Connected Care

Reference Architecture for Interoperability

Architectural requirements for electronic referral, transfer of care and shared care solutions in the New Zealand health sector

Version 1.0

September 2012

Document Control

This document is part of the Reference Architecture for Interoperability. Its title is *Reference Architecture for Interoperability – Connected Care*.

i Revision History

Version	Release Date	Author	Status
0.6	February 2012	Sector Architects Interoperability TWG	Release for review by public sector architects
0.9	June 2012	Sector Architects Interoperability TWG	Release via Health IT Cluster for review by vendors
1.0	September 2012	Sector Architects Interoperability TWG	Updates in response to vendor feedback

Contents

Introduction

This document extends the *Reference Architecture for Interoperability* for the New Zealand health sector, adding to our original principles a set of architectural requirements specific to electronic referral, transfer of care and shared care solutions – or what we term connected care solutions.

Quite generally, connected care is about interconnectedness in healthcare through ready common access to patient information at point of care, whether in community, hospital or residential care settings. The key shared information resources are referral requests, clinical assessments, test results, medication lists, care plans, discharge summaries and other shared health summaries.

From our basic rules, interoperability in the connected care domain then concerns applications being able to produce and consume such resources in the form of standardised datasets and structured documents, and exchange them via web services.

ii Purpose

This specification provides a definitive set of rules for the interoperability of referral, transfer of care and share care solutions. Its first purpose is as a guide to implementers in shaping new solutions of this kind. Second, it scopes new interface standards.

The interoperability reference architecture as a whole exists in support of the *National Health IT Plan*. It is published as part of the health sector architecture library.

iii Scope

This specification is expressly about interoperability in relation to connected care. It builds squarely on the *Health Information Exchange Standard (HISO 10040)* and the earlier foundational parts of the interoperability reference architecture, applying the overarching rules of these parent specifications to this particular domain.

In this specification we describe appropriate mechanisms for referral and transfer of care information flows. And we extend the episodic workings of referral and transfer of care to the more fully collaborative processes of shared care, which feature the requirement for shared access to such resources as the problems list, the medications list and the care plan.

We reiterate prescribed interface standards and outline the necessary datasets, document types and web services in this domain.

We distinguish the process of referral from that of ordering tests. Dealing with test orders and results is a topic in its own right and is not in scope here.

Relevant business process descriptions can be found in the *Referrals, Statuses and Discharges Standard (HISO 10011)* and the referenced shared care white paper endorsed by the National Health IT Board.

iv Key Concepts

The interoperability reference architecture is predicated on a model of information sharing centred on *regional clinical data repositories (R-CDRs).* Under this model, point-of-service systems connect to *health information exchanges (HIEs)*, sharing repository-held data resources via web services. Regional HIEs interconnect to form a nationwide fabric of information exchange services.

We achieve interoperability by standardising on datasets and terminologies, working from internationally accepted specifications wherever possible. Participating systems are required to have clearly defined functional boundaries and interfaces.

Our information flows are by and large document-oriented, with a handful of structured document types serving routine purposes. For example, one of our tenets is that referral requests should be expressed generically – each of them represented as a document following exactly the same template. The same applies for discharge summaries.

More broadly, the ability to produce and consume a standardised *health status summary* is a new requirement for point-of-service systems. This is a dataset that represents the current health status of the individual in terms of allergies, disabilities, problems, test results, medications, and the care plan, among other relevant measures.

Alongside this, we introduce a generic *health event summary*. This dataset captures essential details of any healthcare the patient receives in community, primary care or outpatient settings – it is akin to a hospital discharge summary. It could also convey a response to a request for clinical advice.

We focus on one other important dataset – the transfer of care medications list, which is essentially the same as the medications changes summary of the discharge summary, but warrants its own independent existence for certain use cases.

For each of these datasets there is a document type for information exchange.

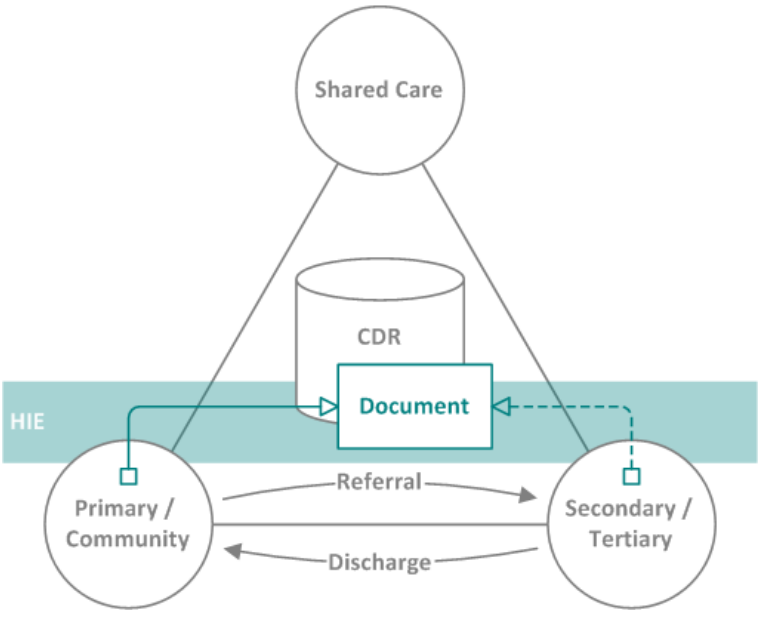

Figure 1 – Sharing transfer of care documents out of an R-CDR

Figure 1 illustrates the foundational concept of information sharing via a repository. In this particular example, a community provider posts a referral document to the repository, where its arrival is notified to intended recipients in the hospital.

Fronted by web services, the repository is our intended vehicle for information sharing in connected care.

v Document Structure

The normative content of this document is presented as sets of directives, under these headings:

1. Structured Documents
2. Document Sharing
3. Referral Services
4. Workflow
5. Identity
6. Privacy and Security
7. Referral Management Systems
8. Referral Data Entry
9. Shared Care Records
10. Business Intelligence

Each directive is presented as a simple statement expressing a particular rule or requirement as it pertains to the desired future state. This is followed by an explanation of that rule or requirement, with a discussion of its rationale and any implications it may have.

An appendix discusses certain key impacts of the reference architecture – the shaping of new standards and new styles of solution:

- HISO Standards Development
- Future State Migration
- Practice Management System Interfaces
- Transfer of Care Medications List

vi References

Related standards, reference architectures and other specifications are the following:

- National Health IT Plan [www.ithealthboard.health.nz/]
- Reference Architecture for Interoperability [www.ithealthboard.health.nz/content/sector-architects-group/]
- HISO 10040 Health Information Exchange Architecture Building Blocks [www.ithealthboard.health.nz/standards-open-for-public-comment/]

- High Level Requirements for Transfer of Care between Health Practitioners [www.ithealthboard.health.nz/]
- Transfer of Care (eDischarge) Processes and Requirements [www.ithealthboard.health.nz/sites/all/files/TOC_0.doc]
- Shared Care White Paper – Enabling Integrated Healthcare [www.ithealthboard.health.nz/content/shared-care-paper-enabling-integrated-health-care/]
- HISO 10011 Referrals, Statuses and Discharges Standard [www.ithealthboard.health.nz/referrals-status-discharges/]
- HISO 10041 Transfer of Care Standard (in development)
- HISO 10043 CDA Common Templates Standard (in development)
- HISO 10014.2 Online Forms Standard [www.ithealthboard.health.nz/online-forms-architecture-technical-specification/]
- HISO 10037 Connected Health Standard [www.ithealthboard.health.nz/connected-health-network-connectivity-standards/
- HISO 10029 Health Information Security Framework [www.ithealthboard.health.nz/hisf/]

1 Structured Documents

The set of directives in this section cover the requirement that connected care information be exchanged in the particular document-oriented style specified by HISO 10040 on HIEs. We outline necessary structured document types for referral, transfer of care and frequent transactions in shared care.

