

# Integrating the Healthcare Enterprise



5

## IHE IT Infrastructure

# XDS Patient Identity Management White Paper

10

## Draft for Public Comment

15

Date: August 10, 2010  
Author: José Mussi (IHE Canada)  
Nathan Domeij (IHE Canada)  
Charles Parisot (ITI Tech Committee)  
Email: iti@ihe.net

25

## TABLE OF CONTENTS

1	Introduction.....	2
2	XDS and Patient Identity Management .....	3
	2.1 XDS Background.....	3
30	2.2 XDS Affinity Domain Patient Identification.....	3
	2.2.1 Impact to the XDS Document Repository .....	4
	2.2.2 Impact to the XDS Document Registry .....	4
	2.3 Changes to Patient ID in Published Documents.....	5
3	Cross-Enterprise Patient Identity Schemes.....	7
35	3.1 Decentralizing Matching (XAD-PID Authoritative).....	7
	3.2 Centralized Matching (Local ID Authoritative).....	10
4	Patient Identity Management Gaps in XDS Specifications .....	12
	4.1 XAD Level Link/Unlink Events.....	12
	4.1.1 Alternative 1 – Using ADT^A31 .....	13
40	4.1.2 Alternatives 2a and 2b – Using PAM .....	14
	4.1.3 Alternative 3 – Using ADT^A43 .....	14
	4.1.4 Use of XDS Metadata Update.....	16
	4.2 Additional Considerations .....	16
	4.2.1 Impact to Submission Sets and Folders .....	16
45	4.2.2 Impact to Document Repositories.....	17
	4.2.3 Impact of Local Merge Events.....	18
5	Recommendations.....	19
	Appendix A: PIX/PDQ Integration Models with XDS.....	19
	A.1 Canadian Interoperable EHR Blueprint.....	19
50	A.2 Other Integration Models.....	23

## 1 Introduction

55 This white paper discusses the technical considerations of using the IHE Cross-Enterprise Document Sharing (XDS) within a Health Information Exchange infrastructure or a jurisdictional shared EHR, in particular in respect to its integration with patient identity management services. The material here presented is based on the most recent XDS integration profiles and associated document content profiles such as XDS-I and XDS-MS specifications.

## 2 XDS and Patient Identity Management

### 2.1 XDS Background

60 The XDS suite of profiles has been created by IHE to support the exchange of a patient’s longitudinal health record across multiple enterprises. Its basic premise is that clinical information will be shared using a “clinical documents model”, where a collection of clinical information pertaining to the same patient is grouped into one or more “documents” and published to a shared infrastructure. These XDS services are composed of one or more document  
65 repositories and a common document registry containing metadata and pointers to all shared documents across these repositories.

Point of service systems that publish documents are known as “document sources” while those who retrieve them are called “document consumers”. The same system can be (and often will be) both a source and consumer of XDS documents. These document source and consumer actors  
70 interact with XDS Document Repository(ies) and a Document Registry Actor. All these “XDS actors” are contained within a single “XDS Affinity Domain”, which establishes a set of conventions about what type of clinical documents, security constraints and other applicable policies must be used by all organizations (i.e., enterprises) that have come together to exchange documents.

### 75 2.2 XDS Affinity Domain Patient Identification

Crucial to the ability to share documents reliably is the need to uniquely and correctly identify the person (i.e., patient or client) to whom the information belongs to. This is a non-trivial problem, as each clinical system that participates in the XDS Affinity Domain may (or more likely will) use different identification means for its patients. The challenge is to find a common,  
80 reliable identification scheme that can be used across the entire XDS Affinity Domain.

The XDS specifications do not attempt to resolve the identification problem, rather it assumes that the XDS Affinity Domain will have some common means to create a unique identifier for persons involved in the domain and allow document sources to find the appropriate patient ID prior to publishing documents to the XDS infrastructure. This identifier is called the XDS  
85 Affinity Domain Patient Identifier (XAD-PID)

Unless there is a shared patient identification scheme such as a regional or national patient identification in place among all Affinity Domain actors, IHE recommends the use of the PIX or PDQ profiles to manage the correlation of identifiers across the XDS Affinity Domain, where a PIX Manager actor or a PDQ Supplier actor provides each document source (and later document  
90 consumers as well) a match between the patient’s local identifier or identity (i.e., that which is known to the Point of Service (POS) application) and the common, Affinity Domain identifier.

With this information in hand, the document source system can reliably publish documents, create submission sets and organize information into XDS folders, all assigned to the same patient identifier.

95 Document consumers must also query the PIX manager to determine the corresponding XAD-PID before submitting any queries to the document registry. As mentioned previously, the XAD-PID is the only identifier that is recognized by the XDS services.

100 A number of variants in not using or combining the use of PIX and/or PDQ are discussed in Appendix A. It is assumed that the reader of this white paper is familiar with one or more of these approaches.

### 2.2.1 Impact to the XDS Document Repository

105 The XDS Document Repository, which first receives the information from the document sources, is mostly agnostic to the XAD-PID contained in each submission set (collection of documents included within an individual Provide and Register Transaction or Register Transaction used to publish documents to the XDS Registry/Repository infrastructure). The job of the XDS Document Repository is basically to receive the documents from the source, store them within its persistence layer, calculate the hash and size of each document, and assign an URI (only if it supports ITI-17) to each object and pass all this information to the document registry. It does not validate or change the patient identifier contained within the “provide and register” transactions.

### 110 2.2.2 Impact to the XDS Document Registry

115 On the other hand, the XAD-PID is a key document attribute for the XDS Document Registry. After all, its role is to organize and group documents that belong to the same person. The registry must ensure that a proper XAD-PID is provided with each submission set. To accomplish this, and before documents are allowed to be registered for any patient, the document registry receives patient identity feeds from the XDS Affinity Domain patient identity source containing information about valid identities for the XDS Affinity Domain. With this information in hand, the document registry can ensure that:

- XAD-PID contained in the submission set is valid and active for the XDS Affinity Domain
- 120 • documents within the submission set belong to the same patient (i.e., the document entry metadata contains the same XAD-PID)
- all documents (document entry metadata) within the same XDS folder (folder metadata) also belong to the same patient

125 The key point about this approach is that the XDS Affinity Domain patient identifier is the authoritative means for identifying patients and grouping documents. Although the local patient identifier can also be provided with each submission set, it is considered a mere attribute of the transaction.

## 2.3 Changes to Patient ID in Published Documents

130 Once a document has been published to the XDS services, there are only two ways where the assigned patient identifier may be changed:

- 135 1. By the original document source – in case where documents were published with the wrong patient identifier, the original document source (or any other authorized document source) can deprecate the original posting and re-submit the same document with the correct patient identifier. The document registry will change the status (using availability status attribute) of the incorrect entry (e.g., deprecated) and add the new information as provided by the document source. This process is integral to the document lifecycle management defined by the XDS specifications as well as the recent Metadata Update Supplement and the XDS infrastructure is not concerned about the reasons for the correction.
- 140 2. By the Patient Identity Source –the XDS specifications (Principles were defined by the now-retired XDS Patient Identity Merge Supplement) currently only allows for patient identity merges. In this scenario, two XDS Affinity Domain patient identifiers are discovered to belong to the same person. Through some process external to the XDS services, one of the two identifiers is chosen to be subsumed (i.e., no longer in use) by the other (i.e., the survivor). A merge notification is sent to the XDS registry by the Patient Identity Source Actor and from that moment on:
  - 145 • All documents that were published through an XDS Provide and Register transaction with the subsumed patient identifier are now joined with documents belonging to the surviving ID.
  - 150 • Any further submission sets communicated through an XDS Provide and Register transaction and referencing a subsumed ID will be rejected by the document registry with an “*XDSUnknownPatientId*” error.
  - 155 • All XDS Stored queries transactions referencing a subsumed patient identifier return no content.
  - All XDS Stored queries transactions referencing a surviving identifier return the entire recorded merge set of documents and return appropriate metadata.

