Interoperability

Reference Architecture

Version 1.0

December 2011

Contents

List of Figures

Contributors

The Interoperability Reference Architecture and related HISO 10040 architecture building blocks were created by the Interoperability Technical Working Group of the New Zealand Health and Disability Sector Architects Group.

The members of the team and individual contributors were as follows.

Contributor	**Organisation**
Dr David Hay	Health Alliance New Zealand
Dr Koray Atalag	Auckland University
Alan Le Maitre	Ministry of Health
Alastair Kenworthy	Ministry of Health

References

ASTM International	http://www.astm.org/
Australian Institute of Health and Welfare (AIHW)	http://www.aihw.gov.au/
Health Level 7 (HL7)	http://www.hl7.org/
HealthBase	http://www.infospace.health.nz/healthbase/
Integrating the Healthcare Enterprise (IHE)	http://www.ihe.net/
International Health Terminology SDO (IHTSDO)	http://www.ihtsdo.org/
International Standards Organisation (ISO)	http://www.iso.org/iso/iso_catalogue.htm
Organization for the Advancement of Structured Information Standards (OASIS)	http://www.oasis-open.org
Object Management Group (OMG)	http://www.omg.org/
openEHR	http://www.openehr.org/
World Wide Web Consortium (W3C)	http://www.w3.org

1 Document Overview

This document is the first edition of the Interoperability Reference Architecture for New Zealand. It describes a high level future state architecture intended to support the interoperability requirements of the National Health IT Plan (2010).

1.1 Background

The National Health IT Plan creates an overarching view of a patient focused, integrated healthcare model, enabling shared care between all providers involved in a patients care, including the patient themselves. To achieve this goal, a high level architecture was created, describing a Shared Care system, supported by a number of core repositories with free information flow between all sectors of the health care system, stating that:

- *To achieve high quality health care and improve patient safety, by 2014 New Zealanders will have a core set of personal health information available electronically to them and their treatment providers regardless of the setting as they access health services.*

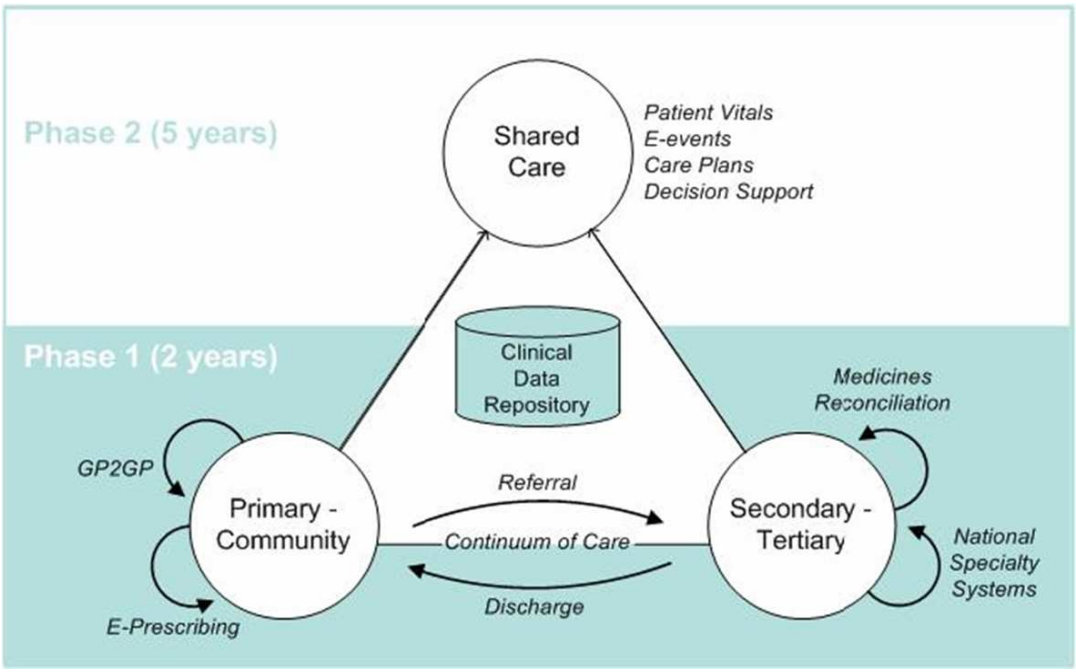

Figure 1 – National Health IT Plan (2010)

This diagram shows:

- The improvement of information flow between and within the main parts (sectors) of the health system, plus the establishment of regional clinical data repositories in the next two years
- The establishment within the next five years of shared care systems that link processes and data across the sectors to provide a seamless interface for health care delivery, and include the patient and family in that delivery

To achieve this goal requires the definition and creation of a number of components, standards to describe information movement between those components, and the agreement by vendors to implement those standards.

1.2 Document Purpose

This document presents the Reference Architecture for health information interoperability within New Zealand, to support the National Health IT Plan. It is a template solution and provides a common vocabulary with which to discuss implementations, often with the aim to stress commonality.

It is therefore intended to be a document that is referred to by solutions architects who are creating solutions to particular business needs within the health interoperability domain.

It is a **future state** document. The architectures and patterns discussed here describe how interoperability can be achieved into the future – not necessarily how it is being done now. It describes a realistic goal for interoperability rather than simply extending the techniques being used today. It is recognized that there will be a transition period between current and future states. It is also accepted that the demands of in-flight projects sometimes require that compromises be made, particularly when the desired infrastructure is not in place. However, it is hoped that all projects will, at the least, reference this document when creating a solution, and put in place components that will move towards the goals described here, or re-use/extend existing components.

It is an **evolving** document. The authors have, as much as possible, looked at current international best practice when making recommendations, however health interoperability is a moving target as implementers gain experience in what works – and what doesn't – when connecting systems together. It is expected and anticipated that the document will be reviewed as thinking advances in this discipline.

It is a **consensus** document. While it is not always possible to get complete agreement between all those involved in health interoperability (or any other human endeavour for that matter) input has been sought from as many participants as possible, and reflected in the contents. However, there is a need to provide leadership and direction, so at times the authors have needed to make decisions that some will disagree with.

Above all it is a **pragmatic** document. New Zealand is not a large country and resources are constrained and so the steps to achieve the vision need to be made in a stepwise fashion.

It is also important to be aware of other work being done internationally, in particular the Australian Personally Controlled Electronic Health Record (PCEHR) project.

The Reference Architecture will be refreshed at regular intervals to ensure it remains relevant.

The Reference Architecture is future state oriented and does not attempt to describe the current state or how to transition from current state to future state. Individual solutions are expected to address the legacy/transition issues that are found to exist.

1.3 Document Scope

The Reference Architecture addresses requirements related to:

- Exchange Content Model
- Health Information Exchange
- Registry-Repository Model
- Regional CDRs
- Model Driven Architecture

1.4 Governance of the Reference Architecture

The Sector Architects Group will take the Reference Architecture through the necessary processes, including reviews from the sector architects group and vendor architects group, satisfying the archi-

tecture governance group requirements and public consultation, to become a HISO approved standard. Where the Reference Architecture has key dependencies on other specifications, these will be made standard in their own right.

The diagram shows the governance model for all reference architectures, including this one.

