry. If the patient being queried does not have an identifier in a domain this indicates that there are no sharable records for that patient in the corresponding community and so no query will be issued to that Responding Gateway. Each Responding Gateway queries their local XDS Registry.
5. The results of the queries are combined into one result and returned to the Document Consumer
6. The Document Consumer retrieves documents of interest by sending a retrieve transaction to the Initiating Gateway which contacts the appropriate other community's Responding Gateway, which contacts the XDS Document Repository in which the document or documents are located.

8 Common Themes in Cross-Community

As the cross-community environment unfolds there are several common themes which come up as people talk about how to exchange patient information among HIE's. Some of the common themes are:

- Finding a community of interest
- Satisfying contractual and legal requirements for sharing patient data
- Policy Bridging
- Audit Log Access
- Identifying capabilities of a community of interest
- Identifying patients of interest
- Querying for data
- Retrieving data
- Sharing coded values and terminology

Query and Retrieve of data have been addressed within the IHE profile XCA, see Section *6.4 Cross-Community Access (XCA)*. For each of the others we will discuss some possible approaches, and places where a future IHE profile would be appropriate.

8.1 Finding a community of interest

In the environments of today finding a community of interest is not challenging since there are few that are enabled in a standards based way for cross-community communication. In the situations where such communication is enabled, the identification of such communities is done via manual exchange of configuration and use of configuration files. As the environment develops we expect that this approach will not be sufficiently effective and a more automated mechanism will be desired. Discussions on this topic generally suggest there will be a registry of communities, called community registry, which can be queried and/or will support the ability to subscribe to changes, so additions, deletions and updates will be pushed to communities. It is envisioned that there will be more than one community registry, perhaps using replication, distributed throughout a nation – and eventually the world. The community registries will replicate and share data. Communities will be able to query community registries, perhaps several, to collect information about other communities. Also a mechansm to subscribe to additions, deletions and updates of communities is needed.

The model described above could be enabled with an IHE profile. Such an IHE profile would enable standards based query of community attributes, as well as subscribe and publish mechansims for changes. See Section *9.2.2 Example use of Cross Community Discovery* for a detailed description of how this might work. The transactions define by this IHE profile would not contain any patient specific information, but only general, mostly public, information about a community. As such, its security and privacy characteristics will be different than those typically applied in an IHE profile. See the next section for further discussion.

8.2 Satisfying contractual and legal requirements for sharing patient data

Prior to sharing of patient healthcare data a community must ensure that all legal requirements, expecially in regards to security and privacy, are satisfied. In the previous section there was no sharing of patient healthcare data so this particular issue does not come up.

Two approaches to satisfying this mechansim have been observed in practice:

- Pair wise legal agreements and negotiations. In this approach, prior to any sharing of data a community negotiates with each individual community to form legal agreements. Those legal agreements form the basis to allow sharing between the communities.
- Collaborative groups. In this approach a group of communities form a collaborative which operates under a particular set of agreements. Often this group of communities is formed and controlled by a government, although this not always the case. In a sense this is just an extension of the pair wise situation, in that the agreements are negotiated prior to patient data sharing. The difference is that the agreements span a group of communities, rather than just two. The key is that this group is fairly fixed in number, with the process of growing or expanding this group being manual.

Fundamentally any sharing of patient data must operate under a clear set of agreements regarding its use and distribution. These agreements can be set up privately, or by law. In either case a community registry could be a source of assurance, in that one potential function of a community registry could be to only index communities which have agreed to a set of policies and agreements regarding sharing of patient healthcare data. Suppose there were a process, defined by contract or law, which was applied by the community registry implementor prior to adding any community to its community registry, and this process included management of the references to the community, to ensure that they remain viable and reputable. This would provide a bootstrapping mechanism upon which communities that interacted with the community registry could reley to enable future communication with other communities. This would remove the manual process currently needed by each community to assure that its sharing across communities is safe. The manuel process would most likely still exist, but would be more centralized at the community registry level and, for many nations, would be a national process which community registries would rely on. The IHE profile discussed in the previous section should consider this model. Most likely the details would be sufficiently different across countries that an IHE profile would not be appropriate. See Section *9.2.2 Example use of Cross Community Discovery* for a detailed description of how this might work.