160 The XDS specification does not specify changes to the internal state of the document registry. Instead it specifies required future behaviors on the part of the transactions listed above. Also, the patient identity merge notification is not propagated to the document repositories, which are not expected in the design of XDS to persist such XAD-PID.. One should note that there are no IHE transactions to the XDS Document Repository that make use of such identifiers (Document are always retrieved by Document UniqueID).

165 Note: The XDS Patient Identity Merge Supplement was based on XDS-a and will be deprecated. It is expected that the activities following the approval of this white paper will also include a review of this matter and new specification for merge events.

However, as will be described in the next sections, there is no current explanation in the XDS specification on how to handle link/unlink events triggered by the XDS Affinity Domain patient identity source.

170

### 3 Cross-Enterprise Patient Identity Schemes

175 This section will focus on two of the various schemes used to manage the XAD-PID that is used when publishing documents to a XDS infrastructure. It is important to understand the differences and similarities among these two approaches in order to understand the impact of patient identity management events to the XDS services. There are basically two classes or models of identity management: central matching or local matching. Both models assume that there is a single source for assigning and maintaining the XAD-PIDs. The difference between the two lies on where the matching between local and common identifiers occurs.

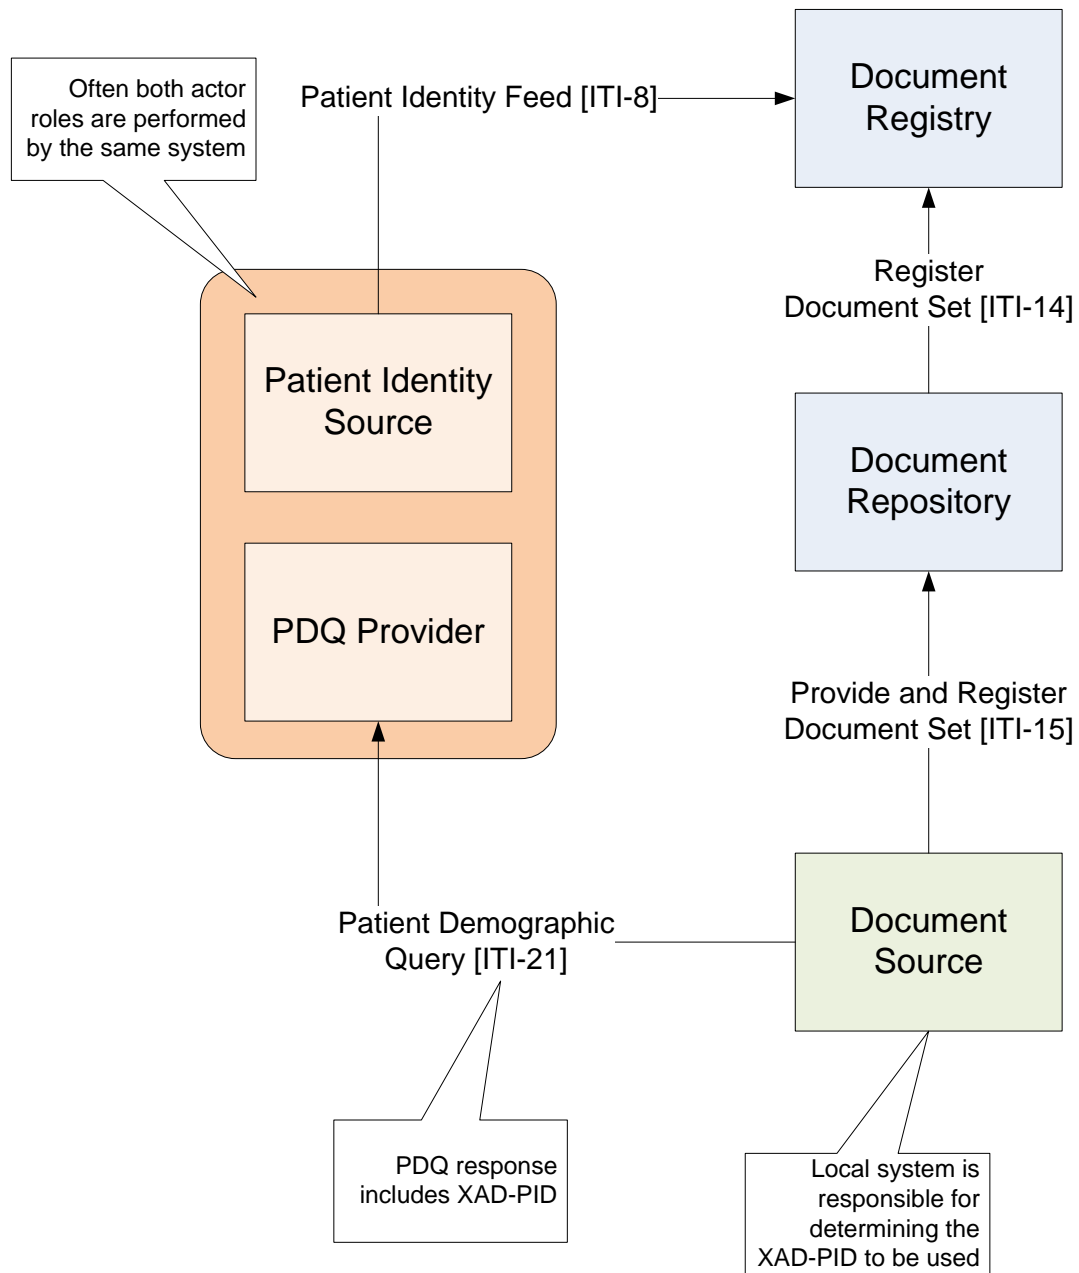
#### 3.1 Decentralizing Matching (XAD-PID Authoritative)

180 In this model, a central patient identity source (e.g., a PDQ Supplier) is used as the patient identity management authority (i.e., Patient Identity Source Actor) for the XDS Affinity Domain but this service is not responsible for determining which XAD-PID should be used in any particular transaction. It is the local system (XDS Document Source or Document Consumer) who takes on the task to properly match their local patient identifier or other patient identity  
185 information with those of the shared XDS Affinity Domain (see figure below for XDS use combined with PDQ).

Note that a similar sequence of events will occur with Document Consumer systems when querying the XDS services.

In this approach two key design elements have to be supported:

- 190
- 1) The mapping or matching between the locally assigned patient identity and the shared one is to be performed before documents are being published or queried and retrieved by the local system.
  - 2) Local systems will use different approaches to determine which XAD-PID should be applied, but it usually requires queries to the PDQ Supplier to find a match from a list of  
195 candidates or validate a given public (i.e., business) identifier, such as when using health or insurance card numbers as the common identifier for the XDS Affinity Domain.