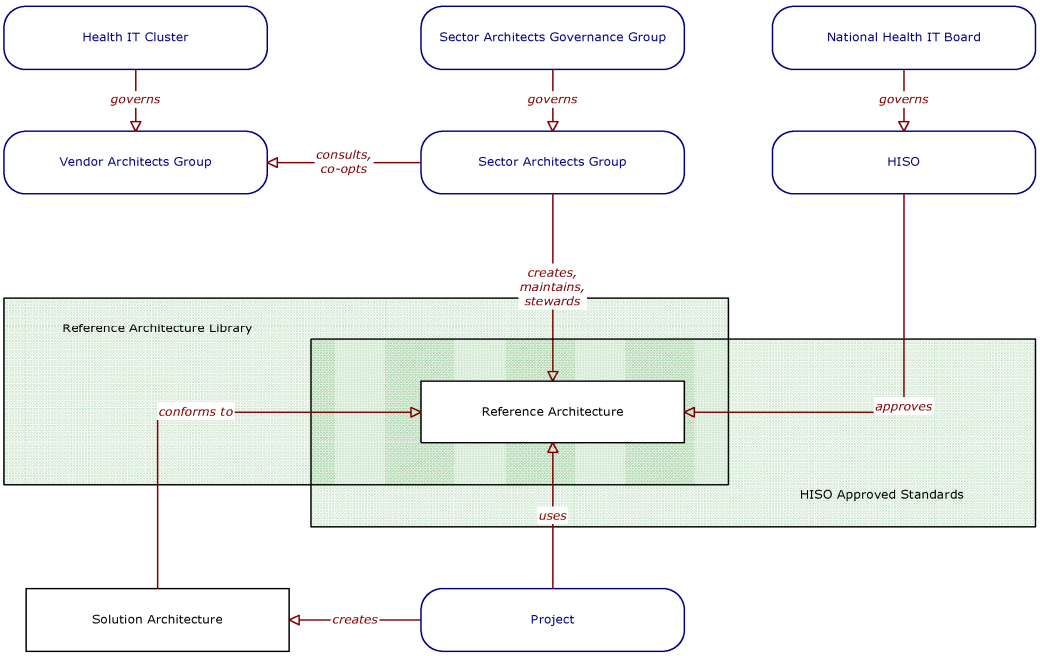

Figure 2 – Governance model for reference architectures

1.5 Document Structure

This document presents the Reference Architecture and is structured as follows.

Chapter 1 (this chapter) is a document overview.

Chapter 2 states the set of principles chosen to underpin the Reference Architecture and its development.

Chapter 3 technical overview of the Reference Architecture in terms of what it sets out to achieve and its essential elements.

Chapter 4 describes the general architecture directives that are to be used across all areas of the reference architecture, including all architecture building blocks Reference Architecture General Directives.

Chapter 5 describes functional model mapping, the mapping shows how the reference architecture will define the requirements for functional interoperability of HIEs and regional CDRs, in line with the EHR-S functional model.

Chapter 6 contains the three Bronze level architecture building blocks that are foundational for interoperability:

1. *HIE CDR Utility Services*, specifying a style of information exchange based on the registry-repository model of the *IHE Cross Enterprise Document Sharing (XDS)* integration profile
2. *HIE Structured Documents*, specifying *HL7 Clinical Document Architecture (CDA)* structured documents as the common currency of information exchange

3. *HIE Content Model*, specifying a common shared content model, with the *ASTM Continuity of Care Record (CCR)* as the basis for core health information, extensible per clinical specialty

Chapter 7 describes the high level scope of the Silver level architecture building blocks.

Chapter 8 describes the high level scope of the Gold level architecture building blocks.

Appendix A is a glossary of terms used and defined by this document.

Appendix B defines the future state architecture of systems and services at all levels in the sector.

Appendix C discusses the service approach of the reference architecture.

Appendix D describes the concept of the health information exchange as the standards based fabric that connects systems across the sector.

Appendix E discusses terminologies and their application.

Appendix F discusses system behaviour and the protocols that enable data exchange and interoperability.

Appendix G discusses the openEHR detailed clinical models.

Appendix H discusses security and privacy requirements for data exchange.

2 Principles

The following high-level principles underpin the Reference Architecture and have guided its development.

1. **Align to national strategy.** The Reference Architecture will align with national standards and business strategies, with priorities defined by national and regional IT plans.
2. **Invest in Information.** We will represent health data for exchange as detailed clinical models that can be represented in different ways independently of any particular information model or serialized representation (structure) and derived directly from business requirements with clinical input. These models may be represented in different ways for different audiences.
3. **Use single content model.** Information for exchange will be defined and represented in a single consistent way at the information model level. Where possible, it will align with national and international standards.
4. **Work with sector.** The development of the Reference Architecture will be in partnership with the sector as represented core groups including: Clinical leadership group, Sector Architecture Group, HISO, National Health IT Board, vendors, PHO, consumer groups and other affected agencies
5. **Align to business needs.** Development of the details of the Reference Architecture will be in conjunction with the prioritized business projects. Prioritization will be set by IT plans embodying those needs. The intent is to ensure clinical and other business engagement.
6. **Use proven standards.** Where there is a relevant national or international standard that is compliant with the overall direction of the Reference Architecture, will meet a particular business/technology requirement and is widely used, we will use that standard. If modifications are required, we will work with the relevant SDO to make the modifications. This approach applies at all levels of the interoperability stack including workflow, payload, security, terminology and transport.
7. **Adopt services approach.** To define the behavioural aspects of interoperability we will use a services approach, where a service can be thought of as a method of encapsulating business functionality behind a clearly defined interface that is technology agnostic and conforms to accepted practices.

3 Technical Overview

This chapter introduces the Reference Architecture in terms of the problem it addresses and the essence of its solution to that problem.

3.1 Problem Statement

The National Health IT Plan sets out to realize an e-health vision of collected personal health information, shared between individuals and their care teams and across settings. Clinical decision support, care plans and other shared care functions will be built on a foundation of orders and results, medication lists and other objective information stored in regional clinical data repositories.

The data flows necessary to underpin this are complex and the situation demands new levels of interoperability. The purpose of the Reference Architecture is to address this problem and to define a new benchmark for systems to be classed as interoperable.

3.2 Interoperability

To say that two systems interoperate means they exhibit coordinated run-time behaviour centred on data exchanges based on an agreed transport protocol, payload specification, information model and workflow model. More generally, systems can be called interoperable when by design they can interoperate with entire classes of system (A system might be a consumer system (such as a GP PMS) or it might be a repository system).

The Reference Architecture distinguishes three levels of interoperability: functional interoperability, semantic interoperability and process interoperability.

- **Functional interoperability** exists where there is a physical method of moving the data between systems. Exchanging a PDF file via email is an example.
- **Semantic interoperability** exists when the recipient system is able to understand the information it has received, and act upon it. For example, a medication list in an electronic discharge summary can be processed to update the patient record in the recipient system.
- **Process interoperability** occurs when there is a process that can occur across systems. For example the discharge summary indicates that a follow-up check of Prothrombin time needs to occur as the patient was started on Warfarin, the recipient system can update its internal structure to carry out the test, and the sending system is able to see the result of that test (or know that it did not occur).

3.3 Key Components

The essence of the solution represented by the Reference Architecture can be defined in the above terms:

- Functional interoperability centres on the IHE Cross Enterprise Document Sharing (XDS) profile, with HL7 Clinical Document Architecture (CDA) payloads and W3C web services transport
- Semantic interoperability centres on an Exchange Content Model derived from the ASTM Continuity of Care Record specification, and use of terminologies such as SNOMED CT and LOINC, with detailed clinical models expressed as ISO 13606 archetypes that extend the Model into specialty areas
- Process interoperability centres on use of Integrating the Healthcare Enterprise (IHE) integration profiles, where those profiles exist and applicable to our clinical domains.

Collectively, the above define a standards-based data services fabric, termed a **Health Information Exchange (HIE)**.

The overall objective of the Reference Architecture is to provide **semantic interoperability** and **process interoperability**, where distributed transactions involving different systems and actors can occur, preserving meaning across the exchange. This will support the business objective of a **shared care** approach to delivering healthcare to an individual; in a way that includes the patient in the care team.

Underpinning the architecture is a **services based approach**, where a service can defined as specific functionality that be invoked using defined interfaces that are implementation agnostic – such as web services. These services can be categorized in a **services taxonomy**. **HSSP** (Healthcare Services Specification Project) is a valuable resource for thinking in services. The **EHR System Functional Model** (ISO/HL7 10781:2009) also helps to categorise services required in interoperability scenarios.

The concept of the **Health information Exchange (HIE)** as a capability or a collection of standard data services with their underpinning concepts of security and privacy provides the logical interface by which applications communicate with each other. An HIE is not a single, nationwide integration engine or ESB, although these systems can implement an HIE.

Integrating the Healthcare Enterprise (IHE) is an organisation with valuable resources representing many person-hours of thinking about clinical workflow and the actors and transactions required to support that workflow.

The **Registry-Repository Model** (e.g. IHE XDS) will provide a solution to the problem of locating information quickly at a regional and national level.

Regional Clinical Data Repositories will follow a registry-repository model; this will support a federated approach, allowing that national systems can be frontline repositories. This approach improves data quality by preserving the authoritative data source. The use of the XDS.b registry will ensure fast response times of patient information and provide granular security of the information.