8.3 Policy Bridging

There are many policies that are expected to be unique within a community. When these communities are connected these policy spaces need to be bridged. The term policy bridging relates to the fact that unified policies is unlikely and thus there needs to be some mechanism that correlates one policy with another. The main place where policy bridging is exposed is in vocabulary.

When a standardized vocabulary can be used it is the best way to assure that the concept will be understood when it is ultimately consumed. Often this isn't the case so current practice is to always qualify any vocabulary term with the value-set that defined its meaning. In this way the ultimate consumption of that vocabulary term can refer back to the original definition at the time that the vocabulary term was used. This ultimate consumption can be eased with an automated policy bridging.

There are many XDS Metadata values that are defined by the XDS Affinity Domain. When two XDS Affinity Domains are bridged, people can figure out what the equivalent vocabularies are. Thus it is possible, though not required by IHE, to automatically convert XDS metadata values. Care must be taken when automatically converting a vocabulary as there is always the possibility that a new version of a document or an amendment may be registered. These updates need to be understandable by the original Document Source.

One example of this is confidentialityCodes. These codes are often specific to an XDS Affinity Domain, and by definition a code that a Document Consumer doesn't understand can not be used to gain access to a document. Thus the ultimate consumption of a document is gated by understanding the confidentialityCodes. In the case of confidentialityCodes any policy bridging should only add equivalent codes and not remove the original codes. The behavior of confidentialityCode in XDS is that they are all in an OR relationship. By maintaining the original code the original intention in the original XDS Affinity Domain is maintained.

Another example is the use of XUA in an XCA environment. XUA provides a way to add an assertion about the user that has caused the query to happen. This assertion is profiled enough to be used by the first level service. In an XCA environment it is possible that this assertion will need to be further propagated to other Responding Gateways. This can be done through simple forwarding, where the secondary services will not be able to validate the assertion. Or may be done through the formal SAML proxy mechanism, where the Initiating Gateway requests a reissuance of a secondary assertion. This proxy mechanism does require that the original assertion allowed proxy, and that the Initiating Gateway has access to the Identity Provider in order to request the secondary assertion. At this time, it is not clear which of these or other mechanisms will be best in an XCA environment.

8.4 Audit Log Access

One of the Policies that need to be worked out before two communities can be bridged is how will security audit logs be handled. There are many good options here and the are supported to different degrees by the IHE Audit Trail and Node Authentication (ATNA) profile.

The ATNA profile includes a set of security relevant events and XML schema defining what to capture in an audit record when these security relevant events happen. All systems are expected to support the recording of all of the security relevant events that might happen on the system. Once an event has happened in the HIE, it will be described in detail in an XML message and communicated to an Audit Record Repository.

The Audit Record Repository is expected to be able to do Filtering, Reporting, Alerting, Alarming, as well as forwarding of events to other Community Audit Record Repositories. The more centralized this audit log analysis can be, the easier it is to provide accountability across the whole Community. The Audit Record Repositories can be centralized or distributed. The approach used by ATNA allows for one or more Audit Record Repositories in the Community. Depending on the policies, each system may have their own Audit Record Repository, there may be a hierarchy, or there may be one for the whole Community.

The following figure shows an Electronic Medical Record producing audit logs within the local Clinic. This local clinic's audit record repository is configured to forward a selected subset of messages that were triggered by activity with the HIE. Further the HIE Audit Record Repository would filter its total audit events and forward a subset to the other Communities. This HIE filtering may be based on the related events, meaning that a query seen at the HIE that did not get forwarded to a community would not need to be forwarded. But a query event that came from the EMR that was forwarded to other communities would need to have the audit events from the EMR forwarded to those other communities. There may be other heuristics that are used to filter and forward.

This may be an automated process, manual reports filed on a regular basis, or only brought together when an incident invokes a log aggregation policy. During an HIE incident investigation there may be need to go back to the Clinic to do a detailed investigation. The HIE policy needs to cover this.