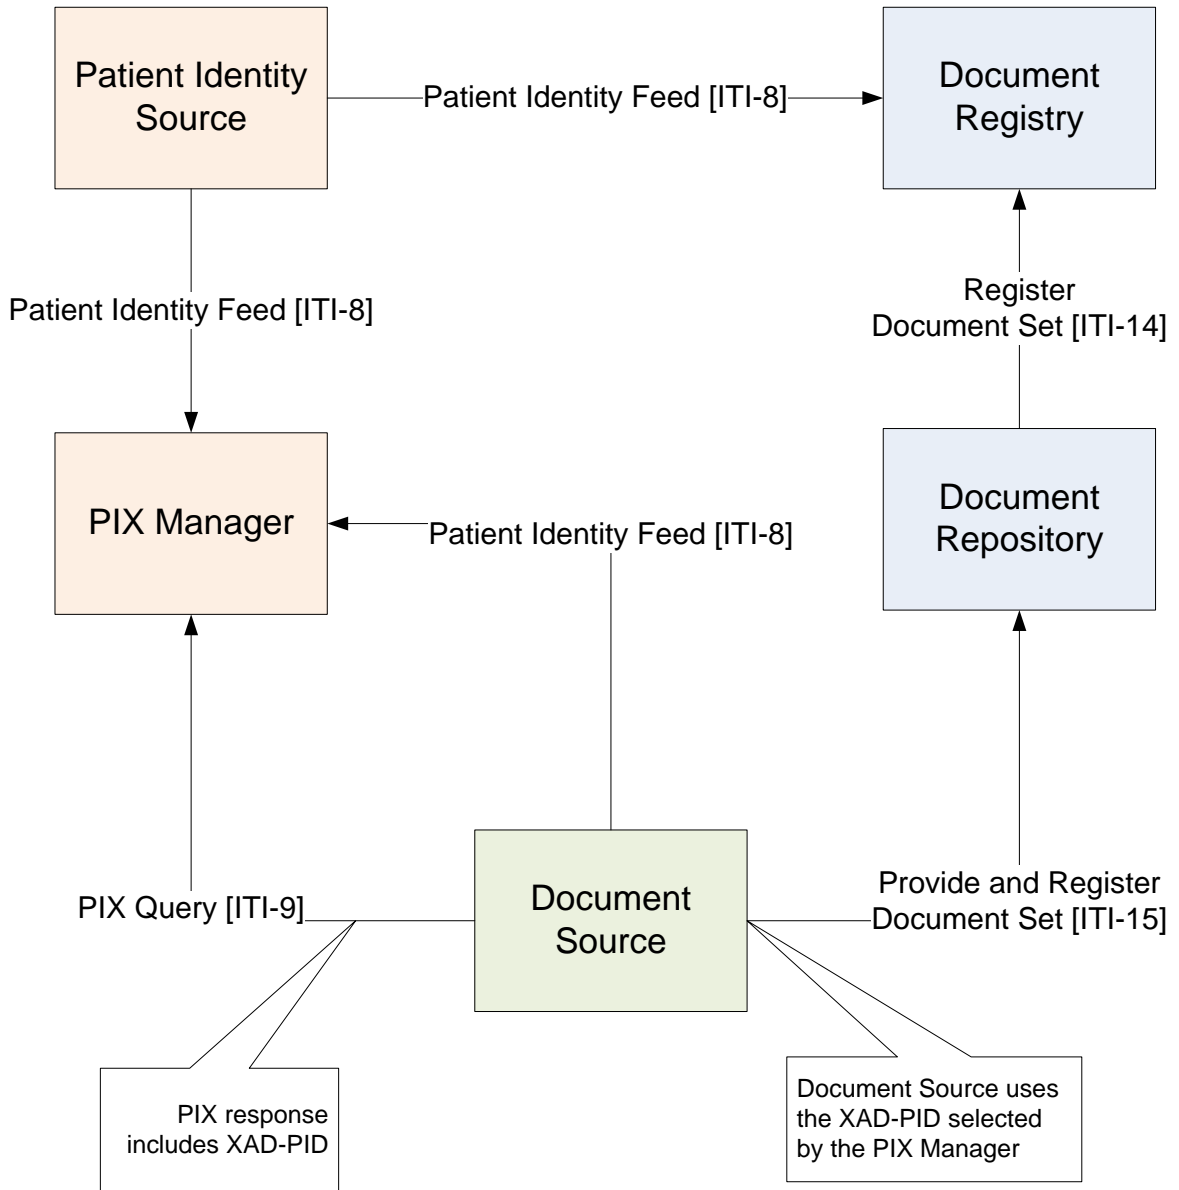
**Figure 3.1-1 Decentralized XAD-PID Matching using PDQ**

200 The XDS Affinity Domains must define the process (automated, and /or manual) for the administration of these common identifiers. Such identity management processes often piggyback on other business processes, such as those put in place to obtain patient consent when a scheme requiring opt-in is used. Other processes associated to health cards (associated or not with health insurance) may also used.

- 205 These examples highlight the key nature of the issue: in this model, the authoritative patient identifier for clinical data published to the XDS Affinity Domain (sometimes called shared EHR) is not the patient ID from the local system, but rather the XAD-PID. The local patient ID can be seen as just an attribute assigned by the local system and in fact may not be required or available in this case.
- 210 Consequently, this model requires that most changes to the XAD-PID associated to any particular document must be submitted by the original source system (or another authorized source). The only exception is that XAD-PID merge events will be managed and triggered by the central patient identity source and notified directly to the XDS Document Registry.

215 **3.2 Centralized Matching (Local ID Authoritative)**

In this other model, the source of XAD-PIDs is assumed by a central service, which also has the responsibility of matching local identifiers with the XAD-PID. This requires the use of a PIX Manager (or similar service) to create and manage the linkage sets between all known identifiers (i.e., all local patient identifiers + XAD-PIDs) as shown below:



**Figure 3.2-1 Centralized XAD-PID Matching**

220

The following conditions must be met for this model to work:

- 225 • The Document Source system must also be a PIX Patient Identity Source to the PIX Manager
- The XDS Affinity Domain will establish the business rules for assigning and managing the valid set of XAD-PIDs (could be provided by an external source or by the PIX Manager, if able to dynamically create XAD-PIDs)
- 230 • The PIX Manager will implement matching rules and algorithms (deterministic, probabilistic or both) that meets the minimum data quality standards established for the XDS Affinity Domain. Manual processes may be required to correct any questionable matches.

235 With the services shown above in place, the Document Source (or Document Consumer) needs only query the PIX Manager using its own local identifier and retrieve the matching XAD-PID. The local system will trust that the correct match has occurred and use that common identifier in all transactions with the XDS infrastructure.

In this model, the authoritative patient identifier for clinical data published to the XDS services is the local patient ID from the document source system; the XAD-PID can be seen as just an attribute, externally assigned by the PIX Manager that can change!

240 The fact that the local system has no knowledge how the XAD-PID was determined or that the relationship sets that define the matches between local identifiers can change means that a new level of communication is required between the PIX Manager and the XDS services.

245 As seen previously, XAD-PID merge events can still occur and will be handled just as described previously. However, in addition, we now must consider that link/unlink events can also occur within the PIX Manager and these are currently not addressed by the XDS specifications. This is one of the issues to be discussed in the next section.