See HISO 10040.1 for general rules on structured documents.

1.1 Connected care information is exchanged primarily in the form of *HL7 Clinical Document Architecture (CDA)* structured documents

The CDA document is our chosen unit of currency for information exchange in the connected care domain, as elsewhere. Referral requests, discharge summaries and shared care record extracts should all be represented as CDA Release 2 documents at system interfaces and in communications. Furthermore, document instances should always have fully structured content, with 'level 3' coded entries present throughout.

We permit attachments and references to other documents, provided the usual CDA linking mechanisms are used.

1.2 The document types in use for connected care include the referral request, discharge summary, health status summary, health event summary and transfer of care medications list – each type conforming to its own definitive set of CDA templates

These are the document types we require as a minimal working set, covering the main connected care use cases. Inevitably, other document types will be added to the list – for example, assessment reports. However, we intend these to be generic document types and never expect to see terribly many of them.

Most of the CDA section templates we require here will be common to multiple document templates – this will constitute a family of templates.

1.3 Referral requests of all kinds are represented using a certain generic document type (especial to this purpose) and conform to the same set of CDA templates

In this context we include requests for advice as a kind of referral.

Under this requirement all referral requests, whatever the target clinical service, are represented as CDA documents of the same generic type. That is, these documents all conform to the very same CDA document template and component section templates.

> Note that these rules do not prevent there being condition-specific referral *forms*, as opposed to documents. Such forms are acceptable as long as the data captured at the user interface can be represented generically behind the scenes. For example, the existence of some clinical indicator might be conveniently presented as a checkbox item on screen, while being represented in the resulting document as simply another SNOMED-coded entry in a list.

In this directive, we are acknowledging differing levels of product capability, and try – by way of a universal specification – to remove point-to-point dependencies as a barrier to interoperability in multi-party referral solutions. Another benefit of a standard-

ised representation is that repository-held referral documents become useful as lasting entries in the longitudinal health record.

Clearly, under these conditions, the actual templates we come up with have to be suitably expressive. CDA is perfectly flexible in this regard, so this is more a matter to consider in designing the underlying dataset and in terminology choices.

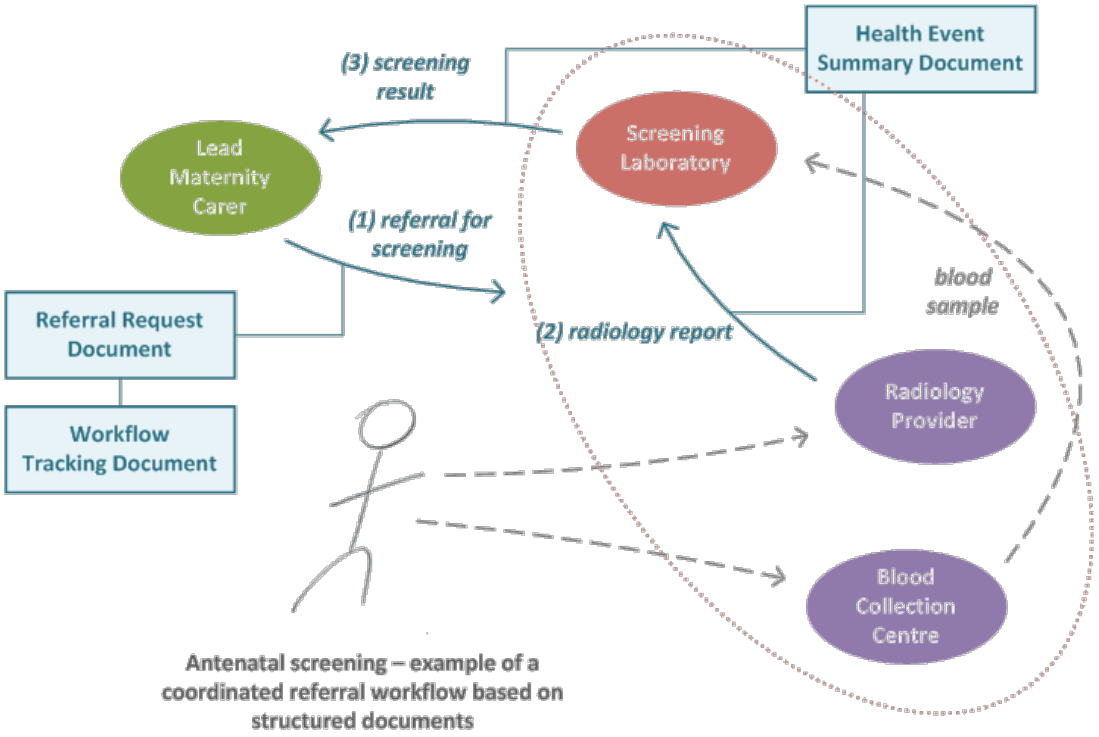

Figure 2 – Example referral process based on a flow of CDA documents

1.4 Discharge summaries of all kinds are represented using a certain generic document type (especial to this purpose) and conform to the same set of CDA templates

This requirement is analogous to the one above on referrals. All discharge summaries are represented as CDA documents of the same generic type, conforming to the same set of CDA templates.

1.5 A standardised health status summary document serves as a multi-purpose shared care extract

The ability to exchange health status summary documents is a minimum requirement of shared care systems. These documents convey the most relevant administrative, demographic and clinical information about the current health status and care of the individual concerned.

The health status summary dataset breaks into the sections listed in the following table. The population of each section in actual document instances is controlled by a set of purposeful restrictions over subject matter, timeframe and level of detail.

References entries are links to selected assessments, clinic letters or test results, for example, the detail of which lies in other repository-held documents.

Demographics	Enrolments	Healthcare Providers
Medications	Test Results	Encounters
Family History	Social History	Problems
Support People	Functional Status	Care Plan
Immunisations	Procedures	References
Allergies and Adverse Reactions		

Table 1 – Sections of the health status summary dataset

The purposes of the health status summary include the following:

- Providing summary information for display at point of care
- Pre-populating referral forms and health event summaries
- Providing input when clinical decision support functions are invoked
- Instantiating shared care records and as a share care record extract

1.6 Health status summaries are represented using a certain generic document type (especial to this purpose) and conform to the same set of CDA templates

While the health status summary document is for now based on the native CDA XML schema, using equivalent but simpler XML schemas in the *Green CDA* style is a strong possibility for the future.

Green CDA permits arbitrary XML schemas in place of the regular CDA schema, provided a mapping – from the former to the latter – has been worked out. Green CDA has the advantage that it does not require a narrative section, freeing fully structured documents such as the health status summary of this overhead.

1.7 There is a standard health event summary document type

We require that there be a standard set of CDA templates for the representation of health event summaries.

1.8 There is a standard transfer of care medications document type

We require that there be a standard set of CDA templates for the representation of the transfer of care medications list. Recall that this derives in large part from the medications changes summary of the discharge summary.

1.9 Connected care documents of all the types defined have standardised datasets that derive from the *ASTM Continuity of Care Record (CCR)*, with SNOMED CT as the default terminology choice

These datasets are all modelled on CCR – adopting reliably constrained and localised versions of that specification's subject area and data element definitions. Data elements have explicit bindings to SNOMED CT wherever suitable reference sets exist; elsewhere, to other approved code sets.

CCR is the chosen logical data reference model not only here but for all our e-health datasets. It covers problems, medications, test results and other information useful whatever the clinical setting. Its generic structures make possible the consistent representation of data from different clinical specialties.

1.10 Connected care dataset specifications are expressed both as openEHR archetypes and as equivalent ISO/IEC 11179 dataset and data element specifications

These are the two methods recommended for representing our data definitions; they are largely complementary. The *openEHR Clinical Knowledge Manager (CKM)* is a suitable tool for the collaborative development of archetypes as expressions of what we call detailed clinical models (DCMs).