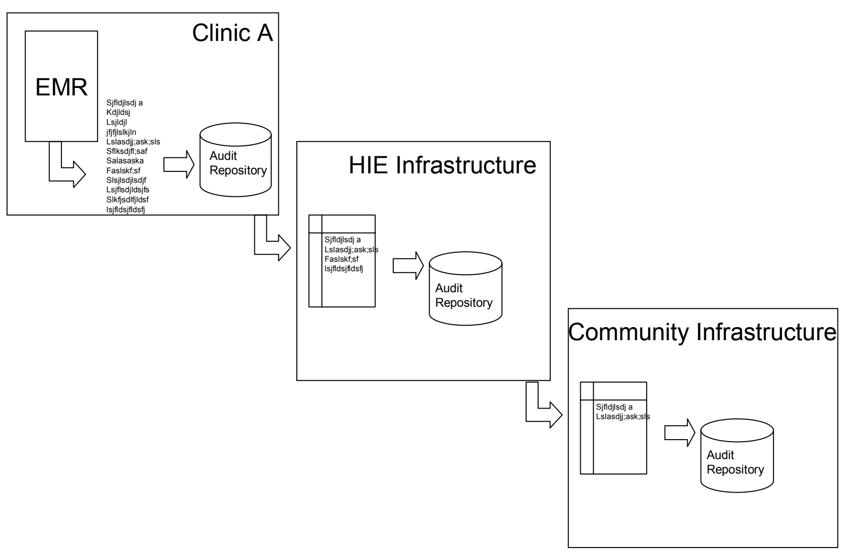

Figure 8.4-1: Audit Flow Down

8.5 Identifying capabilities of a community of interest

Once a community of interest has been located and appropriate agreements regarding sharing of patient healthcare data are in place, a community will need to know the services that are supported by the community of interest. Because we expect this environment to be somewhat dynamic in terms of what interfaces a community is able to support, it seems useful to allow some way of discovering those services or capabilities. This discovery would be over time, for once two communities start sharing data the capabilities on both sides are likely to change. That change would be slow, happening generally over months or years.

One possibility is that the community registry which enabled the sharing could also enable the sharing of capabilities, as well of subscription for updates of capabilities. Alternatively, each community could provide a service for this sharing and subscription. In either case an IHE profile would be useful in enabling the process. At this time it is expected that capabilities will be expressed in the form of WSDL files. See Section *9.2 Cross-Community Discovery* for details.

Another type of capability that communities may choose to share is the coding system that is used for XDS metadata, document formats and coding systems within documents.

8.6 Identifying patients of interest

Prior to sharing patient healthcare data a community needs to be sure that both are talking about the same actual physical person, i.e. that my Jim Smith is the same as your Jim Smith. We expect there will be many different approaches to accomplishing this challenge. For example:

- Use of a national patient identifier
- Cross referencing via a topmost PIX
- Demographics Query – on the fly discovery
- U.S. NHIN style
- Austria work

Each of these will be explored in the following sections

8.6.1 National Patient Identifier

Some countries will make use of a national patient identifier which can be releyed upon to uniquely identify a patient within a country. Although separate identifiers might be used within a community, all cross-community communication would identify the patient via that patient's national patient identifier. This model assumes that the assignment of the identifier is done without significant error or duplication. It risks the exposure of the identifier and suggests that special security and control of the identifier is needed. A IHE profile in support of this approach is not needed.

8.6.2 Cross referencing via a topmost PIX

This approach has been described in detail in section *7 Creating a meta-community using existing IHE profiles*. Existing IHE profiles enable this approach and no further work is required.

8.6.3 PDQ – on the fly discovery

In this approach a community could send a query containing demographics to a community of interest requesting matching patients. A list of patients with matching demographics are returned, along with the other community's patient identifiers. The querying community can then pick the closest match based in its local matching criteria. This approach has complications since demographics often change over time and if the querying community chooses to save the received demographics there is no way to know about changes. Concerns about patient identity merge, link, unmerge, unlink all need to be considered in this environment. Although the PDQ profile satisfies the basic needs of this approach, because of the complicating environment of a cross-community exchange we suggest that an IHE profile is needed to profile a demographics query across communities.