## 4 Patient Identity Management Gaps in XDS Specifications

250 The description of the XAD-PID management schemes in the previous section highlights that there are areas in the XDS specifications that need to be re-visited to address existing gaps. This section will describe these gaps and offer options for their resolution

### 4.1 XAD Level Link/Unlink Events

255 A recent white paper from Canada<sup>1</sup> describes a main consequence of the CR to XDS integration, which is, the need to support changes to the XAD-PID linkage sets. If the XAD-PID assigned to a local ID never changed, then there would be no impact to the XDS specifications, but since this is not the case, how can the XDS infrastructure deal with the dynamic nature of the XAD-PID linkage sets?

260 The white paper provides a very good description of the use cases involved in EHR patient identification as implemented in Alberta. It also provides three suggestions that resolve the problem of how to notify link/unlink events to the document registry. In all cases, the solution requires that the document source system provide in the submission set the local identifier that was used to execute the client resolution (i.e., the PIX query).

265 To illustrate the scenario in discussion, let's assume that a patient presents to a service location in a given XDS Affinity Domain for the first time and that a set of documents from that encounter are published to the XDS infrastructure:

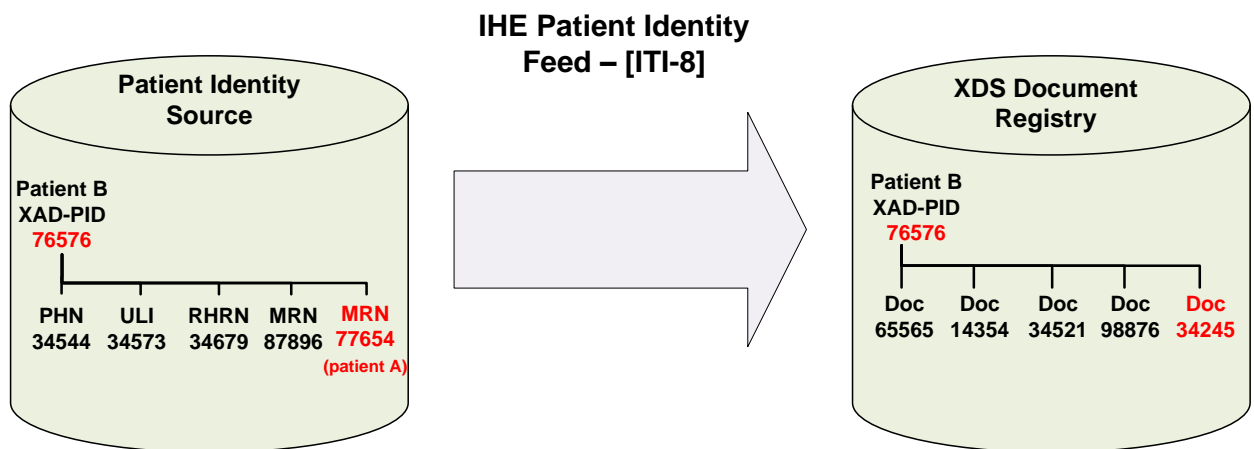


Figure 4.1-1 Initial State of New Document

<sup>1</sup> “Link/Unlink Analysis: pHIE – XDSi”, Alberta NetCare Health Information Exchange, Feb/2010

270 The local patient ID (MRN 77654) is mapped (i.e., linked) by the PIX manager to an existing XAD-PID (XAD-PID 76576). One or more documents are published to XDS using that common identifier.

However, at some later time, it is discovered that Patient A should not have been linked to that XAD-PID in the first place and that in fact, it should have been linked to another identifier as shown below:

275

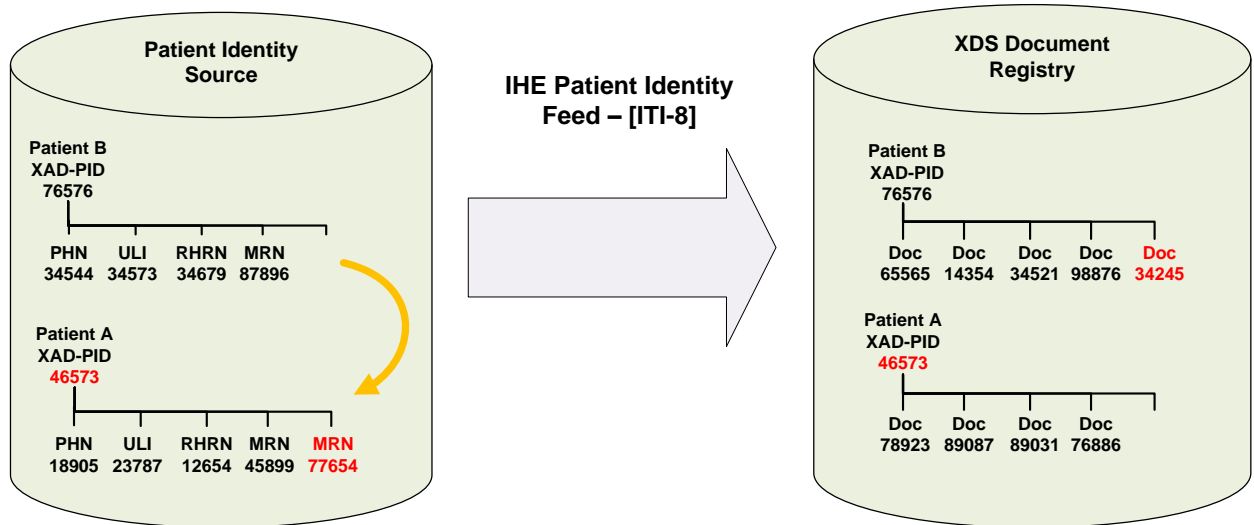


Figure 4.1-2 Link/Unlink MRN

280 In this case, we see that the correct XAD-PID is (XAD-PID 46573) and the change occurs within the XDS Affinity Domain patient ID source. However, the previously published document (DOC 34245) needs to be corrected and reflect this change. Given that the original document source system may not be aware of the link/unlink event, it cannot be expected to deprecate and re-publish the document itself.

285 The solutions proposed in the Canadian white paper try to address the link/unlink problem using existing ITI transactions, one using ITI-8 (currently used for patient identity feed between Patient Identity Source and the XDS registry) and the other two using ITI-30, defined as part of the Patient Administration Management (PAM) profile. A third option was later presented that tries to improve upon those suggestions.

#### 4.1.1 Alternative 1 – Using ADT^A31

290 In the first solution (Option 1), the ADT^A31 – Update Person Information message would be used by the document registry to determine if there has been any change to the linked IDs for a particular XAD-PID. This approach requires that a first message be sent to unlink the local ID from the current XAD-PID and a second message to re-link the same ID to a new XAD-PID. The advantage of this solution is that it uses the existing ITI-8 transaction and message set. On the

295 other hand, using two asynchronous messages to correct a single XAD-PID assignment is a concern since it could leave the registry in an incorrect state if one fails. It would also require the document registry to perform a differential comparison with the information it previously had about that patient and identify any changes to the local patient IDs assigned to the corresponding XAD-PID.

#### 4.1.2 Alternatives 2a and 2b – Using PAM

300 The two other solutions propose using the PAM profile (ITI-30) to manage the synchronization of patient identity information between the PIX manager and the document registry. These recommendations would require in a larger impact to the XDS profile since it changes the set of messages exchanged between the two actors.