Finished archetypes are then submitted a formal registration process according to the ISO/IEC 11179 methodology. The resulting data element specifications are expanded to include contextual details such as relate to provenance, quality and appropriate use of the information.

These methods support an open design process and help foster eventual use of the definitions.

HISO 10040.2 states general requirements for the *HIE Content Model*, which these dataset specifications are naturally part of.

1.11 CDA templates for connected care derive from (a) the corresponding dataset specifications, and (b) the internationally established templates of the *ASTM/HL7 Continuity of Care Document (CCD)* and its close relations

The design of CDA templates for transfer of care, as in other domains, derives from properly formulated and approved dataset specifications. In the interest of fidelity, we require that there be an explicit, granular mapping from templates to the respective datasets.

Templates are based where possible on CCD and the closely related specifications, *HL7 Consolidated CDA (CCDA)* and *IHE Cross Enterprise Sharing of Medical Summaries (XDS-MS)*. We also require that CDA templates be (1) documented in terms of the CDA RMIM, and (2) formulated as sets of conformance statements that can be easily translated to Schematron assertions.

1.12 CDA templates for transfer of care contribute to a national template library

Newly developed CDA templates will be submitted for inclusion in a national template library, whatever form that takes. Our recommendation would be for a HISO-governed registration process, with the templates incorporated into formal HISO specifications.

Templates in computable form as Schematron rule sets will likely be made available, for the benefit of implementers, in an online interface-testing tool.

2 Document Sharing

Directives in this section describe the use of document repositories as the primary means of sharing transfer of care information, with one such repository per region serving as the recognised distribution point for all referral and discharge documents.

HISO 10040.1 states our general requirements for document sharing in terms of the *IHE Cross Enterprise Document Sharing (XDS)* integration profile. XDS itself is based on *OASIS ebXML* for e-business transactions and *WS-I Basic Profile* web services.

2.1 Transfer of care documents are exchanged under a publish-subscribe model out of an XDS-enabled document repository in the R-CDR

The basic requirement here is that all referral and transfer of care documents should reach their audiences via placement in a designated repository. The creator of a document publishes it in this way, allowing its availability to be notified to the authorised recipients – who may be subscribing members of the patient's care team or other named individuals or service organisations. The document can then be accessed in situ or downloaded to a local system, as users choose.

This is about leveraging the regional information platform to enable secure, effective and reliable information sharing in transfer of care.

Notification mechanisms are key – to acknowledge the receipt of new referrals, provide referral status updates, and to advise of the availability of new discharge documents. Other notifications that a practitioner might receive include advice of hospital admission or of enrolment in a shared care programme.

Notification services under the publish-subscribe model implement the *IHE Document Metadata Subscription (DSUB)* integration profile, which means use of SOAP web services as specified by *OASIS WS-Notification*. RESTful notification based on secure RSS or ATOM feeds is an acceptable alternative. In either case, R-CDRs must have filtering mechanisms (XDS prescribes XPath as the query language) allowing incoming documents to be tested against notification criteria.

As a point-to-point counterpart to these publish-subscribe options, the *IHE Notification of Document Availability (NAV)* integration profile caters for SMTP email notifications. However, this is not a recommended approach – except perhaps in the consumer space.

2.2 Transfer of care dataset integrity is preserved between document submission to the repository and subsequent retrieval

We want to allow for the possibility of repositories not storing referral and discharge documents in native CDA format. While native storage means that all of the XML elements contained in the documents could be reproduced with unblemished fidelity between submission and retrieval, this would be an unreasonably stringent requirement. We don't wish to force the use of native XML database management systems when relational systems are superior in other ways.

Instead we impose the lesser requirement of dataset integrity. This requires the dataset for any particular instance of transfer of care to be reproducible as the very same set of data elements, with the same values and encodings that it was at the time of submission.

This is then a measure that ensures that content is not compromised by our repository based dissemination process, even when storing documents natively is an obstacle.

2.3 Transfer of care documents are retained in an R-CDR document repository as elements of the patient's longitudinal health record

This requires that every referral or discharge document submitted become available – preferably immediately – as a member of the complete set of documents held about the patient as a lasting record. Naturally, the actual use of transfer of care documents would be subject to specific information governance rules – as would the retention period.

A common national *XDS affinity domain policy* will have rules requiring documents relating to the same instance of transfer of care to be stored in the same XDS folder in the repository. The parent document and its attachments will be conveyed in the same XDS submission set on entry to the repository. The affinity domain policy will also map from a standardised CDA document header to XDS registry metadata entries.

2.4 Repository-held transfer of care documents are individually addressable by URL

XDS requires repository documents to be uniquely numbered but – for the sake of RESTful web applications and inter-document references – each document should also be assigned a URL. The ebXML specification describes the recommended scheme.

3 Referral Services

This section describes a base set of transfer of care web services – mostly for referral – that clinicians and consumers can expect to have available to them ubiquitously at point-of-service.

3.1 The transfer of care interface exposed to point-of-service systems is a set of XDS-enabled web services, including forms services, submission services, notification services and lookup/retrieval services

Referral web services are XDS-enabled in order to exchange documents with the R-CDR. These web services implement *XDS Document Source* and *XDS Document Consumer* actors and transaction types, enabling store and retrieval operations against target repositories.

Figure 3 presents a template solution for provider access to referral web services, which include the following:

- *Referral forms service* – Generating pre-populated HTML forms for the given patient, provider and clinical service (offering menu selections where necessary)
- *Referral submission service* – Creating CDA documents from completed referral forms and submitting them to the repository
- *Referral notification service* – Notifying recipients of incoming referral or discharge documents to the repository (the PMS can poll this service and updates itself accordingly)
- *Referral lookup/retrieval service* – Retrieving referral or discharge documents from the repository and (optionally) rendering them for display

Patient identity, location, provider directory, clinical pathway and decision support services are orchestrated with these services.