8.7 Sharing coded values and terminology

In order to effectively share data a common understanding of coding systems and meaning is required. This includes the metadata used in the query and response via XCA as well as the content of the document. Translation from one coding system to another can be challenging, as exact matches for codes is sometimes not available. Translations from finer grain to coarser grain terminology can be accomplished, but from coarse grain to fine grain is not easy to do automatically. Current systems have generally approached this by requireing all coding to be in a common format. Internally a community may use a different set of codes, but those must be translated prior to cross-community communication.

8.7.1 XDS metadata

Effective communication is possible without any agreement on coding of XDS metadata because some of metadata elements use a common coding system or are not coded. We believe the following XDS metadata elements are usable across communities without translation or specialized understanding.

- author
- availabilityStatus
- comments
- creationTime
- formatCode
- languageCode
- legalAuthenticator
- serviceStartTime

- serviceStopTime
- sourcePatientId
- sourcePatientInfo
- title

The following XDS metadata elements are coded and in order to be usable across communities would need to either use a common coding system or need to be translated. Note that patientId is covered in section *8.6 Identifying patients of interest*.

- classCode
- confidentialityCode -
- eventCodeList
- healthcareFacilityTypeCode
- patientId
- practiceSettingCode
- typeCode

8.7.2 Document Content

Document format is defined by the XDS metadata element formatCode, which is expected to hold a value that identifies an existing IHE Content Profile which defines the document format. In an ideal world every community would share only documents whose format is defined in an IHE Content Profile and all communities would understand all documents abiding by any IHE Content Profile. It is unclear at this point to what degree the ideal situation will hold true. Communities may define proprietary formats or may only use a small number of document formats. If the receiver of the document is unable to effectively view and process the document it becomes significantly less useful. If acceptance of IHE Content Profile formats is not good enough there will be more work involved in understanding content of documents.

Beyond the format of a document, some documents, like CDA, have coded values inside which also may need translation across coding systems.

9 Proposed Future IHE profiles

Given the discussion in Section 8, this section outlines a set of profiles to be considered for the next few years. These profiles will support both meta-community creation and lateral communication.

1. Cross-Community Patient Identification – this profile will define actors which allow for discovery of matching patients in communities of interest. Initially this profile would be based on patient demographic query, but as other approaches come to light it could be expanded beyond a demographics query.
2. Cross-Community Discovery – this profile defines mechanisms for automatically discovering communities of interest, based on community level attributes, and discovery of services to automate the protocol and policy decisions required prior to exchanging data with a newly discovered community.

We believe that Cross-Community Patient Identification is necessary to enable current industry activities and should be addressed in the coming year.

9.1 Cross-Community Patient Identification

This profile would extend the existing XCA profile to support patient identity matching. Profiled support of a patient demographics query between XCA Intiating and Responding Gateway would be a primary purpose of the profile. Alternate methods of exchanging patient identities should be investigated and integrated into the profile if appropriate. In some cases, grouping with existing actors may be sufficient. For instance, to support patient demographics query it may be necessary only to declare a grouping of a Initiating Gateway with a Patient Demographics Consumer and a Responding Gateway with a Patient Demographics Supplier. If the grouping behavior is sufficient one reason to document this grouping behavior is because it perhaps should be required support of Initiating and Respondiong Gateways.

9.2 Cross-Community Discovery

9.2.1 Overview

This section describes a workflow to support automated lateral communication of patient health information. Because lateral transactions do not have an overriding organizational mechanism (as in meta-communities) there is no common place to build up patient data existence locations. The transactions described in this section support using community attributes to identify communities holding patient records. The goal is that a community (or meta-community) is able find patient data locations for patients of interest to the community. There will be no national or global index of patient data locations, but for any patient there may be several communities which are identifying and possibly saving the location of records for that patient.

9.2.2 Example use of Cross Community Discovery

The following diagram presents a high level overview of the process of identifying communities holding records for a particular patient.