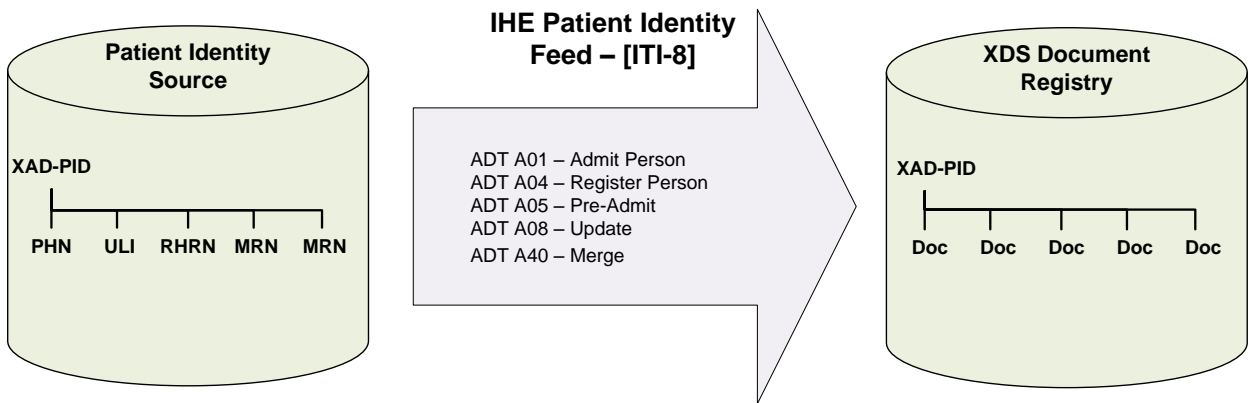
Variants for this approach are described based on the two PAM options:

- 305 1. Option 2a: Using merge – in this option, the link/unlink information is done through ADT^A47 – Change Patient Identifier List, very similar to how the CR uses ADT^A43 – Move Patient Information Event, to notify the EHR of link/unlink occurrences. This approach only requires a single message to complete the change.
- 310 2. Option 2b: Using Link/Unlink – in this option, two separate messages, ADT^A24 – Link Patient Information and ADT^37 – Unlink Patient Information are required. Similar to option 1 described previously, this approach uses two asynchronous messages and could create a situation where the XDS registry is left in an incorrect state.

315 The three solutions proposed in the document address the same problem: how the Patient Identity Source should notify the XDS document registry of changes to which XAD-PID is assigned to a particular local ID. In this sense, all three solutions can work, although preference should be made to Option 2a which does not rely on asynchronous pairs of messages.

#### 4.1.3 Alternative 3 – Using ADT^A43

Another option that could be simpler, is to add a new message to ITI-8 shown below:



320

**Figure 4.1.3-1 IHE ITI-8**

In this option, ADT^A43 would be added to the set of messages in ITI-8 or introduced in a new transaction. This notification message would be used strictly to inform the document registry that a new XAD-PID has been assigned to an existing local ID, using a message structure as shown below:

325

```
MSH|^~\&|EMPI|EHRI|CH|20100126113930||ADT^A43|20100126000022537083|D|2.4
EVN|A43|20100126113925
PID|46573^ECID^EMPI|77654^MRN^CH|||||||||||||||||||||||||||||||||||||N
MRG|77654^MRN^CH||76576^ECID^EMPI
```

**Figure 4.1.3-2 - Sample ADT^A43**

This approach has the same advantage as Option 2a (i.e., single message) and smaller impact as it does not require a change in the XDS integration profile from using ITI-8 to ITI-30 to synchronize patient identities. Since support for ADT^A43 would be optional, it would also not require changes to any XDS implementation that does not involve link/unlink events as described.

330

Note: the use of ITI-8 is only one option, here described because it already describes the set of messages exchanged between the Patient Identity Source and the XDS Document Registry. However, since this same transaction is also used in a number of other IHE integration profiles, it may be more appropriate to look at other alternatives, such as ITI-30 or maybe creating a new ITI transaction altogether.

335

#### 4.1.4 Use of XDS Metadata Update

340 A recent option, part of the 2010 workitems, that was considered is the use of the XDS Metadata  
Update (MU)<sup>2</sup> which has provisions for updating the PatientID information in folder and  
document entry objects. The use of the transaction Update Document Set [ITI-57] to resolve  
link/unlink events, although feasible, would result in a much more complex interaction model  
between the PIX Manager and the XDS Document Registry. Since Update Document Set must  
345 address each individual object impacted by the new association between local identifier and  
XAD-PID, it would require that the PIX Manager perform a Registry Stored Query [ITI-18] (as  
is the case with the Document Administrator) and create a (possibly) large submission set with  
all objects (folders and document entries) that would need to be changed. This type of capability  
is non trivial and likely well beyond the current capabilities of systems that implement the PIX  
Manager role. The use of existing HL7 v2 ADT messages, such as ADT^A43, seems like a much  
350 simpler approach, although it does transfer the complexity to the document registry actor.

However, as discussed in the next section, the impact to the document registry of patient ID changes is well documented in the XDS Metadata Update Supplement and should be considered when finalizing the decision on the link/unlink support.

## 4.2 Additional Considerations

355 However, the white paper does not address some key impacts to the XDS implementation that  
will occur regardless of which messaging approach is adopted. Although it does note the need for  
document sources to include the local identifier with each submission set, there are three other  
issues that need further clarification:

### 4.2.1 Impact to Submission Sets and Folders

360 Link/Unlink events can cause inconsistencies in submission sets or folders, where documents  
from two (now) different patients are grouped together. The issue here is that when folders and  
submission sets are first created, they are validated to ensure that all documents they contain  
belong to the same patient (i.e., they have the same XDS Affinity Domain patient ID). After a  
link/unlink event, it is possible that this condition is no longer valid, as one or more documents  
365 that used to belong to a particular XAD-PID is now changed to a different identifier.

One solution would be to include a new (optional) restriction where documents within  
submission sets and folders must have the same local patient ID. This would ensure that  
regardless of what happens in a link/unlink scenario, that integrity of the objects in respect  
to patient identity is preserved. However, this may be too restrictive in regards to folders, where  
370 some implementations are using them to combine documents from various sources.

---

<sup>2</sup> IHE IT Infrastructure Technical Framework Supplement, XDS Metadata Update (MU) – Draft for Public Comment (June 11, 2010)

375 A more pragmatic option would be to maintain the current rules (i.e., documents from the same patient, regardless of source) but add a new behaviour to the registry. If, in consequence of a link/unlink event, one or more documents within a given folder no longer have the same XDS Affinity Domain patient ID than these must be logically removed from the folder. In addition, the XDS document registry must generate an exception event (possibly through a report) to notify that such action has taken place and allow for subsequent corrective measures by the document source.

This is similar to what is recommended in XDS MU (page 17):

380 *“An update that changes the Patient ID attribute of a Folder or DocumentEntry is more complicated. The rules for consistency of Patient ID between Folder and member DocumentEntries require that if a DocumentEntry gets a new Patient ID then it must be removed from the Folder (assuming the Folder does not change). The association propagation rules implemented by the Document Registry actor do not handle this type of update. The Document Administrator must calculate and submit all of the detail changes required.”*

385

The key difference is that unlike what is recommended above, the burden of removing the affected documents from folders when patient ID surface would fall upon the Document Registry and not the system that triggered the link/unlink event (e.g., PIX Manager)

390 This change should not create an issue with submission sets, since documents from the same document source would likely have the same local ID. However, the same may not be true in regards to XDS folders. Since folders are not widely used, it is difficult to assess if this restriction would create implementation barriers.