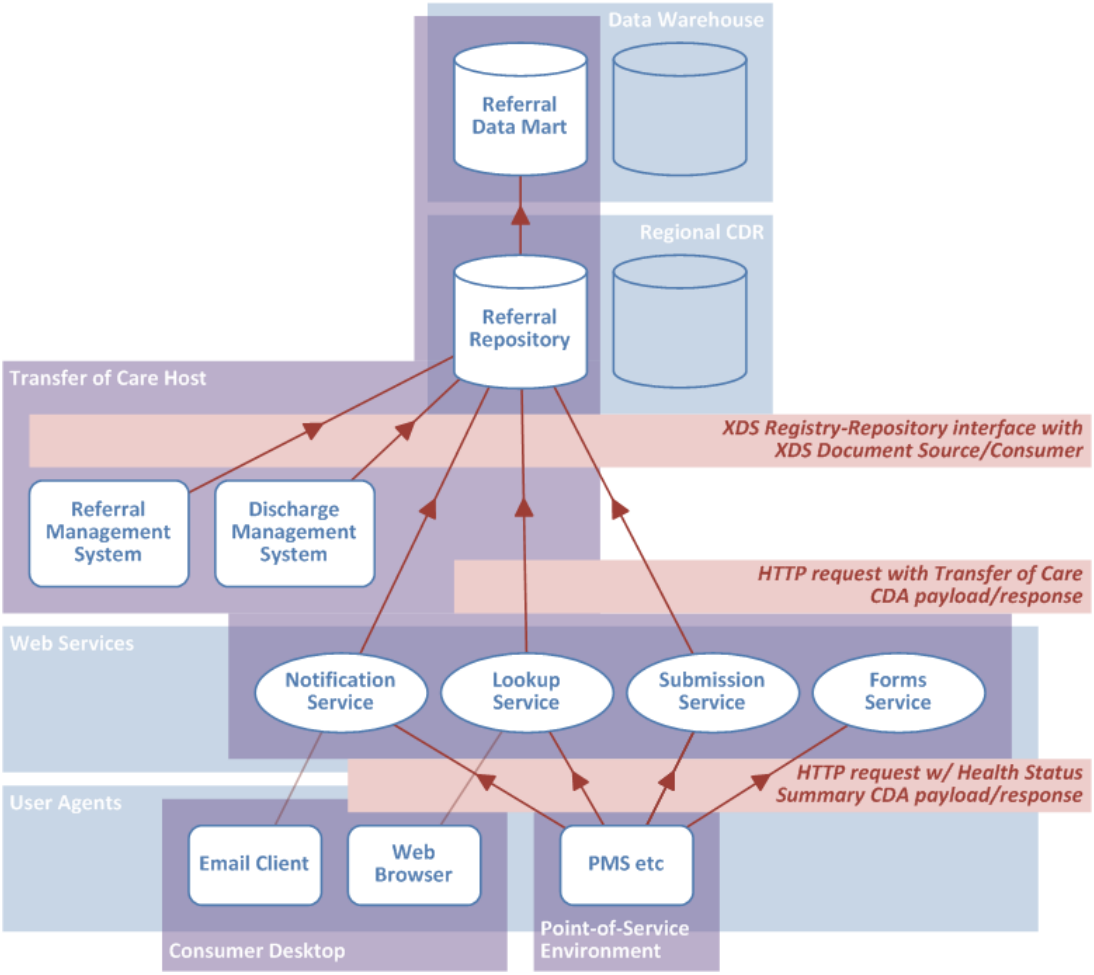

Figure 3 – Point-of-service access to transfer of care web services

Consumer access to transfer of care web services is also shown. Patients receive notifications via email of incoming referral or discharge documents, which can then be viewed using a lookup service via a secure portal.

3.2 Point-of-service systems stand in a one-way client-server relationship to transfer of care web services

In other words, we endorse an orthodox client-server architecture that clearly distinguishes client systems and servers, with all service calls being directed outward from the point-of-service environment and none back the other way. That is, the point-of-service system (or its proxy in the local environment) is always strictly an initiator of HTTP requests and is never in the position of having to respond to such requests. Furthermore, we can say that point-of-service systems should not themselves be positioned as shared repositories.

There are two reasons for these rules. The first is that the point-of-service systems we're talking about cannot be expected to have the level of performance, uptime and security expected of web service hosts. The second is about topology. Systems built on

peer-to-peer relationships, especially when there are many nodes, tend to be complex and hard to operate. In contrast, hub-and-spoke arrangements of many clients and few servers are naturally easy and predictable.

3.3 Transfer of care web services are published at addresses in the *health.nz* domain on the Connected Health data network

Referral and discharge web services that have passed certification can be issued addresses in the health.nz domain by the domain moderator (presently the Ministry of Health). This promotes user confidence in the fitness for purpose and security of the web services made available, which in turn depends on accreditation of the organisations developing and hosting those web services.

This rule also creates a useful checkpoint in operational governance over the set of web services in production.

4 Workflow

This section describes the required mechanism for workflow coordination in transfer of care. Referral and discharge are often simple processes but the advent of clinical pathways has created opportunities to optimise workflows. For example, in certain situations a radiology referral may be updated and forwarded directly to a specialist (without recourse to the referrer) on the basis of the diagnostic result. We need to able to accommodate such lateral shifts in information and responsibility.

4.1 Human workflow around transfer of care is based on the *IHE Cross Enterprise Document Workflow (XDW)* integration profile, controlled out of a companion document to the referral or discharge

XDW – based on the *OASIS WS-HumanTask* web services profile – supports distributed workflow management around such processes as referral and discharge. Scheduled and completed activity is tracked in a CDA document that is updated using XDS methods, with multiple participating systems playing their parts in distributed fashion. The XDW document is stored in the same XDS repository folder as the associated referral or discharge document.

The workflow document and the related referral or discharge document are mutually referencing.

We anticipate that XDW will be incorporated in the next edition of the XDS-MS profile, which will provide further clues as to how it should be implemented.

5 Identity

This section states directives for identification of participants to transfer of care processes, including consumers, practitioners and other members of the healthcare workforce. In particular, it is necessary to be able to address the various parties to a referral or discharge document and ascribe authorship.

5.1 Transfer of care web services use patient identity web services styled on the *IHE Patient Demographic Query HL7 v3 (PDQV3)* integration profile, with the *National Health Index (NHI)* as the patient identity source system

PDQV3 specifies a set of HL7 v3 RIM based web services for access to patient identity and demographic details. The NHI system will implement PDQV3 and will play the key role of patient identity source system in our XDS-compliant R-CDRs.

5.2 Transfer of care web services use provider identity web services based on the *IHE Health Provider Directory (HPD)* integration profile, with the *Health Provider Index (HPI)* as the provider identity source system

HPD has a parallel purpose to PDQV3, specifying HL7 v3 based web services for access to provider identity and demographic details. The HPI system will implement HPD.

5.3 Participants to transfer of care processes are each identified by URIs in the *health.nz* domain – for use as electronic addresses

The health.nz domain provides a suitable basis for a vendor-neutral addressing scheme that confers a measure of authenticity to communications.

6 Privacy and Security

This section states required privacy enforcement and security measures around information exchange in transfer of care processes.

6.1 Privacy preferences are recorded and enacted in conformance with the *IHE Basic Patient Privacy Consents (BPPC)* integration profile, with transfer of care documents in the repository subject to fine-grained access controls

Authorisation needs to be patient-practitioner specific in order to recognise shared care teams and the individuals who are entitled to receive particular transfer of care documents.

A consent model governs who can have access to transfer of care documents in the repository, subject to break glass and auditing provisions.

6.2 Referral management systems and web service platforms are secured in conformance with the *IHE Audit Trail and Node Authentication (ATNA)* integration profile

ATNA states requirements for security around systems and web services implementing XDS actors – document source, consumer, registry, repository and so on. The need here is to properly secure all the parts of what is a distributed system.

6.3 Authentication and access control over transfer of care web services conforms to the *SAML 2.0 Web Browser Single Sign-On Profile*

This profile separates the role of service provider (of the service the client is trying to access) from that of identity provider (which accepts and authenticates the client's credentials and issues a security assertion token that the client can forward to the service). SAML 2.0 is a recommendation of the *IHE Enterprise User Authentication (EUA)* integration profile.

This is a lightweight, browser-friendly and readily implementable method of achieving application layer security over access to web services. It avoids overreliance on the security of point-of-service applications.

6.4 Network access to connected care web services is via the Connected Health private data network, secured by use of the *Transport Layer Security (TLS)* protocol

Use of the Connected Health private data network (interconnected VPNs from participating telcos) is recommended for access at point-of-service to connected care web services. It is recommended that TCP/IP sessions have the further protection of TLS encryption to ensure confidentiality and data integrity.

Hosts are required to have X.509 digital certificates, issued by a recognised certificate authority.

It could be argued that, in order to safeguard authenticity, communications should carry digital signatures. However, this is not as yet a requirement.

7 Referral Management Systems

Referral management systems are the document tracking and workflow systems used to enable the smooth processing of referral requests and responses. They have functions supporting the referral grading process and administrative functions too. They interface to clinical portal, patient administration and scheduling systems. They capture information about the patient journey.

Directives in this section address the integration of referral management systems.

7.1 The regional information platform includes a referral management system as the coordination point for all referral traffic within the region

Having a single referral management system in any region should be sufficient, one that caters for all referrals sent within the region, whatever their kind or source and destination. This means having the ability to cover all situations, without regard to primary/secondary and public/private distinctions (subject to the extent that private providers are prepared to be involved).

With health services becoming progressively more interconnected, having a common referral management system within each region is one enabler of that.