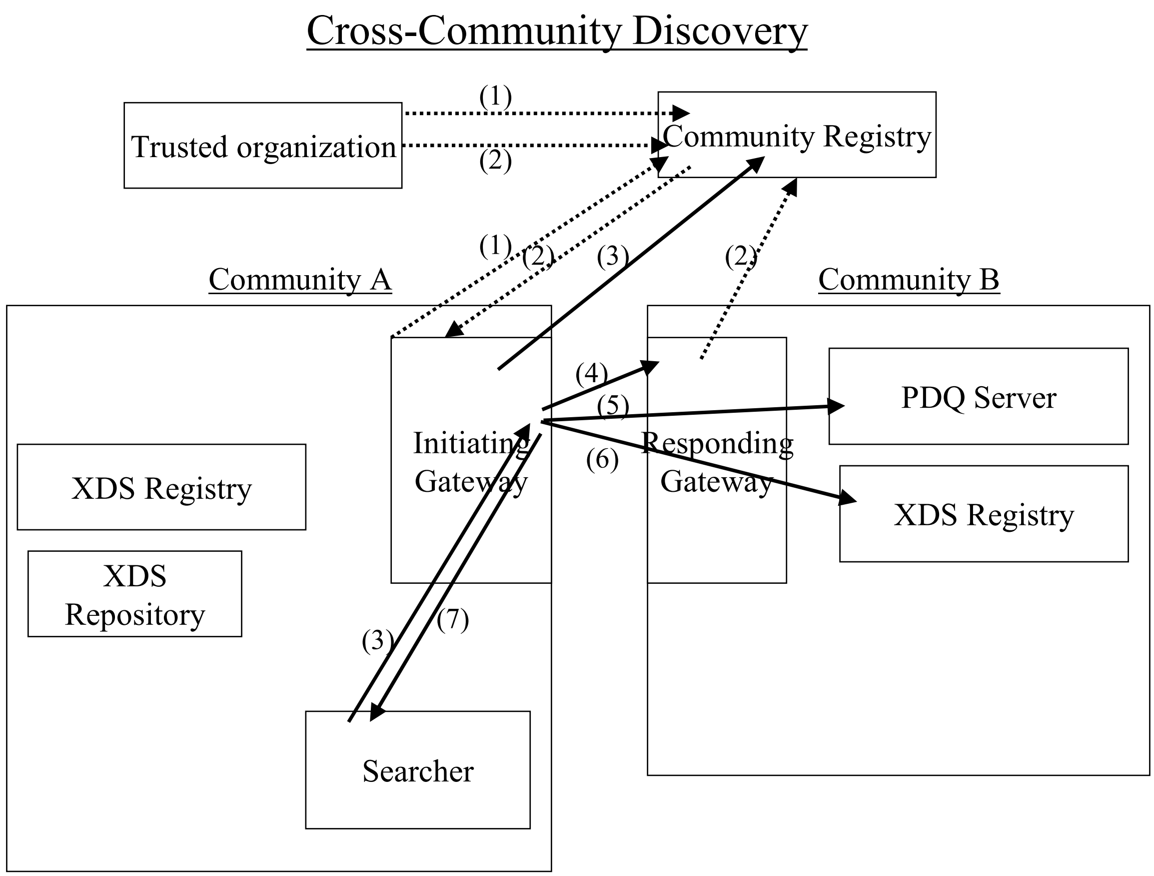

Figure 9.2-1: Cross-Community Discovery

1. Prior to any cross-community communication, each community is verified by a trusted organization. That verification includes agreements about policies for data sharing and certification that the community is valid and can support the environment effectively. Once verified the trusted organization feeds the community attributes to the community registry. In the diagram (1) sends attributes about Community A and by doing so has validated that Community A has been verified for sharing under the policies defined by the trusted organization feeding the Community Registry. Community A then queries the community registry to pull the current list of communities, and requests that all future changes are pushed. Community A establishes a local cache of other communities indexed by the community registry.
2. At some time later Community B is verified and its attributes are fed to the Community Registry. Community B then queries the community registry to pull the list of

communities for its use. Because Community A requested updates, the information about Community B is sent to Community A which updates its local cache.