#### **4.2.2 Impact to Document Repositories**

395 Link/Unlink event notifications are not propagated to document repositories. The issue here is that document repositories may also contain a copy of the metadata published with each submission set. If this is the case, than existing record would no longer be correct after the link/unlink event occurs.

400 This situation is similar to the patient demographic information kept in the metadata record. It is not guaranteed to be correct over the lifetime of a document, and only reflects the original state of the record. If this is acceptable, and document consumers are aware of possible discrepancies, than no further changes are required. Otherwise, the same message used by the document registry to handle link/unlink events will have to be passed on to all document repositories in a XDS Affinity Domain.

405 Note: Many of the current document repository implementations do not persist the document metadata, in which case this issue is not applicable.

### 4.2.3 Impact of Local Merge Events

Merge/Unmerge of local patient Identifiers need to be propagated to the document registry. Since the XDS services are now made aware and rely upon the correctness of the local patient identifiers, it should be informed of any merge/unmerge events that effect that information.

- 410 Here is the scenario: two local identifiers (Lid-A and Lid-B) have been in use for some time but have been determined to belong to the same patient (i.e., source system duplicates). A merge event occurs in the local system and is notified to the PIX Manager, where Lid-A is subsumed by Lid-B. Within the PIX Manager, Lid-B is the only surviving ID but the XDS infrastructure still has one or more documents attached to Lid-A. Since (presumably) Lid-A and Lid-B were linked
- 415 to the same XAD-PID (let's say XAD-PID-123), this situation is not a problem as all queries to the registry would use that same XAD-PID number. However, if in the (not very likely) situation where Lid-B is now determined to be linked to the wrong XAD-PID (i.e., it should be linked to XAD-PID-987), we would have a problem. The link/unlink solutions here described would see a notification that all documents belonging to Lid-B should be now assigned to XAD-PID-987.
- 420 Nothing would be said about Lid-A since it has been for all purposes subsumed. At this point, the document registry would be left in an incorrect state, where Lid-A documents remain attached to XAD-PID-123 while Lid-B documents are moved to XAD-PID-987.

- 425 There are different ways to address this problem; most require the propagation of local merge/unmerge events to the XDS infrastructure. This will add a new level of requirements to the XDS integration model and may only be required in very rare instances. Alternatively, one may define a solution where the client identity source needs to include all merged identifiers (i.e., all subsumed local Ids) when creating a notification of a link/unlink event.

## 5 Recommendations

430 This white paper has discussed the various implementation schemes that are used for patient identification management and how they can impact the XDS profiles. In particular, it has addressed the scenario where changes to the link between a local identifier and the XDS Affinity Domain patient identifier (XAD-PID) need to be propagated to the XDS infrastructure services (i.e., document registry). This has been identified as an issue in several XDS implementations in Canada.

435 Considering the various options presented, the option that seems to address the problem with the least impact to the existing XDS profiles is Alternative 3, that recommends to introduce a new optional message (ADT^43) either in an existing Transaction (e.g., Patient Identity Feed [ITI-8] transaction) or as a new Transaction.

440 This message would inform that a local ID (LID) that was previously linked to XAD-PID(a) is now linked to XAD-PID(b). The receipt of this message by the document registry would result in a global change within the registry of all objects that contain the {XAD-PID(a), LID} should be changed to {XAD-PID(b), LID}.

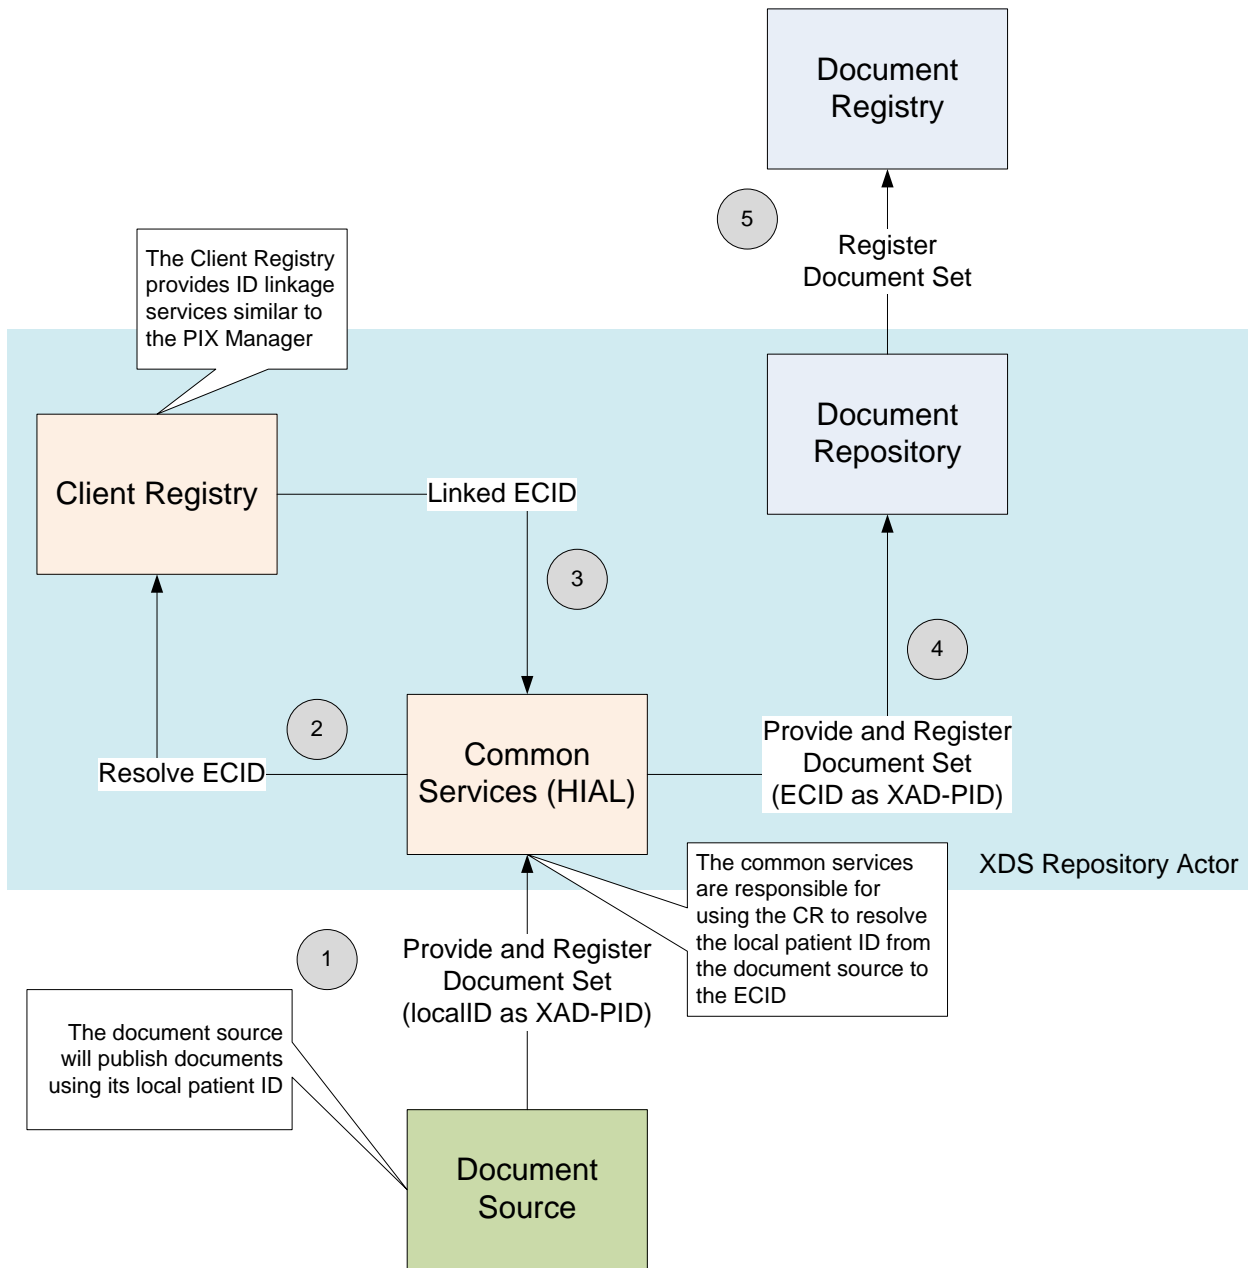
445 In addition, this white paper further recommends additional discussion in respect to how to handle the impact of link/unlinks on folders and submission sets previously published to the XDS infrastructure.