7.2 Referral management systems are XDS-enabled and integrate with the document repository in the R-CDR

Figure 3 shows the close relationship of the referral management system to the regional document repository. In fact, the repository plays an integral part in this by serving as the primary data store for incoming referral documents – that is, the referral management system doesn't need to duplicate these documents in a database of its own.

8 Referral Data Entry

This section presents a recommended architecture for referral data entry – one based on the use of web forms. Characteristically, web forms exist independently of calling applications in the client environment, without any reliance on shared data structures or functions. Instead, all interaction is via published service interfaces.

Of course, there are perfectly legitimate alternatives to this modular style – take the obvious example of an application with inbuilt data entry forms and functions. However, the purview of directives in this section is the web forms style only.

Our architectural objectives here are solutions that demonstrate:

- *User interface conformity* – enabling forms to be deployed with the same look and feel into different settings
- *Systems interface standardisation* – for plug-and-play integration of forms services and their clients
- *Flexibility in software development and hosting arrangements* – promoting best of breed solutions and as a counter to vendor lock-in
- *Adaptability to changing requirements* – for fast and cost-effective delivery of functional enhancements
- *Reusability and configurability* – for reuse from one health service to another, permitting minor changes in livery or addressing, for example

8.1 Referral forms are HTML web forms, produced on demand by transfer of care web services and pre-populated from a CDA document containing a health status summary dataset

The crux of this directive is the prescribed method of pre-populating web forms. It makes sense to be able to pre-populate a form with whatever information is already held about the patient, saving the need to look it up and re-enter it. While the referrer's own system is an obvious source, we cannot expect such systems to be able to expose a suitably robust interface to external web services.

Which leads us to the strictly client-server solution recommended here. The calling application queries its own database to create a health status summary document, posting this to the selected forms service. (This service might very well draw on repository-held patient information too).

In order to behave in this way, applications must have an HTTP client component and a certain level of CDA- and HTML-processing capability. This can be inbuilt functionality of the base product itself or locally installed adapters can provide it.

Figure 4 illustrates submission of a GP-created referral, with the support of a typical practice management system (PMS).

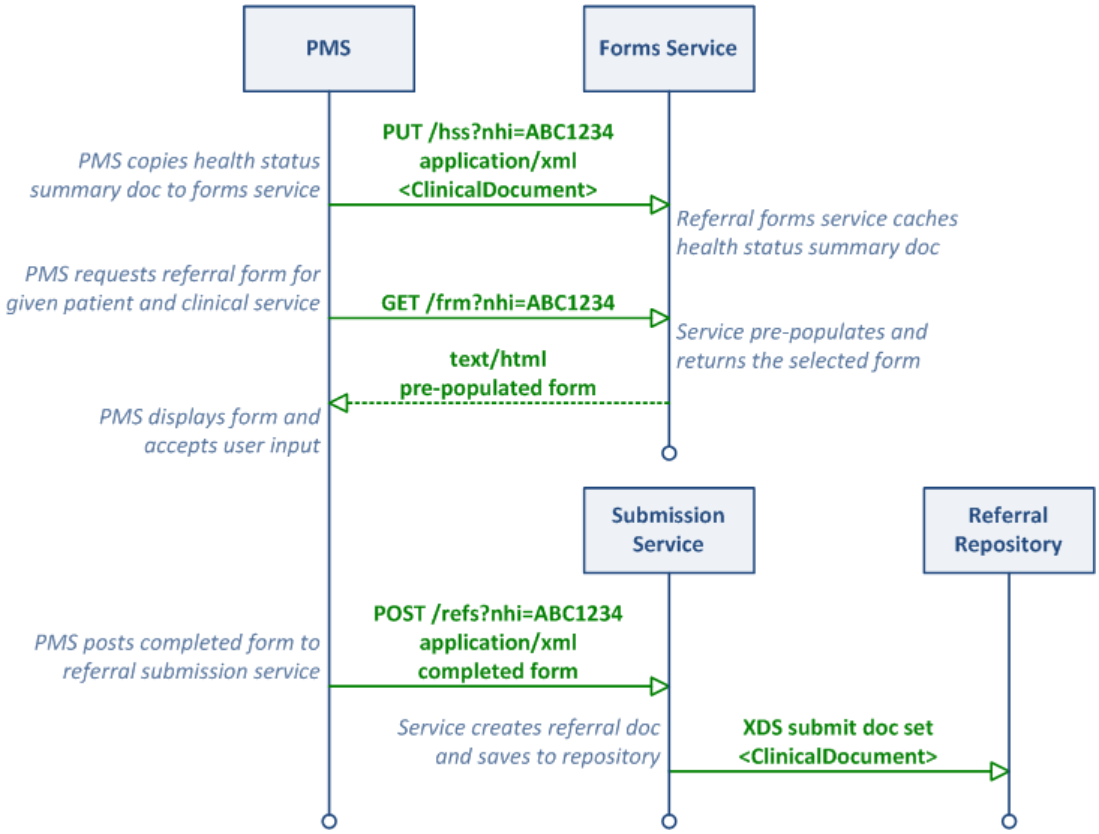

Figure 4 – Example of referral submission

The key interactions are (expanding a little on the diagram):

- The PMS invokes a chosen referral forms service and posts a health status summary document for the given patient
- The forms service pre-populates and serves the requested form (invoking other services as necessary)
- The user completes the form (invoking other services in order to reference any test results etc) and posts it to a submission service
- The submission service creates a referral document and sends it and any attachments as an XDS submission set to the local R-CDR document repository
- Recipients are notified of arrival of the document

The principle of incremental interoperability applies here, meaning that the forms service should be able to work with all health status summaries, even when the calling application doesn't have the capability (or authorisation from the user) to produce a fully complete document. Taking this to the extreme, services should still be able to function in the absence of any pre-populating data whatsoever, such as when the user's only tool is a web browser – a significant use case.

Figure 5 illustrates the notification process – in this case, a GP being notified as one of the recipients of a discharge summary.