An **Exchange Content Model** to which all data exchange refers, and is extensible to specialist domains will lead to semantic interoperability. **Archetypes** are used to define this model, which can be represented in other formats such as UML. This may lead to the possibility of **Model Driven Development** of artefacts. **Standardised Terminology and Reference sets** complete the picture.

HL7 **CDA documents** form the common coin of exchange. Mapping to the Exchange Content Model, CDA documents carry the clinical information being transferred, with the services around them providing the workflow.

The **HL7 Services Aware Interoperability Framework (SAIF**) can be used to organize the analysis and design process and its outputs.

Tooling and metadata repositories will be important, whether for designing artefacts, providing a searchable store of existing artefacts and/or concepts or as run-time adapters such as the GP2GP toolkit.

Finally, **effective governance** over all aspects and artefacts involved in interoperability, as well as **clinician and user engagement** will be critical to ensure that we do not have a multitude of services and artefacts with overlapping and duplicate responsibilities.

3.4 Interoperability Standards Model

The table below shows the layers and standards in the standards model applicable to the Reference Architecture.

Terminology	**SNOMED CT*, LOINC*, ICD-10***
Exchange Content Model (ECM)	**CCR** - Logical Framework based on ASTM Continuity of Care Record **ISO 11179*** - Used to express DCM's (Authoritative Source) **ISO 13606/openEHR** - Used to express DCM's Emerging Standard: ISO 13972 - Used for methodologies for DCM's
Workflow	**WS-HumanTask, BPEL, BPMN, UML 2.0** (IHE XDW)
Record Locator	**ebXML** (IHE XDS.b and supporting profiles)
Forms	**HTML5 using AJAX** - Used for Online Forms
Document Exchange	**HL7 CDA R2, MTOM/XOP** (Containment Standards: HL7 V2.5*, RSD HL7 V2.3, RSD HL7 V2.4*, LAB HL7 V2.1) (Emerging Standards: HL7 CDA R3, HL7 V3)
Transport	**SOAP over HTTP, REST** and **DICOM** (IHE XDS-I) WebServices Only
Network	**TCP/IP** over Connected Health Network*

* HISO Standard

Figure 3 – Interoperability Standards Model

4 Reference Architecture General Directives

This section contains the general architecture directives that apply across all areas of the Reference Architecture, including the architecture building blocks.

4.1 Services

Directives under this heading cover:

- Service oriented architecture
- Service contract
- Service taxonomy
- Service specifications
- Unity of purpose
- Service composition
- Service conformity
- Consistent addressing scheme
- Minimal touch points

Directive	Service Oriented Architecture
Statement	The Reference Architecture requires that the SOA approach be applied to inter-operability and that a governance structure be created to ensure that services are created and published correctly, and used by all participants.
Rationale	SOA is industry best practice.

Directive	Service contracts
Statement	Each service has a standard published interface exposed as a service contract
Rationale	The service contract will describe the technical and operational requirements of the service.

Directive	Service taxonomy
Statement	Services added to the Service Taxonomy will comply with the Taxonomy Matrix
Rationale	All services must be shown in the matrix and approved to ensure compatibility with other services in the matrix

Directive	Service specifications
Statement	Service specifications shall be well defined and clearly scoped and with well understood requirements and responsibilities. Services will be specified sufficiently to address interoperability requirements at all levels – functional, semantic, and process interoperability.
Rationale	The set of standardized services must conform to common parameters. The set of standardized services need to provide for various levels of interoperability.

Directive	Service unit of purpose
Statement	Services shall have a unity of purpose, fulfilling specific, related requirements.
Rationale	Services need to be able to be used in different ways and combinations to produce different results; an example of this is when many services combined provide workflow to satisfy a use case.

Directive	Service composition
Statement	Services shall be composable, i.e. multiple services can be composed with one another to create new services.
Rationale	Services need to be able to be used in different ways and combinations to produce different results, an example of this is when many services combined provide workflow to satisfy a use case

Directive	Service conformity
Statement	It must be possible to replace a conformant service implementation with another one meeting the very same conformance profile while maintaining functionality of the system
Rationale	This is how plug-and-play can be achieved

Directive	Consistent addressing scheme
Statement	Consistent addressing scheme based on services is to be used
Rationale	The addressing schema of the services will be held in a registry and will need to be consistent to allow services to be located

Directive	Minimal touch points
Statement	The Consumer service will have minimal touch points to obtain information
Rationale	To reduce application interoperability complexity

4.2 Health Information Exchange

Directives under this heading cover:

- Certification of touch points
- Standard published interface
- HIE deployment methods
- Standardized services
- Workflow
- Security
- Audit functions

Directive	Certification of touch points
Statement	All HIE touch points need to be certified
Rationale	The HIE functions in native mode only, HIE reliability will be compromised if any application does not conform to the HIE native mode

Directive	Standard published interface
Statement	Each data service has a standard published interface
Rationale	The HIE functions in native mode only, HIE functionality will be compromised if any application does not conform to the HIE native mode

Directive	HIE deployment methods
Statement	There are four permitted deployment methods; the methods are (1) dedicated, (2) external, (3) internal and (4) virtualized. (These deployment methods are described in detail in Health Information Exchange appendix.)
Rationale	To ensure consistency and reliability of the HIE

Directive	HIE functions in native mode only
Statement	The HIE will only perform functions that are derived from the standardized services of the services taxonomy
Rationale	To ensure reliability of the HIE, only standardized services can be used

Directive	HIE provides workflow
Statement	The HIE can provide workflow as required across the set standardized services, within the HIE fabric
Rationale	Workflow is required to order the service business functions within the HIE

Directive	HIE provides security
Statement	All services must have security, this function is actioned within the HIE fabric
Rationale	Security is a mandatory requirement when using HIE services

Directive	HIE provides auditing
Statement	All services must have auditing, this function is actioned within the HIE fabric
Rationale	Auditing is a mandatory requirement when using HIE services

4.3 HIE Adapters

Directives under this heading cover:

- HIE external adapters
- Native connections
- Adapter deployment methods
- Adapter certification

Directive	Adapters must be external to HIE
Statement	The HIE fabric must not be compromised, by internal adapters
Rationale	To maintain reliability and performance the HIE does not host adapters

Directive	Adapters provide native connection to HIE
Statement	All applications that interface with the HIE must do so natively, if the application cannot do this then an adapter is required.
Rationale	The purpose of adapters is to allow applications to interface to the HIE that cannot interface natively

Directive	HIE adapter deployment methods
Statement	There are three permitted deployment methods; the methods are 1. Consumer Application, 2. Provider Application and 3. Middleware (These deployment methods are described in detail in Health Information Exchange Appendix)
Rationale	To ensure consistency and reliability of the HIE

Directive	Certified adapters
Statement	Adapters must be certified
Rationale	To maintain reliability and performance of the HIE

4.4 Behaviour

Directives under this heading cover:

- Analysis and design methodology
- Functional model
- Behaviour modelling
- Workflow implementation
- Technical frameworks
- Localization

Directive	Analysis and design methodology
Statement	The HL7 Services Aware Interoperability Framework (SAIF) shall be used in analysis and design relating to interoperability.
Rationale	Although still a work in progress, the SAIF framework offers a way to align all the artefacts involved in interoperability, providing traceability throughout the entire stack back to the business objectives for interoperability.

Directive	Functional model
Statement	Define the required interoperability functions in terms of the ISO/HL7 10781:2009 Electronic Health Record System Functional Model (EHR-S FM) and Personal Health Record System Functional Model (PHR-S FM) where possible.
Rationale	These functional models are international standards providing rich sets of definitions and conformance criteria with local applicability. They are compatible with and complementary to the SAIF methodology. EHR-System FM defines the required set of shared care functions, for example.

Directive	Behaviour modelling
Statement	The required behaviour of any solution should be expressed in UML models, using the diagram types indicated. This includes use of the Object Constraint Language (OCL) for business rules specification.
Rationale	Formal specifications are essential to properly engineered solutions. UML is the most widely used modelling language worldwide, tools are freely available and skills exist locally.