## Appendix A: PIX/PDQ Integration Models with XDS

### A.1 Canadian Interoperable EHR Blueprint

450 In the Canadian interoperable EHR architecture, the responsibility for managing and cross referencing client (i.e., patient) identities often falls with a central EMPI services known as the Client Registry (CR). The CR collects registration information from various patient identity domains in its jurisdiction and groups these records together through a combination of automatic algorithms and manual linking events. These sets of linked IDs are given a unique identifier, known as the Enterprise Client Identifier (ECID), which is analogous to the IHE XDS Affinity Domain patient identifier (XAD-PID).

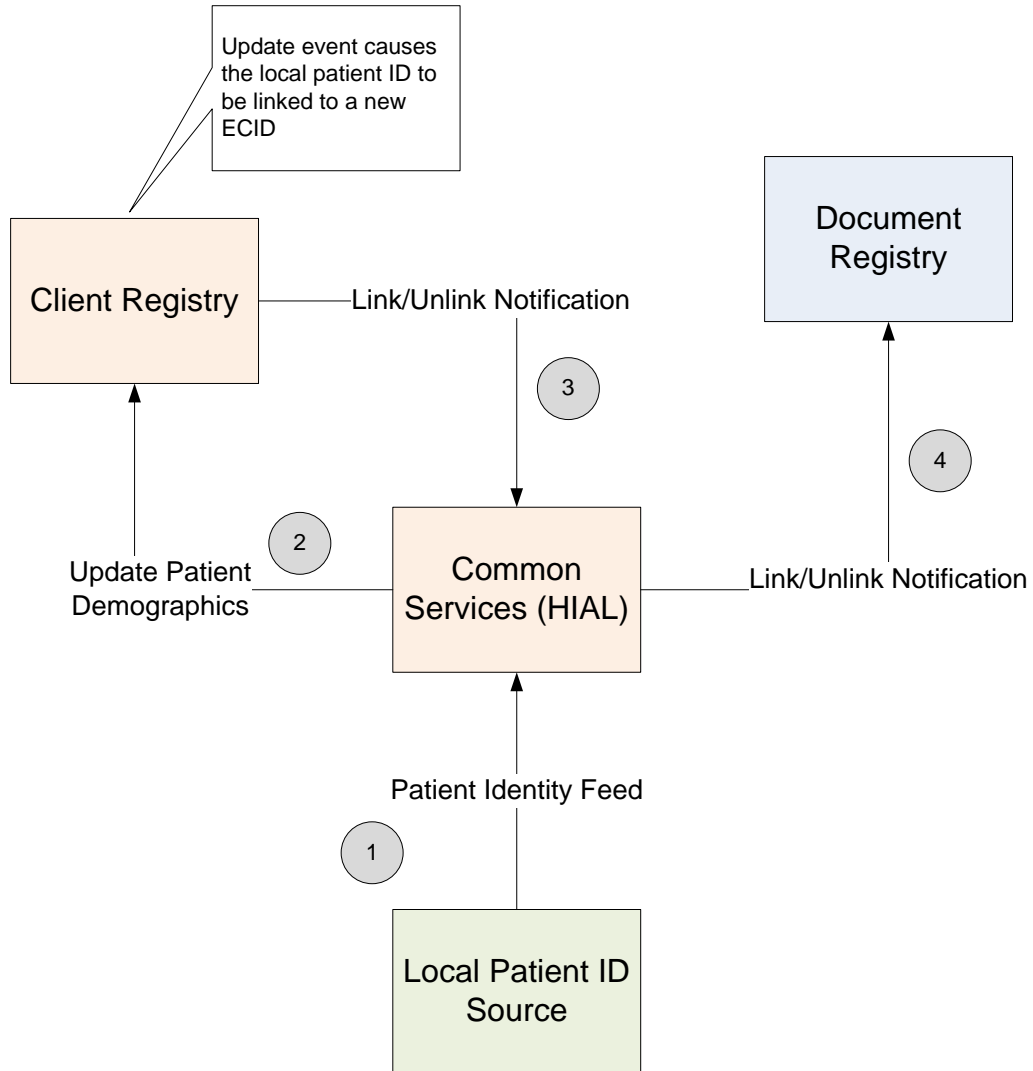
455 The ECID can either be a new ID created by the CR itself (i.e., shadow ECID) or can be assigned (through jurisdictional policy) by one of the patient identity domain (i.e., provincial health card number). Regardless of the model, the CR is the only authoritative source of ECIDs for the jurisdictional EHR infrastructure. In most cases, when documents are published to the XDS the actual resolution of local ID to ECID is not performed by the document source, but rather by  
460 other common services (i.e., the HIAL) in the XDS Affinity Domain.



**Figure A.1-1 Centralized XAD-PID Matching**

465 For the most part, the jurisdictional CR role is very close to the services provided by a combined PIX/PDQ manager. The most significant difference is that the relationship local ID → ECID is not guaranteed to be permanent. In fact, as described in the Alberta Link/Unlink white paper, it can very often change because a new local patient ID will be given an initial ECID when it is first published to the CR. It may or may not remain the same, depending on the result of the

470 matching algorithm and possibly manual linking events. If a match is found, the local ID is moved from the original ECID to another, pre-existing ECID. In this case, the original ECID is deprecated and no longer in use<sup>3</sup>. There are also other scenarios where, for different reasons, the original match of a local ID may be wrong and a new ECID assignment is required.



**Figure A.1-2 Link/Unlink Events**

475 In all these cases, the impact to the XDS services occurs if there is one or more published documents belonging to the local patient ID in question. Since this patient identifier is now

---

<sup>3</sup> This is similar to a patient identity merge described previously since the original ECID was a singleton, that is, it only contained one linked local ID.

linked (i.e., assigned) to a different ECID, the XDS services needs to be informed about the event and its information corrected.

480 When these change events occur, the CR will trigger notification messages to any system that relies on the ECID. Typically, these notifications will be forwarded to another jurisdictional service, called the Health Information Access Layer (HIAL), responsible for routing such messages to all appropriate “listeners”. The XDS document registry would need to be one of the recipients of the link/unlink notifications.