3. A new actor is defined, Searcher, which initiates the next phase. This actor and interactions could be profiled, or could be considered internal workings of an Initiating Gateway which require no profiling. This actor could potentially be linked with other client type actors like Document Consumer. The Searcher initiates a Find transaction providing the patient demographics and a set of attributes to be used to look for communities which might hold information about the patient. Community A can either look in its local cache for appropriate communities to contact, or could query the Community Registry. In either case, Community B is a potential match for having data about the patient in question.
4. The Initiating Gateway queries the capabilities of Community B to determine if communication is feasible. A community may not support the particular protocol required by the gateway, or it may have security or privacy restrictions that the gateway does not qualify for. If the gateway cannot communicate with a community for any of these or other reasons it notes this, via logging or error message to the Searcher. If any other communities matched the attributes the Initiating Gateway queries the capabilities of other communities.
5. In this example the community is an XDS Affinity Domain and as part of the capabilities states that the Responding Gateway supports the Patient Identification service. The Bridge queries this service to get the correct patient identifier using demographics provided by the Searcher application. If no patient is found the Initiating Gateway continues with the next community assuming that there is no available data in this community.
6. Assuming a patient is identified in the previous step, the Initiating Gateway uses the patient identifier returned and the capabilities previously retrieved to query the Responding Gateway for document entries for the patient.
7. The Initiating Gateway returns information about what was found, or was not found, to the originator.

9.3 Summary

The long term vision for Cross-Community Information Exchange includes IHE profiles to enable hierarchical and lateral communication strategies. The Initiating and Responding Gateway actors facilitate all transactions into and out of the community. The Cross-Community Access profile supports retrieving patient health records from external communities. The correlation of patients is enabled by via the Cross-Community Patient Identification Profile. And the discovery and automated configuration of inter-community lateral communications is enabled via the Cross-Community Discovery profile.

A.1 Appendix - Use Cases

This section lists the use cases discussed in preparation for this paper. There is a long list of use cases that could be viewed as illustrating the problems to be solved by a Cross Community Information Exchange profile. This list is not necessarily complete and not all use cases are addressed by the technical objects defined in the paper.

Use Case: Multiple primary residences

This use case describes the situation where a patient maintains more than one principal residence. Generally the principal residences are not geographically close so the medical data generated while in each residence would be created by separate institutions.

A common example of this use case is what is described in the United States as the Snow Bird. This is a person who maintains two residences, one in the northern part of the USA for use during the hot summer months and one in the south for the colder winter. If a patient lives in Florida in the winter and in New York in the summer, this patient will likely have medical records in both places which need to be shared. If the patient is managing a long term medical condition, like diabetes, it will be important as she moves from New York to Florida and back that the background and related testing associated with management of the medical condition is readily accessible to the local physician.

A variation of this case involves a patient who lives on the border between two communities or works and lives at some distance. A patient who lives in Greenwich, Connecticut, which is on the border between the U.S. states of Connecticut and New York, may access health facilities both in New York and Connecticut. If that patient works in Manhattan, he would also access health organizations in New York City. All of these disparate areas, although closer geographically than New York and Florida would probably hold patient information in separate domains thereby requiring the same kind of cross domain sharing as the snowbird case described above.

Use Case: Between two XDS Affinity Domains

A doctor's office is on the borderline of two XDS Affinity Domains and frequently deals with both. In this case the doctor may choose to be a member of both XDS Affinity Domains, submitting and retrieving documents from both - separately. This can be accomplished with existing profiles.

A patient lives on the border between different states or regions. This patient will likely access medical services in multiple regions.

A different example that fall in a similar use case is that of a Provider member of several XDS Affinity Domains that wish to locate the specific XDS Affinity Domain to which a Consumer has his PHR Service Provider located.

Use Case: Patient Move

A healthcare facility may need to do a one time transfer of information from one XDS Affinity Domain to another, perhaps because of a patient move. This might be a good application for the XDM (Cross-Enterprise Document Media Interchange) or XDR (Cross-Enterprise Document Reliable Interchange) profiles.