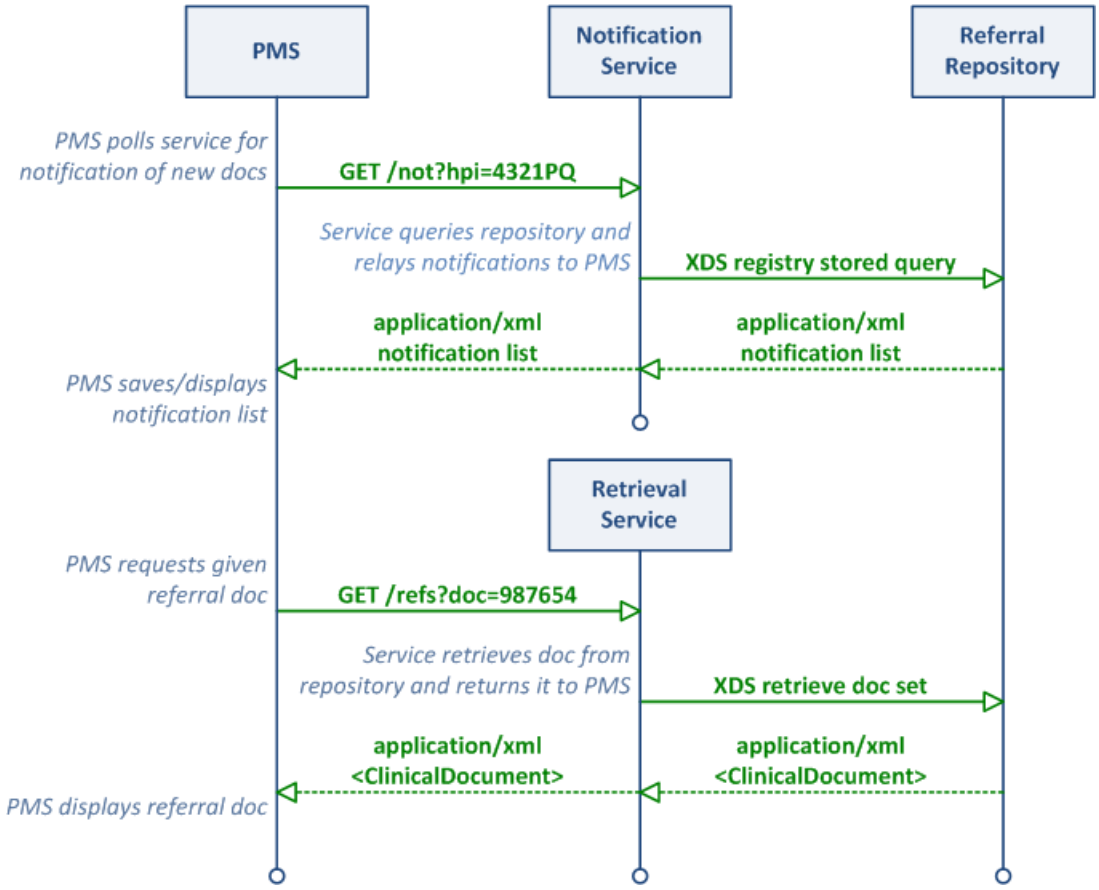

Figure 5 – Example of referral notification

We expect that forms services will offer sundry functions such as parking a form.

Form designs and program code are open source and are collected – at least conceptually – in a national forms library.

Although the health status summary is an extremely versatile dataset that will support most referral types, there will be other situations where the data requirements are somewhat different – referrals outside the health sector, for example. There is no reason that such requirements couldn't be treated in essentially the same way, if a small number of additional document types were defined.

8.2 Referral (or discharge) data is coded into CDA documents in such a way that recipients are able to display completed forms in the same format as created

The aim here is to honour the forms paradigm by preserving the look of completed forms from one end of the referral process to the other – mimicking paper forms in this respect.

Achieving this involves taking special care over how we populate the narrative blocks (one per section) that contain the CDA document's displayable elements. The simple markup involved can be easily rendered with an XSLT stylesheet into HTML lists, tables and formatted text. The requirement for conformity at either end can then be met by using the same stylesheets. Having PDF equivalents to CDA documents is not recommended.

Figure 6 illustrates the recommended information cycle. The point-of-service application automatically creates a health status summary document, which is used to pre-populate a data entry form. The form accepts user input, is posted, and populates a transfer of care document, which is eventually transformed for display.

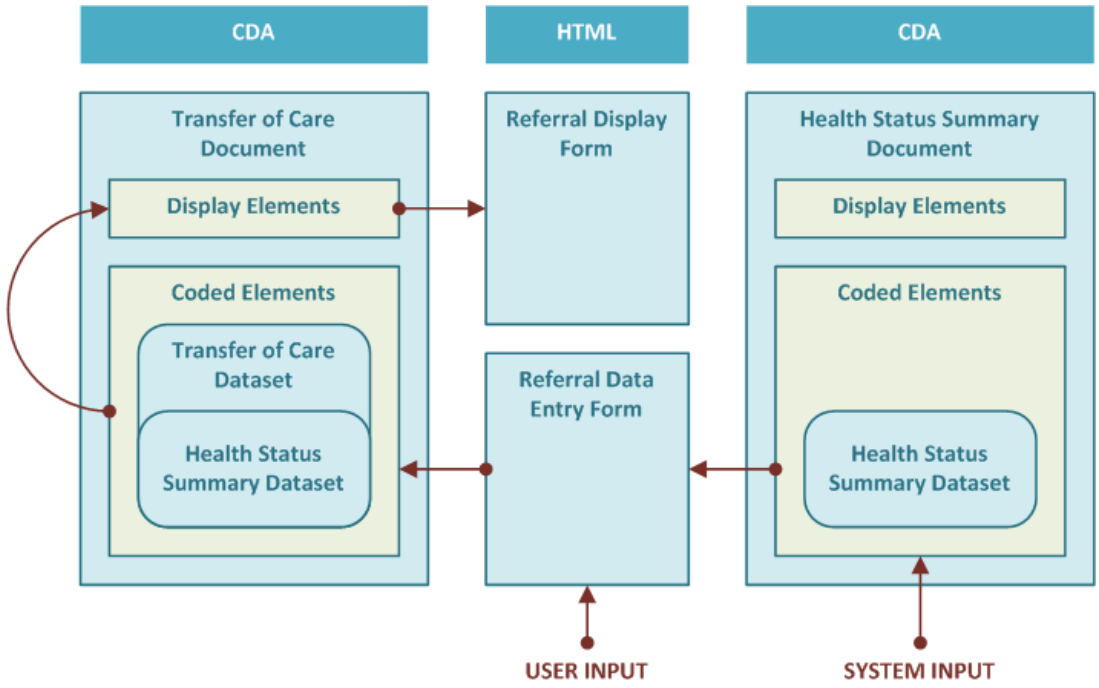

Figure 6 – Information cycle between documents and forms

Note that all display elements (narrative block entries) in the two CDA documents are derived from coded elements.

8.3 Forms are designed for cross browser support and use across a range of device types

Today, a reasonable expectation is that applications should have support for current versions of Internet Explorer, Chrome, Firefox and Safari browsers. Beyond that, applications should increasingly be designed with tablets and smartphones in mind, as much as the traditional forms of desktop PCs and laptops.

9 Shared Care Records

Our ideas on interoperability in transfer of care extend to the broader requirements of shared care. This section establishes certain fundamentals for interoperability of shared care records and systems.

9.1 Shared care records are composite datasets mastered in R-CDRs, exposed via XDS-based web services as CDA documents to point-of-service systems

The requirement here is that shared care records be maintained independently of any one application, stored securely in R-CDRs as commonly accessible data resources. Their content may be distributed across a number of repositories, with XDS as the linking mechanism.

Shared care records in whole or in part (if components are individually addressable) are transacted as standardised CDA documents. There are standard CDA templates for the medications list, for example.

Updating a shared care record involves checking out that record (or component of it), making the change, and checking the record back in. XDS- and CDA-enabled web services or adapters mediate this process where point-of-service applications do not have the capability natively.

Figure 7 shows an example of a shared care record as a composite dataset stored in distributed fashion across several XDS-enabled repositories. The registry is the central location where a comprehensive index to content is kept; it responds to client queries by returning pointers to the requested repository entries.