485 These examples highlight the key nature of the issue: in the CR model, the authoritative patient identifier for clinical data published to the EHR (and this would include XDS documents) is the local patient ID from the POS source system; the ECID can be seen as just an attribute, externally assigned by the Client Registry that can change!

Note: There are other different EHR client identification models applied in Canada:

- 490 1. In some cases (i.e., the province of Quebec), the ECID is a public identifier that must be resolved by each EHR source or consumer system, much as is described in the XDS interaction model. This approach does not require the support of link/unlink and can easily adopt XDS solutions as specified.
- 495 2. In other places (i.e., province of Manitoba), the CR manages the linked list as described but the ECID is not used by EHR repositories. Records are created in the EHR just with the local identifiers and linked together, via the CR linked set, dynamically at query time. Link/unlink events are handled exclusively by the CR and repositories need not be notified. However, this model does not fit well with the XDS specifications which require a single patient identification domain. It is likely that the ECID would be used for XDS documents, bringing us back to the model in discussion.

500

## A.2 Other Integration Models

505 This section provides an inventory of a number of approaches for patient identification management in a document sharing environment. Each one of these approaches is described at a high-level. This list may need to be expanded as further experience and diverse implementation strategies are being considered.

The patient identification management strategies have a variety of characteristics. Three discriminatory characteristics are used in the table below to distinguish among the various approaches:

510 A - Source for the patient ID used as Affinity Domain Patient ID (XAD-PID)

515 This differentiating characteristic is related to the process/actor that is in charge of assigning patient identifiers used for the affinity domain (patient id used by the XDS Registry in the Metadata attribute). Example of such sources, may be the system managing the citizen identifiers associated with national healthcard, a participant registration systems where patient's consent for health information exchange is managed.

B - Use or non-use of Local Patient IDs

520 This differentiating characteristic is related to the role played by local patient IDs (i.e., those used internally by Document Sources and Document Consumers). For example, the patient ID used locally may be a regional patient identifier, hence no local identifiers are visible (not used) within the document sharing environment. In other cases there may be a cross-referencing service (e.g., an MPI or PIX Manager) that links local patient IDs (used) and affinity domain IDs.

C – Designation of the Authoritative Patient ID

525 When multiple patient identifying schemes need to coexist, their respective “strength” or “stability” need to be qualified and understood, so that when updates and corrections are applied, consistency rules can be applied along with any changes or updates.

530

<b>Characteristics</b> <b>Approaches</b>	<b>Source of for the patient ID used as Affinity Domain patient ID (XAD-PID)</b>	<b>Use or non-use of Local Patient IDs</b>	<b>Designation of the Authoritative Patient ID</b>	<b>IHE Profiles Used</b>
<b>1 -Affinity Domain MPI</b>	Internal MPI algorithm	Use	Local Patient ID	XDS, PIX
<b>2 -Affinity Domain MPI with external ID Registration</b>	Affinity Domain central assignment service (e.g., opt-in consent management, health card issuance, registration service).	Use	Local Patient ID	XDS, PIX
<b>3 - Decentralized Matching</b>	Affinity Domain central assignment service (e.g., opt-in consent management, health card issuance, registration service).	Non-use	XAD-PID	XDS, PDQ
<b>4 - Multi Community Federated Matching</b>	Each Community assigns XAD-PID	N/A	N/A	XCA, XCPD
<b>5 – Single Community Federated matching</b>	No XAD-PID assigned	Use	Local Patient ID	XDS and either PIX or implicit XCPD.
<b>6 - Multi-Community Shared Patient ID</b>	Cross-community central assignment service (e.g., opt-in consent management, health card issuance, registration service).	N/A	N/A	XCA, XCPD (only location discovery)

### 1 -Affinity Domain MPI

535 In this approach, the MPI is used and is capable not only to cross-referencing patient IDs from different local patient ID domains (e.g., hospitals), but to assign these linked IDs its own XDS Affinity Domain Patient Identifier (often an internal one). In this case it acts as a source of the affinity domain ID that will be used by the XDS Registry (e.g., received via ITI-44)

In this case, the "PIX Manager" Actor and the "XDS Affinity Domain patient ID Feed" Actors are grouped together (they are supported by the same system, let's call it an MPI system). It is very important to understand that IHE specifies an integration profile as transactions between

540 Actors. In this situation one is combining the use of two profiles (PIX and XDS) where Actors of these profiles are grouped:

- a. "PIX Manager" is grouped with XDS "patient ID Source"
- b. "PIX Consumer" is grouped with XDS "Doc Source"
- c. "PIX Consumer" is grouped with XDS "Doc Consumer"

## 545 **2 -Affinity Domain MPI with external ID Registration**

In this approach, there is an independent source of identities for the affinity domain. For example, associated with the process of the patient granting patient consent, before any documents are shared (may support IHE BPPC). Such a consent creation application requests the patient to provide its demographics, and it becomes a source of affinity domain identities.

550 This affinity domain patient id is fed to both the PIX Manager and the XDS registry. The PIX manager registers this ID and its tasks to link it with local patient IDs whenever they are fed to the PIX Manager/MPI. The local patient ids and the affinity domain patient IDs once cross-referenced by the MPI, become accessible via the PIX Manager to the edge systems that may query (ITI-45) based on their local ID, and get back the shared patient ID.

## 555 **3 - Decentralized Matching**

In this approach, there is an independent source of identities for the affinity domain and these patient identifiers are used as local patient identifiers by Document Sources and Document Consumers. In this approach Document Sources and Document Consumers may be grouped with a PDQ Demographics Consumer Actor if there is a need to look-up the Affinity Domain ID (XAD-PID) at the time of document publication or query. Example of use of such patient identifiers are shared patient ID such as those provided by the use of a national health card, or some form of fully deployed shared patient identification.

560

## **4 –Multi Community Federated Matching**

565 In this approach one has several XDS Affinity Domains that are federated through the use of XCA and XCPD where XCPD is used to match XDA-PID across the various affinity domains.

## **5 -Single Community Federated matching**

570 In this approach one has a degenerated case of the multi-community federated matching: (1) the XDS Document Registries across the various affinity domains are supported by a single Document Registry system and (2) each affinity domain has only one system supporting a Document Source/Consumer. XDS Registry serving an Affinity Domain which is configured in a specific manner, which allows the use of XDS without a shared affinity domain patient ID. Each Document Source or Document Consumer is using its local patient ID as the XAD-PID in its

575 interactions with the XDS Repository/Registry. This results in considering each edge system acting as a document source (or document consumer) as a “mini-affinity domain” where the local patient ID is the one used by the XDS Registry. Although there is only one XDS Document Registry system, it actually is made of a federated (within one system) of the XDS Document Registries across all mini-affinity domains. one is actually creating a virtually federated affinity domain, where, each document consumer queries its “own” registry, which includes an “implicit XCA/XCPD” initiating gateway that discovers the matching patient IDs in 580 the other XDS Registries of the mini XDS affinity domains and federates the query responses. The same result can be achieved with the virtual use of PIX to match across the mini-affinity domain patient IDs. A consequence of this specific configuration of XDS Actors allows the use of XDS without a shared affinity domain patient ID.

### **6 - Multi-Community Shared Patient ID**

585 This approach is similar to Approach 4 - Multi-Community Federated Patient ID, with one major difference. XCPD need not be used to match XAD-PID across affinity domains, as all affinity domains are using the same shared patient id issued by an external source. Only the location of patient information of XCPD may be used.

590