Use Case: Vacationer

Vacationer: A patient is traveling and goes to the hospital. The hospital needs to access records from the patient's home region.

Once the treatment is complete the patient will want to have the records available to his home community. The transfer of the records to the home community might be a good application of XDM or XDR.

Use Case: Mergers, acquisitions, divestitures.

Merger

A large healthcare corporation, Bigco, acquires a small local hospital, Smallco. Before the records integration transition is complete, the medical records from Smallco reside locally at Smallco. When a Smallco client is admitted to one of the other facilities of Bigco, their records appear to be stored locally at that other facility. But in fact, the records are being transferred as necessary from Smallco. There is also a transitional activity going on to transfer the Smallco records into the Bigco EHR. This does not happen instantaneously. In fact, it takes several weeks for the complete transition. During this period, a facility must track which patient records are kept in which EHR facility.

The transition can be managed by creation of one additional affinity domain that incorporates only Bigco and Smallco. The use the XDS affinity domain mechanisms to track locations of information during the transition to a single EHR. During this period there is a kind of federation between this internally motivated transitional affinity domain and the other affinity domains.

Divestitures

A healthcare facility splits into two new organizations: Newco1 and Newco2. At the start of the split, all records are being kept in a single EHR facility. For a period of time while organizational changes are taking place, the two organizations share that single EHR facility. But then, they set up their own internal EHR facilities. There is a gradual transition of medical records from the old EHR facility into the new separate facilities. During this transition, the three EHR facilities (Old EHR, Newco1, and Newco2) act as a federated system to track the current location of patient records. A transitional affinity domain can be used to manage this transition. During this period, both Newco1 and Newco2 are members of it and of all the other affinity domains that Oldco had been a member in.

This does raise the related issue of managing the transition when a healthcare facility decides to withdraw from an affinity domain.

Use Case: Transitory Alliances

A medical facility and a research facility form a short term alliance for a research project. The two organizations would like to share information during this short period of time only.

Use Case: Surveillance, CDC

Multiple Local, State, Federal agencies are interested in the collection of largely overlapping information produced in the course of delivering care (e.g. laboratory or case information). Organizing this information in pseudonimized "digital objects" or documents shared between the many sources of clinical information through an XDS Affinity Domain that serves the various Public Health agencies. .

Use Case: Specialty treatment

A large medical institution (located in Boston for example) frequently has patients from outside the region come for specialized treatment. The medical institution needs access to records created prior to the specialized treatment and the home region needs access to records related to the specialized treatment.

Use Case: PHR Services

Personal Health Records systems will want to import data from and export data to a variety of EHRs where the consumer receives care. As those EHRs may belong to multiple XDS Affinity Domains, other than the XDS Affinity Domain to which the PHR Service Provider may belong.

Use Case: Remote Consulting.

One hospital would like to consult a specialist in a remote hospital. In order for this consultation to take place a number of patient records should be extracted and made available for one or more interaction between one or more clinicians across the two institutions. After some time (e.g. a couple weeks) this shared information is no longer needed and may be deprecated. XDS is in clinical use today between two Italian hospitals, one Spanish hospital, and a Dutch hospital. These hospitals may belong to different XDS Affinity Domains and access to this shared information may require a targeted federation.

A.2 Appendix - Strategies for solving the selected Use Case

In this section we will narrow the discussion to the use case involving multiple primary residences. In this use case, medical data is being collected in two or more geographically separate areas and sharing of that data is required in order to properly treat a patient's on-going medical condition. The strategies outlined below are high level mechanisms for sharing patient data among communities.

A.2.1 Distributed

This is the typical data sharing mechanism currently used. Data is created within many communities and sharing of medical data is done using an ad-hoc method defined at the time that the sharing is needed. Typical methods for sharing are:

- The patient collects physical copies of appropriate documents and carries them to the receiving organization (known as the "sneakernet" method of sharing by some).
- U.S. Mail or FAX.
- Collection of data on some electronic media (USB device or CD) which is then physically transported to the receiving organization – hoping that the format can be interpreted by the receiving organization.