Directive	Workflow implementation
Statement	Workflow will be implemented based on the IHE XDW profile, the profile is dependant on the XDS registry - repository model
Rationale	The XDW profile uses three important standards. The WS-HumanTask specification supports the ability to allow any application to create human tasks in a SOA environment. BPEL is an open specification with good vendor support internationally and has a tie-in to UML modelling. BPMN extends BPEL.

Directive	Technical frameworks
Statement	IHE technical frameworks and integration profiles should be used where available. When analysing the requirements in an interoperability use case, it is recommended that the IHE suite of profiles should be first examined if there are existing profiles that match the requirements. If there is a partial match, then the project should engage with IHE to further refine the profile for the benefit of all.
Rationale	IHE profiles have widespread use internationally, contain a significant amount of input from clinicians and technical people, and offer considerable reuse locally; they provide at the very least a starting point.

Directive	Localization
Statement	Local standards should be traceable to international standards and as far as possible faithful to them.
Rationale	This is the rule of adopt first, adapt only if you really have to. Local standards are their international counterparts should be more the same than different.

4.5 Security

Directives under this heading cover:

- Authentication
- Audit
- Credential management
- Role-Based Security
- Confidentiality
- Document integrity
- Non-repudiation
- Consent
- Secure communication channel

Directive	Authentication
Statement	HIE transactions and interactions shall be authenticated, including all access to R-CDRs, national systems and other participant systems. Authentication shall be at the level of individuals. Multi-factor authentication should be implemented where practical.
Rationale	Required for information security and governance compliance.

Directive	Audit
Statement	The IHE Audit Trail and Node Identification (ATNA) profile should be considered the basis for audit of information exchange.
Rationale	This profile has a reasonably comprehensive set of functions and the profile as a whole interrelates well with XDS and other IHE ITI profiles.

Directive	Credential management
Statement	There shall be national identity and access management standards, systems and processes.
Rationale	Required for Information Security and Governance compliance

Directive	Authorisation
Statement	Authorisation should be on the basis of a nationally defined set of roles, but implemented by the application/facility. This may require extensions locally.
Rationale	Required for Information Security and Governance compliance

Directive	Role-based security
Statement	There shall be role-based security
Rationale	Required for Information Security and Governance compliance

Directive	Confidentiality
Statement	Participant systems and exchange services shall ensure confidentiality in the exchange of health information.
Rationale	This protects against health information being disclosed, whether intentionally or unintentionally.

Directive	Document integrity
Statement	Participant systems and exchange services shall ensure document integrity in the exchange of health information.
Rationale	This validates to all parties that information has not been tampered with.

Directive	Non repudiation
Statement	Participant systems and exchange services shall support non-repudiation of source in the exchange of health information.
Rationale	This ensures that the true source and integrity of any piece of health informa-

	tion is always verifiable by all parties.

Directive	Consent directives
Statement	Participant systems and exchange services shall have support for consent directives in the exchange of health information. This includes the ability to record consent directives, attach them to subject health information as necessary, and at all times respect and enforce them. At a minimum this will include mechanisms for adherence to the information sharing and privacy rules of the Health Act, Privacy Act, Health Information Privacy Code and other applicable statutes. It will also address consent directives made by the individual, in terms of a defined consent model.
Rationale	To ensure that health information is collected, accessed, used and disclosed only with the individual's consent, and that these protections extend to health information exchange.

Directive	Secure communication channel
Statement	HIEs shall always use Connected Health accredited private networks for exchange of health information.
Rationale	Connected Health is the approved standard for interconnected private health networks. HIEs will not be exposed to the public internet.

4.6 Point-to-Point Messaging

Directives under this heading cover:

- Messaging point-to-point

Directive	Document messaging point-to-point
Statement	Document messaging point-to-point shall adhere to XDR/XDM integration profiles.
Rationale	The XDR/XDM profiles belong to the same family as XDS and can be used to put some rigour around point-to-point messaging for applications such as GP2GP.

5 Functional Model Mapping

The table below shows the ISO/HL7 10781:2009 EHR System (EHR-S) Functional Model and the mapping to the interoperability architecture building blocks. The mapping shows how the Reference Architecture will define the requirements for functional interoperability of HIEs and R-CDRs, in line with the EHR-S functional model. The ABBs are divided into three phases: Gold, Silver and Bronze. In this document, the Bronze phase ABBs have been completed; the Silver and Gold phases are only scoped.

EHR-S Section	EHR-S Subsections		Interoperability Architecture Building Blocks	
Direct Care	DC.1	Care Management	Silver.3 Gold.1	Shared Diagnostics Ordering and Reporting Task Services, Shared Care Records
	DC.2	Clinical Decision Support	Gold.1	Shared Care Records
	DC.3	Operations Management and Communications		Out of Scope
Supportive	S.1	Clinical Support	Bronze.1 Silver.4	HIE CDR Utility Services, Health Provider Entity Services
	S.2	Measurement, Analysis, Research and Reports	Gold.2	Population Based Services
	S.3	Administrative and Financial		Out Of Scope
Information Infrastructure	IN.1	Security	Silver.5	HIE Security Utility Services
	IN.2	Health Record Information and Management	Bronze.1 Bronze.2	HIE CDR Utility Services, HIE Content Model, HIE Structured Documents

			Bronze.3	
	IN.3	Registry and Directory Services	Bronze.1 Silver.4	HIE CDR Utility Services, Health Provider Entity Services
	IN.4	Standard Terminologies and Terminology Services	Bronze.2 Silver.2	HIE Content Model, Terminology Utility Service
	IN.5	Standards-Based Interoperability	Bronze.1 Bronze.2 Bronze.3 Silver.1	HIE CDR Utility Services, HIE Content Model, HIE Structured Documents, Point to Point Messaging Utility Services
	IN.6	Business Rules Management	Gold.3	Business Rules and Workflow Management
	IN.7	Workflow Management	Gold.3	Business Rules and Workflow Management

6 Bronze Level Architecture Building Blocks

There are three key building blocks that we are proposing as being essential for interoperability.

1. *HIE CDR Utility Services*, specifying a style of information exchange based on the registry-repository model of the *IHE Cross Enterprise Document Sharing (XDS)* integration profile
2. *HIE Structured Documents*, specifying *HL7 Clinical Document Architecture (CDA)* structured documents as the common currency of information exchange
3. HIE Content Model, specifying a common shared content model, with the *ASTM Continuity of Care Record (CCR)* as the basis for core health information, extensible per clinical specialty

To use an everyday analogy, if we consider the first of these to define the postal system, then the second defines the envelope, and the third the contents of the envelope.

6.1 Specifications

Each building block is formulated as a separate document. They combine to make a coherent whole, but are individually more or less discrete and independent.

They are technical specifications. They comprise architectural principles and requirements, and reference standards. They are reasonably terse, going into detail only where necessary to avoid ambiguity.

They are future state documents. They define a future state that can be achieved using the standards and technology of today, but without undue deference to today's systems and their limitations.

They are pragmatic documents. The future state they describe is readily achievable. They are designed for uptake and not to place unreasonable demands on software vendors and implementers.

They are evolutionary specifications. Interoperability is a moving target as the industry gains experience. The specifications can be expected to change in their detail, although not their essence.

6.2 Stakeholders

There are a number of stakeholders to this work.

Regional Information Systems Groups are key stakeholders. In turn, so too are there customers, the health workers and organisations around the country.

Software vendors are key stakeholders. They are being asked to support the new standards in their products.

Some in-flight and upcoming projects will be affected. This includes R-CDR projects, the Auckland Region eReferral Project, the Health Identity Project, and the National View of Cancer Project.

The building blocks do not directly affect consumers. Any connection is via the National Health IT Plan and Regional Information Systems Plans, for which the building blocks provide some of the technical substance.

Information governance work is underway as a separate although related exercise to define how information should be used, collected and so forth, in business terms.

Solution architects are the primary audience for the specifications. The building blocks are theirs to use in creating new solutions.

Australia's Personally Controlled Electronic Health Record (PCEHR) project is well aligned to our work. We are watching several others, too.

6.3 Related HISO Standards

Certain HISO standards are impacted by this work and will require review or at least the overall fit needs to be worked out.

The building blocks note in some detail where overlaps and conflicts occur.