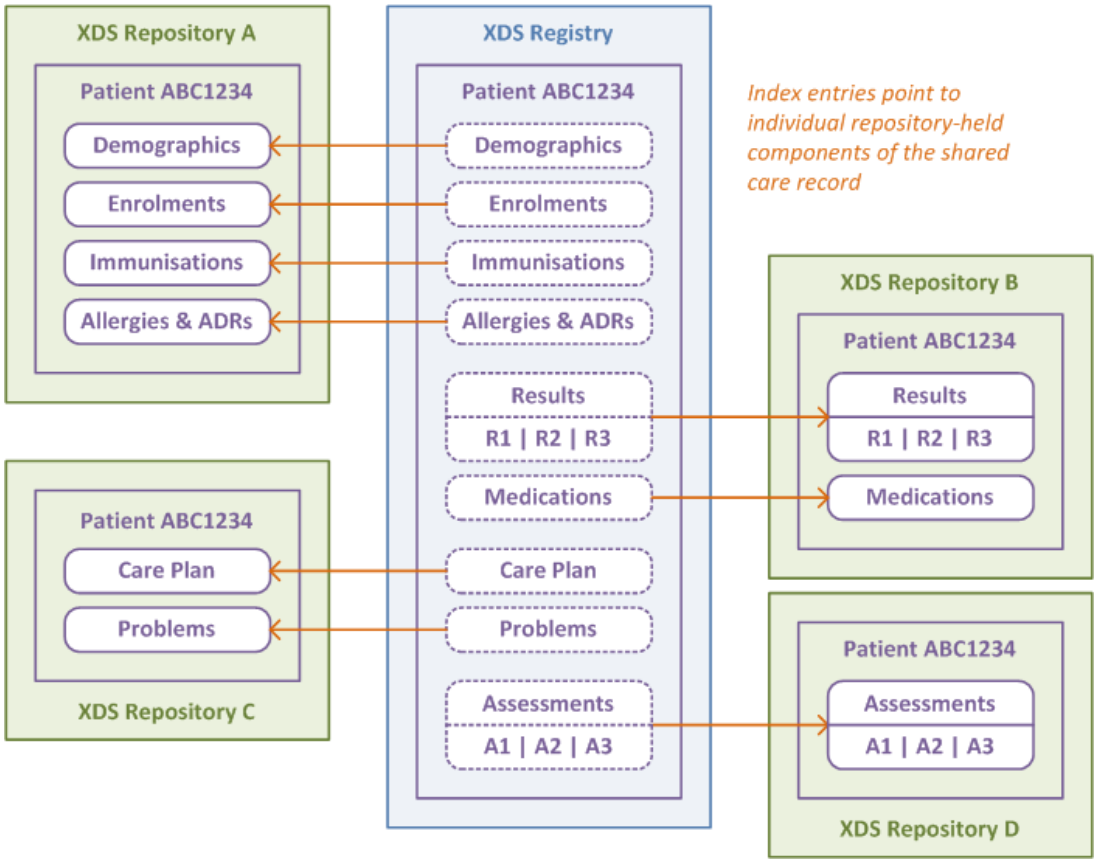

Figure 7 – Repository-held shared care records

As XDS permits, content may either be copied into a dedicated repository (e.g. a laboratory information system submits a test result) or it may be left at source and simply registered for subsequent retrieval on demand (e.g. sharing a PACS image).

The diagram shows the registry having pointers to entire sets of results or assessments; however, individual results etc can just as easily be indexed.

10 Business Intelligence

This section states requirements around the use of a business intelligence platform for the purposes of health service performance monitoring and population health data analysis in relation to transfer of care and patient flow. Business intelligence as a discipline employs the tools of statistical analysis, data mining and predictive modelling in order to derive knowledge from process and case related data.

10.1 The business intelligence platform for transfer of care is a data mart in the regional and/or national data warehouse

It is assumed that regions will each have their own data warehouses or, failing this, that there is a common national data warehouse at their disposal. In every region, the target data warehouse is populated with data extracted as a matter of course from all referral and discharge documents trafficked within that region, as they arrive in the R-CDR. These records become integrated in the data warehouse with booking and admission information from the same episodes of care, forming a comprehensive patient flow dataset. Data marts and cubes structure this information for analytical use.

Information governance rules are required to ensure that safeguards such as anonymization are in place.

10.2 Extracts from operational environments into the data warehouse conform to a standardised patient flow dataset

There is a standard minimum dataset for the capture of patient flow information, which naturally includes referral and discharge information. Referral and discharge solutions, including R-CDRS as their primary data collection points, are required to be able to capture and export a specified set of properly coded data elements, in an agreed format.

Appendix

In this appendix we look at certain important impacts of the reference architecture and considerations for its uptake.

The topics covered are:

- Standards Development
- Future State Migration
- Practice Management System Interfaces
- Transfer of Care Medications List

i Standards Development

The set of HISO standards relating to connected care is developing in line with the precepts of our reference architectures for interoperability and the information strategy of the National Health IT Plan.

The forecast changes are:

- The RSD standard – based on HL7 version 2 – will be replaced by a new CDA-based *Connected Care Information Standard*, with wider applicability
- This new standard will fully specify important new datasets and corresponding CDA document types, covering referral requests, discharge summaries, health status summaries and standalone medications changes summaries
- Connected care data models will be expressed as openEHR archetypes and ISO/IEC 11179 dataset specifications
- A supporting *CDA Common Templates Standard* will be published
- The 'lightweight forms' approach will be promoted as a CDA-based alternative to Online Forms (HISO 10014.2) for simple applications
- A common national XDS affinity domain policy will be developed in support of R-CDR implementations
- A new patient flow dataset will be specified in support of business intelligence

ii Future State Migration

This appendix outlines issues of migration from the current state in transfer of care solutions to the future state described by this reference architecture.

The key transitions required are:

- Repository-based information sharing replacing point-to-point messaging
- CDA documents replacing HL7 v2 messages as the currency of data exchange
- HTTP web services replacing the HL7 v2 transport protocol
- Repositories becoming CDA- and XDS-enabled
- Point-of-service systems becoming CDA-enabled

Organisations will have the opportunity to address these matters as new transfer of care and shared care solutions are developed regionally and nationally.

iii Practice Management System Interfaces

The reference architecture presents ideas about the implementation of transfer of care functions as web services and the ways in which point-of-service systems are interfaced to those services. This appendix looks at the practical implications and offers a template solution.

Figure 8 shows the PMS as a consumer of HIE web services in the recommended style of this reference architecture (left-hand side of diagram). The HTTP client invoked by the PMS is naturally pivotal here, both as an agent of the user's interactive session and used in periodic polling of notification services. The CDA toolkit (first developed as an adapter for GP2GP patient notes transfer) facilitates the information flows between the PMS and these agents. A protocol handler allows locally stored files to be picked up.

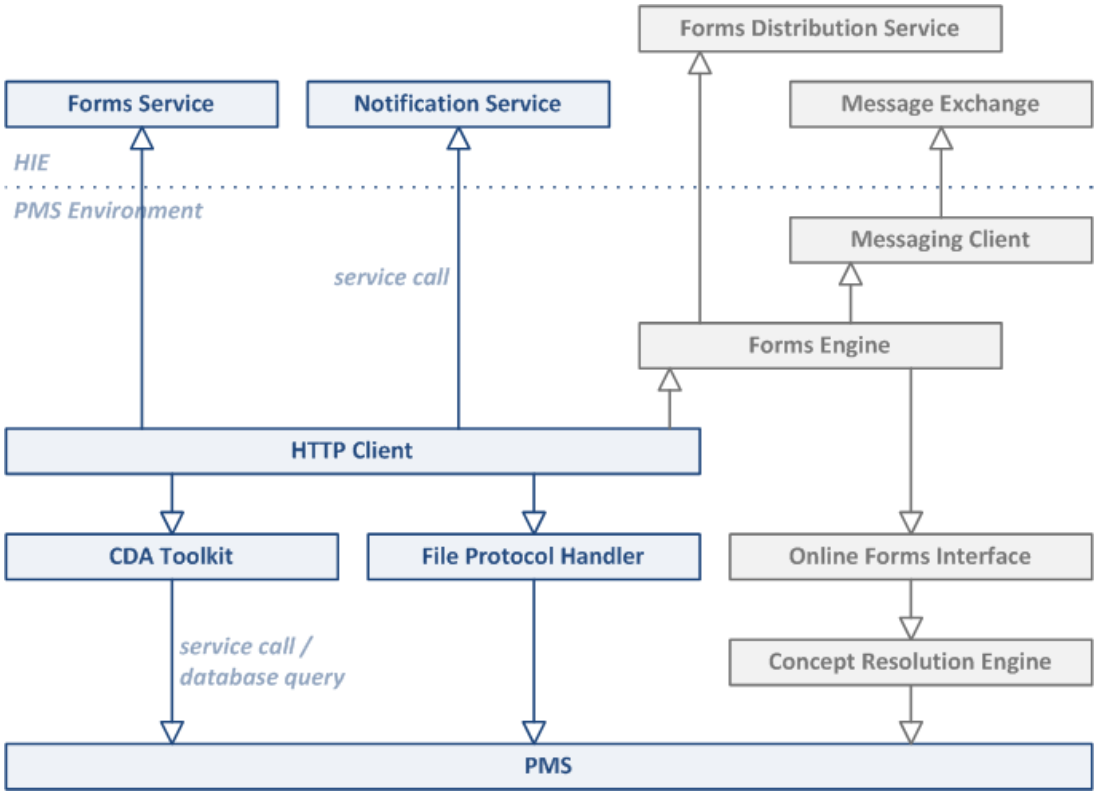

Figure 8 – PMS interface to web services

This pattern contrasts with that of some existing referral solutions built to the Online Forms standard and predating HISO 10040 (right-hand side of diagram). Local to the PMS environment are a forms (servlet) engine, and a messaging client. The forms engine downloads an HTML form template from a forms distribution service, pre-populates this form by calling a local web service fronting the PMS, and uses a local messaging client to send the completed form to a message exchange where the XML (but non-CDA) payload is packaged as an HL7 v2 RSD message for delivery.