- Creation of a special purpose sharing arrangement that connects them electronically through e-mail or the Internet using some locally defined protocol.

These methods, and probably more, all must be negotiated again and again and none are available without a significant level of cooperation from all parties involved.

A.2.2 Distributed, community indexed

In order to encourage the finding of information, an inter-community indexing service is designed and provided which collects metadata about each community (such as geographical region and/or a list of facilities within the community) and defines the mechanism for the sharing of medical information between communities. This allows, in our use case, the community in New York to search for the appropriate community in Florida and query and retrieve data from the Florida community. This strategy assumes that the retrieving organization has some understanding of the community from which it will be pulling data, or the retrieving organization has sufficient resources to query all communities in the index.

This strategy supports peer-to-peer communication in situations where there is no common cross referencing based on patient. It is believed that a national patient cross referencing service will not be acceptable for some nationalities, thus leading to this form of generalized, public search.

A.2.3 Distributed, patient indexed

To encourage finding of information for a particular patient, an indexing service is provided which takes information about a patient (demographics most likely, another variation could use patient identifiers in some form) and return a list of communities which are expected to have information about that patient. The service or some other mechanism would also define the protocol to be used to query and retrieve data from a community once it has been located. The mechanism for identifying a patient is a common problem in all the strategies. It is assumed that demographics will be used to identify a patient unless the country uses a national patient id.

See Section *7 Creating a meta-community using existing IHE profiles* for one implementation of this approach.

A.2.4 Collected metadata, distributed data

To promote timely searches for relevant information, all the metadata about the patient records is collected within one community. This community is referred to as the home or preferred community. A mechanism for identifying the preferred community for a particular patient is needed, although the patient can generally provide this him/herself except in emergency situations.

By collecting metadata about a patient in a single place, a single query can identify all the documents that are available for that patient. This metadata is general, non-clinical (except in some situations and countries) information about the patient record (e.g. XDS Registry metadata). The patient data is still stored at the creating sites and retrieved from there if desired.

A standardized method for communicating queries and retrievals between communities is needed as well as a mechanism for getting the metadata to the home community; either by having other communities push it, or having the home community pull it. At issue is also how the metadata is kept up to date in case of modification.

This strategy assumes that the patient has agreed to have his data consolidated/collected.

A.2.5 Collected data

To encourage quick access in the place of most frequent use, all data (not just metadata) is collected in one community, the home or preferred community. This is equivalent to the above, except that the data is moved or copied to one community.

This strategy introduces the issue of who the real owner of the data is the creator or the home community storage system. If data is updated how is that managed?

A.2.6 Multi-home collected

Either of the collected strategies described above can be extended by using more than one home community and duplicating the data in both. This increases the questions regarding duplicate data, ownership, and processing of changes.

A.3 Appendix – National Domain projects that reference IHE

A.3.1 United States HITSP

The Healthcare Information Technology Standards Panel (HITSP) organization references IHE profiles for many of its national Interoperability Specifications. Figure A.3-1 presents the mapping from HITSP contstructs to IHE profiles.

Related Documents	Document Description	IHE
TP13	Manage Sharing of Documents Transaction Package	XDS.b, XCA
T15	Collect and Communicate Security Audit Trail Transaction	ATNA
T16	Consistent Time Transaction	CT
T17	Secured Communication Channel Transaction	ATNA
C19	Entity Identity Assertion Component	XUA
TP21	Query for Existing Data	QED
TP22	Patient ID Cross-Referencing Transaction Package	PIX
T23	Patient Demographics Query Transaction	PDQ
C26	Non-Repudiation Component	DSG
T29	Notification of Document Availability Transaction	NAV
T31	Document Reliable Interchange	XDR
T33	Transfer of Documents on Media	XDM
TP49	Sharing Radiology Results Transaction Package	XDS-I
C62	Unstructured Document Component	XDS-SD
T64	Identify Communication Recipients (service)	PWP
T66	Terminology Service	SVS
T67	Document Referral Request Transaction	DRR

Figure A.3-1 HITSP mapping to IHE