In summary, there is some impact to the following in particular:

- HISO 10014.1 Data Concept Repository Processes
- HISO 10014.2 Online Forms
- HISO 12345 Referrals, Status and Discharge
- HISO 12345 Cancer Standard
- HISO 12345 Palliative Care Standard
- HISO Connected Health Connectivity Standards
- **HISO 10040.1 Health Information Exchange Clinical Data Repository Utility Services**
- **HISO 10040.2 Health Information Exchange Content Model**
- **HISO 10040.3 Health Information Exchange Structured Documents**

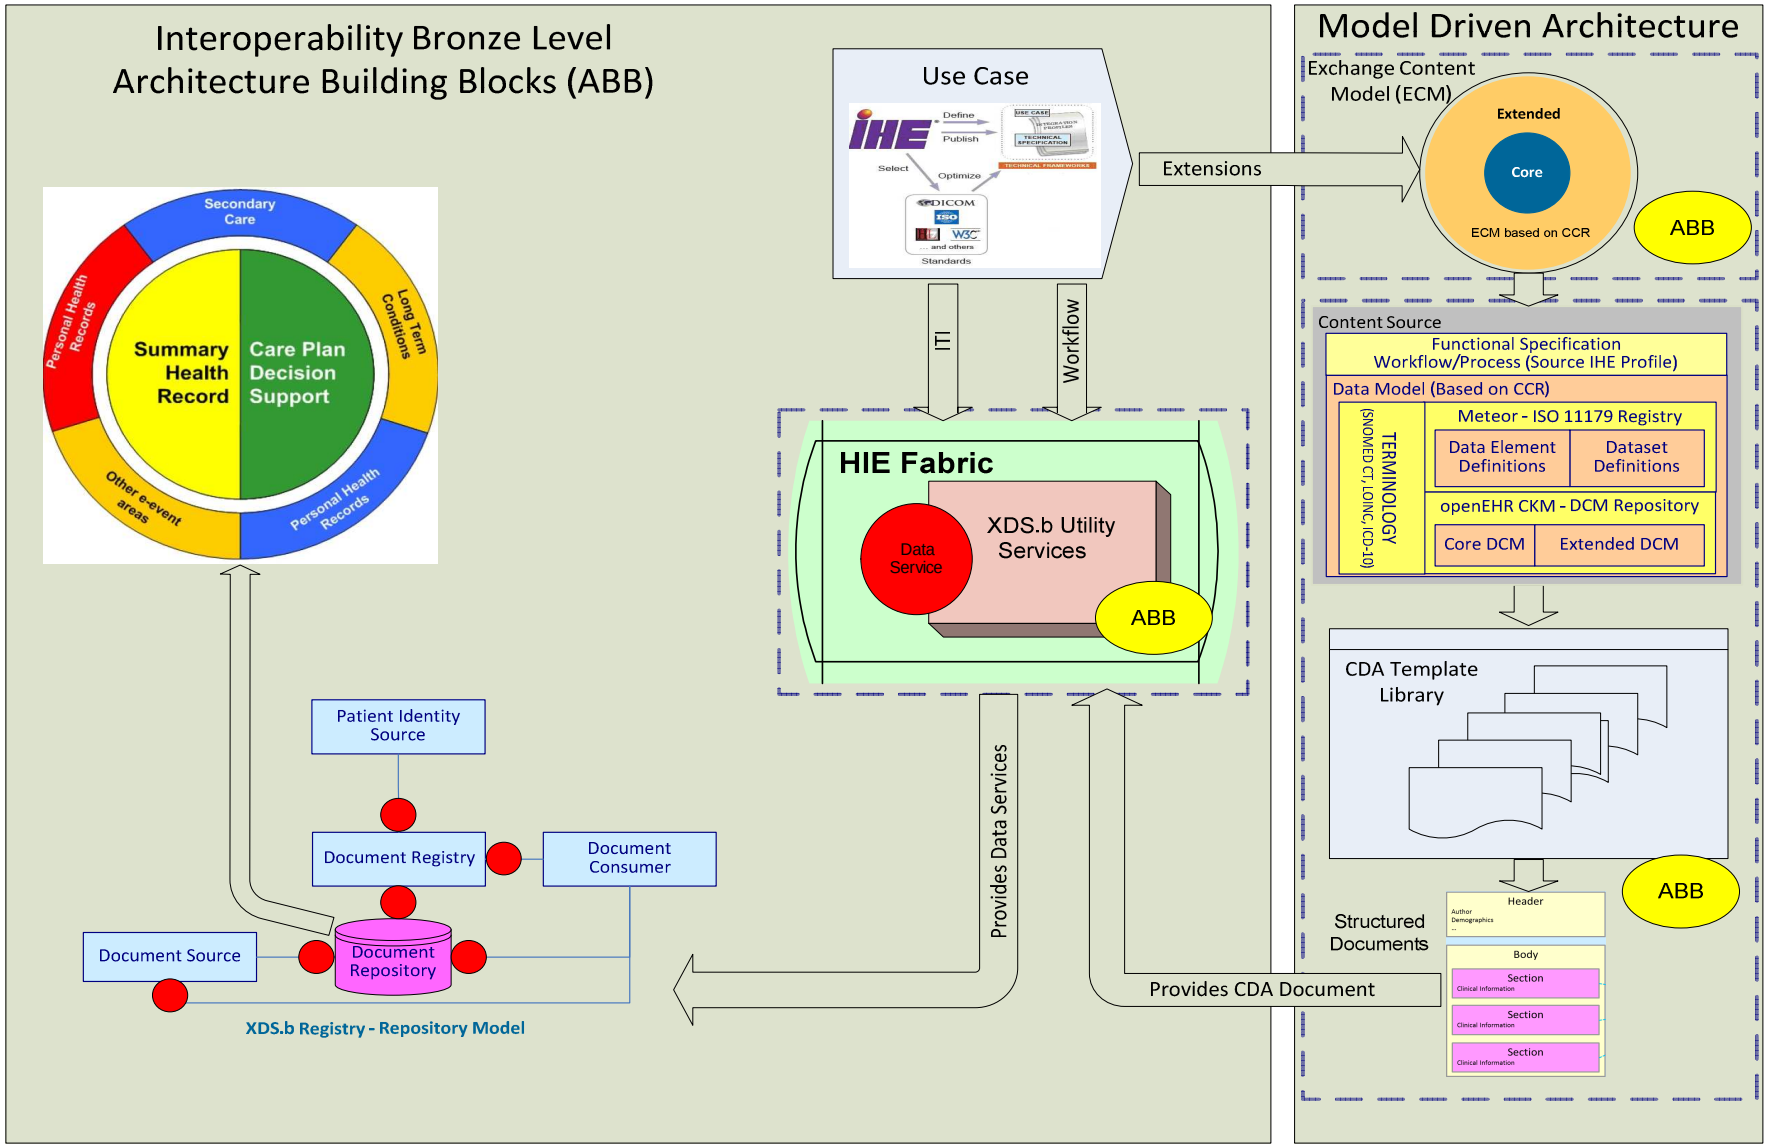

Figure 4 – Bronze Level ABBs of the Reference Architecture

6.4 HIE CDR Utility Services (Bronze.1) – HISO 10040.1

This section presents the *Health Information Exchange Clinical Data Repository Utility Services* architecture building block, a set of architectural requirements for information exchange based on R-CDRs and the registry-repository model.

The building block comprises architectural principles and requirements, organised under these headings:

- Utility services
- Registry-repository model
- Regional CDRs
- Transport services
- Identity services
- Security
- Document and image management
- Network requirements
- Terminology services

The required services are based on integration profiles in the IHE IT Infrastructure technical framework and collectively define a registry-repository model of information exchange.

XDS.b utility services are required to enable R-CDRs as described by the National IT Plan; they are considered Bronze level services. The services represent the registry-repository mandatory services for R-CDRs. This section describes the components and requirements related to this registry-repository model of information sharing.

The diagram shows documents moving between producers and consumers via a registry-repository hub: documents are stored in one or other repository while pointers to those documents are stored along with other metadata in a central registry. Consumers query the registry to locate documents and then retrieve them directly from the repositories identified.