Clearly, aspects of this approach have been overtaken by our new rules for interoperability. Under the more truly online and loosely coupled approach described by this reference architecture, forms are served externally out of the HIE rather than locally in the PMS environment. The PMS uses the toolkit to obtain local patient data as a CDA document before invoking a forms service with this payload. The user completes the pre-populated form returned by the chosen service, with the resulting dataset submitted via another service as a CDA document into a shared repository.

The above approach avoids the problems inherent when the PMS is directly exposed to – and even driven by – external software. Its other advantage is to avoid the complexity involved in having a forms engine in the local environment. Yet this solution is no less functional.

Very often the user will want to attach supporting documents to a referral request. The referral application then needs to allow the various documents available to be filtered and previewed, before final selections are linked to the request. This is easy enough when clinical documents are stored predominantly in the R-CDR; however, in the short term, we need the same functions over locally stored documents. Probably the cleanest way to provide these functions to a web form is by having the PMS implement a file protocol handler, which is our recommendation.

As a final comment on PMS interfaces, we acknowledge that there are alternatives to our very loosely coupled approach. The most interesting is for a style of client-side integration that has the HTTP client accessing PMS functions via a standardised JavaScript interface.

iv Transfer of Care Medications List

Transfer of care involves a shift in who is responsible for overseeing the patient's medications. Clinicians need to be able to review and update the medications list variously when they prescribe, dispense, perform medicines reconciliation, or see the patient at other times. Patients, too, will want some control over the medications list.

The concept of *My List of Medicines* is developing in answer to this need for an authoritative, widely accessible view of the patient's medications. In the following, we look at how this can be realised in terms of the interoperability reference architecture.

Figure 9 shows the information flows of the recommended solution. The medications list is mastered as a repository-held dataset. It is accompanied by a set of prescription and dispensing records, which are captured by a prescription broker – a web service called the *New Zealand ePrescription Service* – that interfaces with prescribers' practice management systems and with community pharmacy dispensing systems.

Typically, a managed medications list will be maintained for every patient enrolled in a shared care programme. List updates are ordinarily the responsibility of the prescriber, although the pharmacist and the nurse may also play a part. At any rate, these updates are never completed in simple automatic response to prescribing or dispensing transactions, but always as a deliberate action by a clinician. Accuracy of the list is paramount and we recognise that transactional records cannot tell the whole story.

All access to the medications list occurs via a second web service. This service has check out and check in controls of some kind (say, optimistic locking), permitting the list to be downloaded and edited with the user's own (certified) practice management or dispensing software. The modified list is then uploaded to the repository as an auditable new version. The payload of these transactions is a standardised CDA document – the *Transfer of Care Medications Document*.

The medications list service can also be interfaced to consumer portals. In the consumer use case, the transfer of care medications document is used to derive (by applying an XSLT stylesheet, for example) a displayable 'yellow card'.

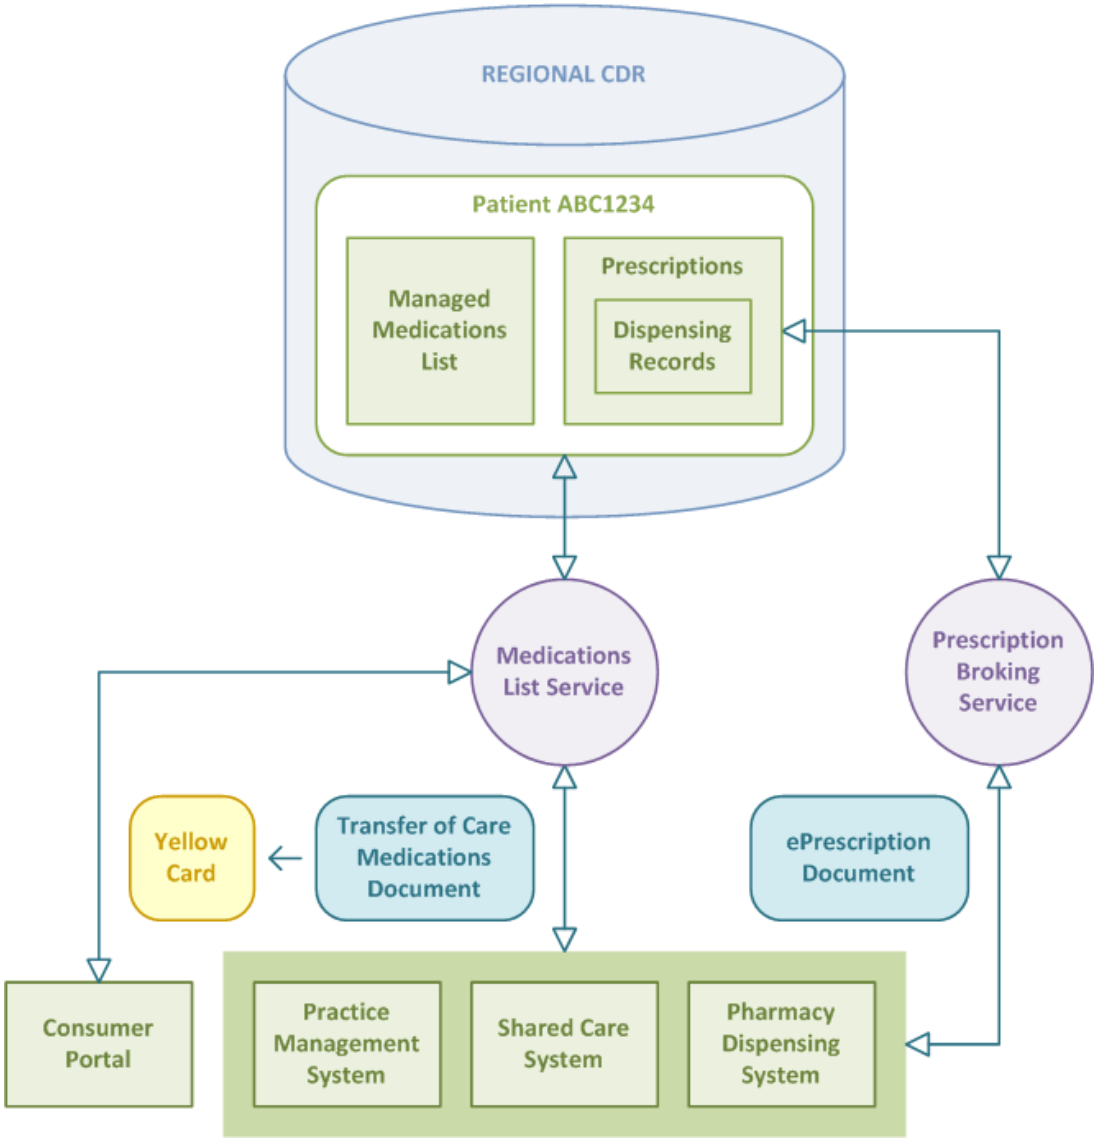

Figure 9 – Medications information flows