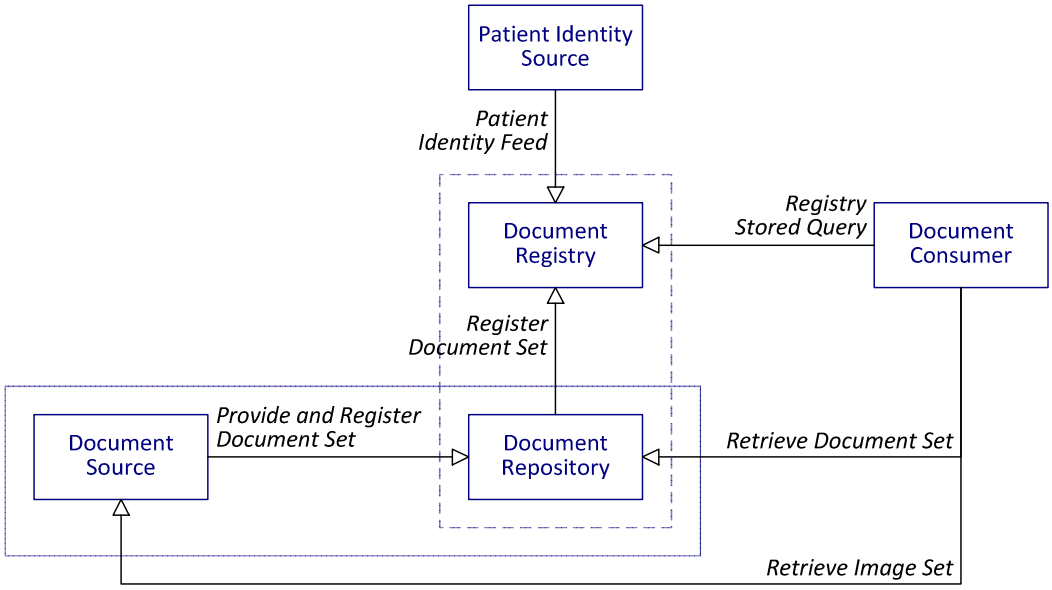

Figure 5 – IHE XDS integration profile

Document orientation and the registry-repository model are central to the IHE approach to interoperability. The diagram shows the actors and transactions of the XDS.b integration profile. Dashed lines indicate that the document repository can be collapsed into the same entity as either the document source or document registry.

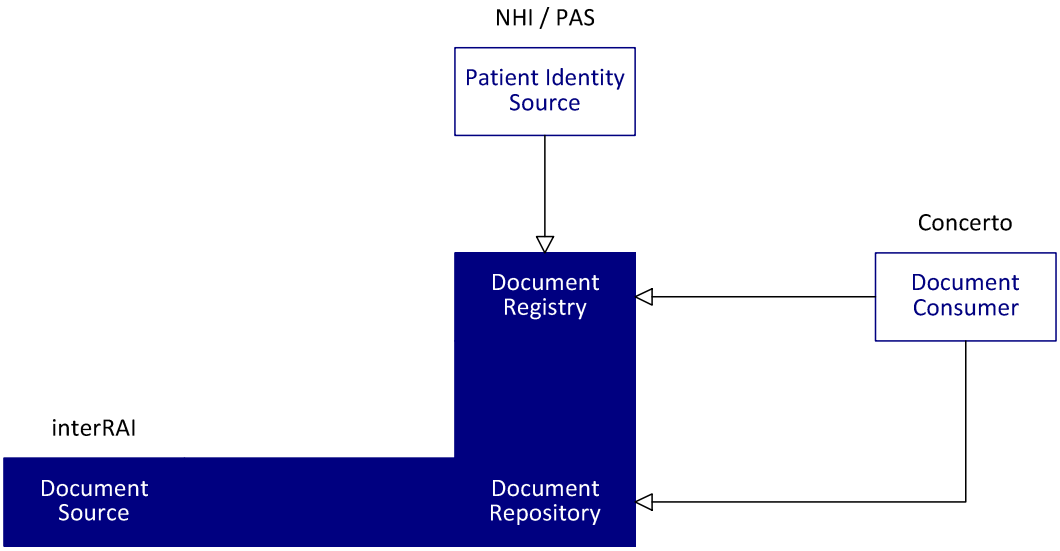

Figure 6 – IHE XDS.b and InterRAI

The interRAI national host system is being interfaced to Concerto to enable clinical workstation users to view home care assessment reports. The solution does not really follow the registry-repository model because as can be seen the document source, repository and registry are not clearly distinguished, but are collapsed together in the interRAI application. This has the major disadvantage that at runtime every time a new patient is brought into context in the clinical workstation, a query has to be directed at the interRAI system to look for assessments, when for most patients there won't be any.

6.4.1 Utility Services

Figure 7 shows the HIE as a standards-based fabric across which participant systems exchange information via services. R-CDRs and certain other national and regional systems produce services and point-of-care systems consume them. An example of this is a service that enables a General Practitioner (GP) Practice Management System (PMS) to update a medication list in an R-CDR.

The diagram shows the three service layers that are:

- Utility services (for basic operations)
- Entity services (that compose utility services for operations on representations of business entities – e.g. patient records)
- Task services (that compose entity services at the task or process level)

This building block addresses the utility services layer. It states requirements for a particular set of document sharing services in this layer. The services are based on the IHE XDS.b profile and associated profiles: PIX and PDQ identity services, XDS.b registry-repository services and ATNA audit services.

There are other kinds of utility services, but these are out of scope for this building block. Additional utility services will be defined in other standards as the requirements arise.

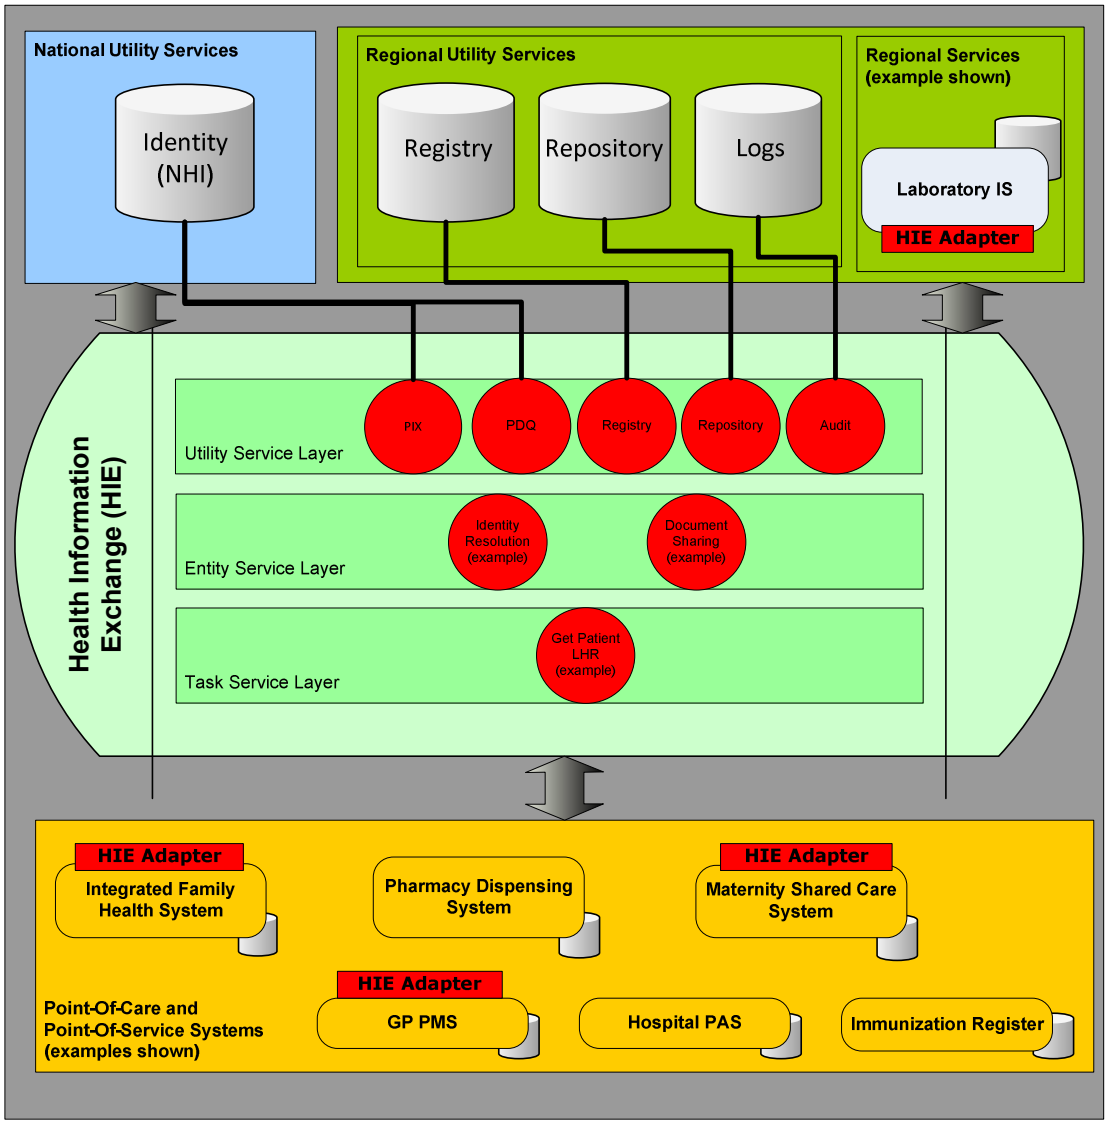

Figure 7 – HIE Utility Services

6.4.2 Registry-Repository Model

Figure 8 depicts a registry-repository model of information exchange, showing the various actors and their interactions. This particular model is defined by the IHE Cross Enterprise Document Sharing (XDS.b) integration profile and is a key requirement for the HIE.

Directive	HIE Transport
Statement	HIE transport shall follow the registry-repository model specified by the IHE XDS.b integration profile
Rationale	XDS.b is a standards-based specification for document-oriented information sharing that can be used to locate and retrieve patient information, with both regional and national views, while permitting regional independence in infra-structure and operations. XDS.b is based on standards ebXML, SOAP, MTOM/XOP, WS-I Basic profile, HTTP 1.1. Refer to the glossary for definitions.

Directive	Clinical Document Sharing
Statement	Clinical document sharing shall adhere to the registry-repository model as defined by XDS.b.
Rationale	XDS.b is a widely used profile internationally and ideally suited to the efficient sharing of clinical documents – both static and created dynamically.

Directive	Locating Services
Statement	Services will require an XDS.b registry for location requirements
Rationale	To store the WSDL documents for service location purposes. A UDDI registry was considered, but as the XDS.b registry has this capability. UDDI registry is an unnecessary extra system and UDDI is an end of life specification.

The registry-repository diagram only shows the connections between systems that are required for the IHE XDS.b profile. The Point of Care (PoC) Document Consumer can be any HIE-connected client application, for example a clinical portal or PMS application. The PoC Document Source or Image Source is any application that captures or stores patient information.

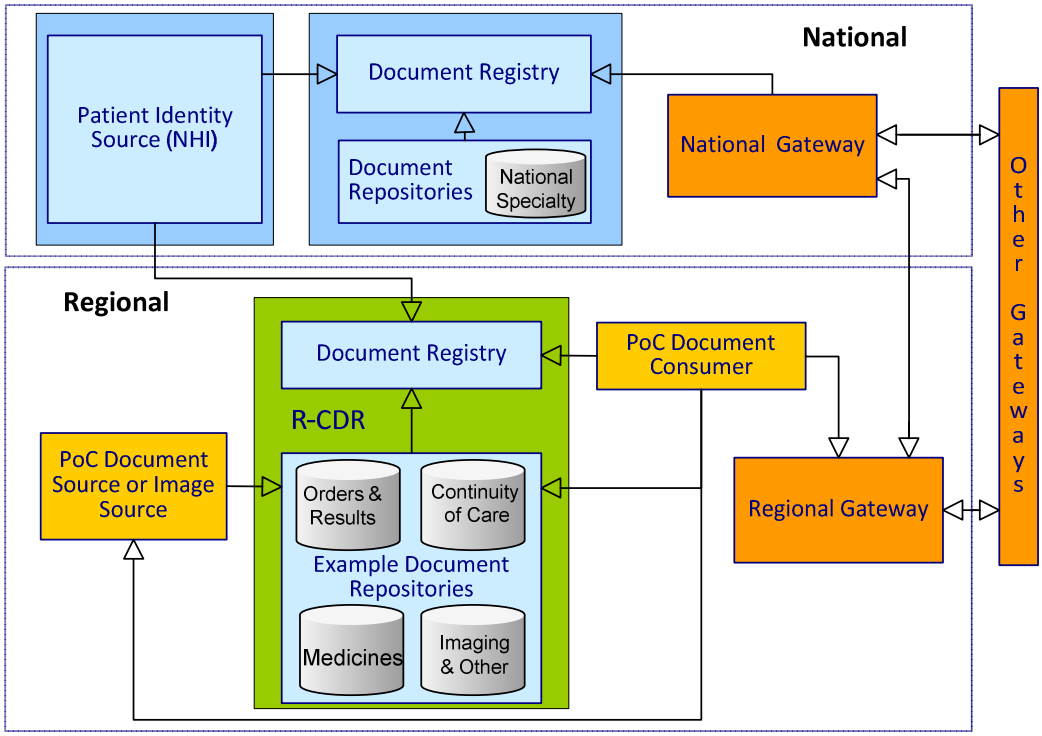

Figure 8 – Registry-Repository Model

Using XDS, stored patient information can include textual documents, coded documents and images. In the case of images, the actual image data is only stored in the source system (typically PACS - Pic-

ture Archiving and Communication System), while the repository stores an image manifest file using a Key Object Selection (KOS) format as a pointer. The registry then stores a pointer to the manifest.

The national registry and repositories collects national information, e.g. oncology and is made up of multiple repositories held by multiple specialist systems.

The affinity domain is a component of the XDS.b profile. The affinity domain defines the boundaries that the XDS.b registry will use, in this standard the R-CDR and affinity domain boundaries are the same.

There will be multiple affinity domains used by this standard. Figure 8 shows a regional affinity domain and a national affinity domain. Every region will have its own affinity domain and R-CDR. The private sector may also create compliant affinity domains. The purpose of the gateways (IHE XCA enabled) shown is to interconnect the affinity domains, to enable information sharing between R-CDRs and allowing for a single national patient view.

6.4.3 Regional CDRs

Figure 9 shows the R-CDR ecosystem in terms of XDS.b affinity domain components.

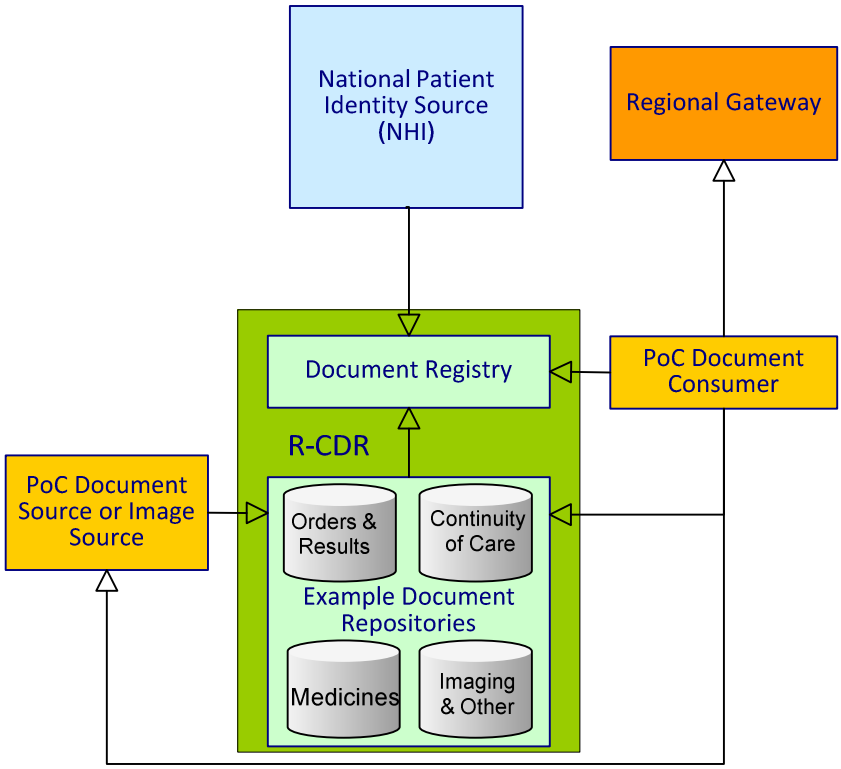

Figure 9 – R-CDR Components

Directive	Affinity Domain Per Region
Statement	There shall be one IHE XDS.b affinity domain per region. Each region will have federated XDS.b-enabled repositories with a single XDS.b registry. Each registry shall have a single affinity domain
Rationale	To support R-CDR requirements

Directive	Single Regional Registry
Statement	R-CDR shall include a single XDS.b registry. Every R-CDR is required to include a single XDS.b registry providing a record locator service over the content of the component CDRs.
Rationale	IHE XDS.b profile requires a single registry per affinity domain

Directive	XDS.b R-CDR
Statement	R-CDRs shall be XDS.b enabled registry-repositories model implementations. R-CDRs may comprise multiple component CDRs, each of which is required to implement the XDS.b integration profile and use the common XDS.b registry for that region.
Rationale	Repositories are required to be XDS.b enabled to allow them to be placed in the R-CDR ecosystem

Directive	R-CDR Purpose
Statement	R-CDRs may be general purpose or special purpose. The regional repositories can store clinical information in a configuration that is a best fit for the region
Rationale	The IHE XDS.b profile allows for multiple repositories per affinity domain, allowing regions flexibility on how the clinical information is held.

Directive	R-CDR Interconnection
Statement	R-CDRs shall interconnect using the IHE XCA integration profile. The regional gateway is an IHE Cross Community Access (XCA) enabled gateway; the gateway is used to find patient information outside the local region or affinity domain. The IHE XCA integration profile creates a network of communities by supporting the means to query and retrieve patient information held in other regions or affinity domains. Cross Gateway Query [ITI-38] and Cross Gateway Retrieve [ITI-39] are required transaction types.
Rationale	The XCA profile allows for the gateways across regional XDS.b affinity domains or other XDS.b affinity domains to find patient specific information, enabling national patient information to be presented to any consumer application.

The regional gateway is a Cross Community Access XCA-enabled gateway; the gateway is used to find patient information outside the local region or affinity domain. The XCA profile allows for the gateways across regional XDS.b affinity domains or other XDS.b affinity domains to find patient spe-

cific information, enabling national patient information to be presented to any consumer application.

The IHE XCA integration profile creates a network of communities by supporting the means to query and retrieve patient information held in other regions or affinity domains. Cross Gateway Query [ITI-38] and Cross Gateway Retrieve [ITI-39] are required transaction types.

Directive	Single Affinity Domain Policy
Statement	There shall be a single nationally agreed IHE XDS.b affinity domain policy. R-CDRs shall be configured and used in accord with a common, nationally agreed XDS.b affinity domain policy. The community of participants in the use of an R-CDR will form a regional XDS.b affinity domain. This community will be required to adhere to a common, nationally agreed affinity domain policy in configuring and using the R-CDR. This policy will apply equally in all regions individually and at national level. The policy will require governance over its content and application.
Rationale	The policy will describe a set of common rules and agreements, including formatting, naming conventions, working policies, document attributes and how to store data in document repositories

Directive	HIE Participant Systems Adapters
Statement	HIE participant systems may have service adapters. Participant systems are permitted to use adapters in order to produce and consume services on the HIE. That is, applications are not required to have native support for HIE protocols, but may instead use adapters.
Rationale	That is, applications are not required to have native support for HIE protocols, but may instead use adapters.

Online Forms

The HIE does not depend on or require the use of eForms, nor does it prevent their use. The HISO 10014.2 Online Forms Architecture Technical Specification concerns operations within a single user's session, while the HIE is about information exchange between separate participant systems. The two specifications are essentially independent and complementary.

6.4.4 Transport Services

Directive	HIE Standard Transport
Statement	The standard transport mechanism of the HIE shall be web services. Either SOAP over HTTP web services or REST web services are acceptable. Both styles of web services are supported by the IHE XDS.b transport protocol.
Rationale	Both styles of web services are supported by the IHE XDS.b transport protocol

Directive	HL7 v2 Containment
Statement	Use of HL7 v2 for transport is in containment. HL7 v2 may only be used when constraints on a particular solution make the use of web services impractical. The development of the building blocks is based on using HL7 v3 Clinical Document Architecture (CDA) R2. In this instance, users are encouraged to migrate to HL7 v3 CDA R2.
Rationale	HL7 v2 does not support the services based semantic interoperability, described in the Reference Architecture

6.4.5 Identity Services

The diagram shows required identity services.

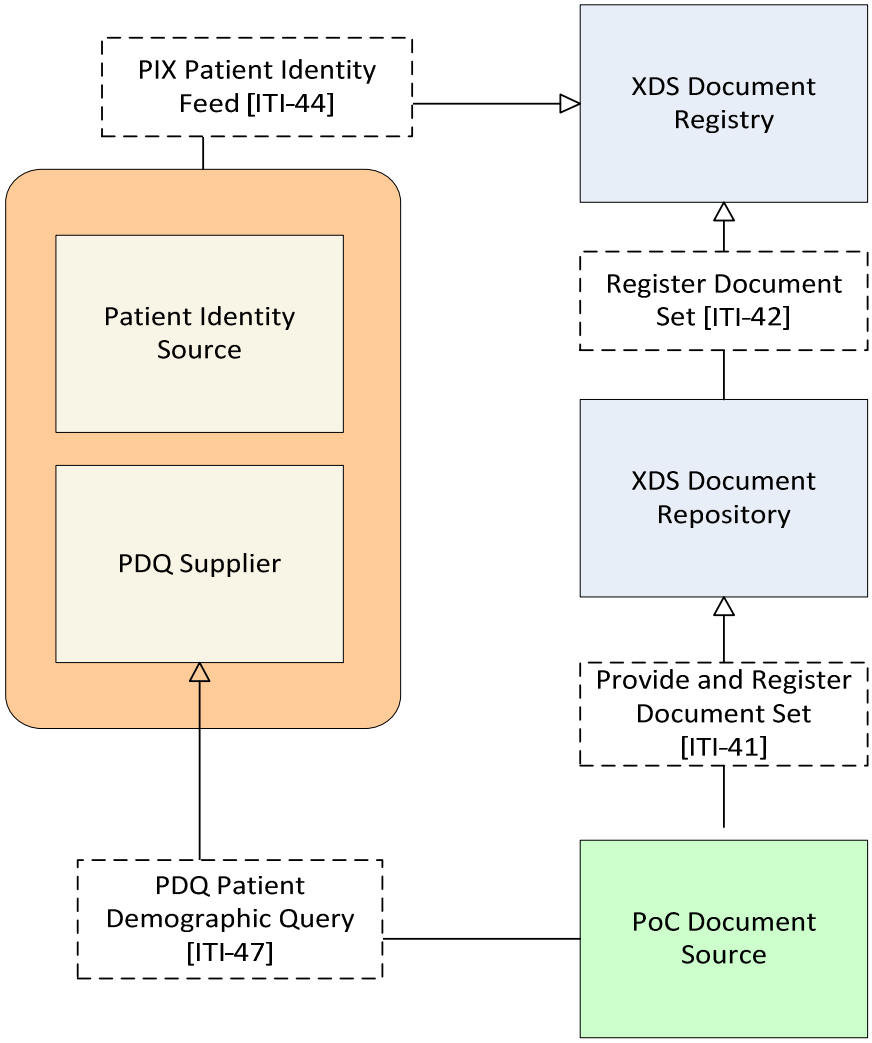

Figure 10 – Patient Identity Transactions

Directive	Authoritative Source Of Patient Identity
Statement	The authoritative source of patient identity information shall be the NHI. This rule will apply in all regions and affinity domains. This means that repositories and registries receive patient identity information from the National Health Index (NHI) and not from the regional Patient Administration system (PAS) or other regional patient index. It is important to note: there will be a nationally agreed, single affinity domain policy, using the NHI for the patient ID attribute
Rationale	Using the IHE XDS.b profile requires a patient ID to store information, an important requirement of using the IHE XDS.b profile affinity domain policy, is a unique and authoritative patient identifier

Directive	Patient Identity Source
Statement	Regional IHE XDS.b affinity domains shall use the NHI as the patient identity and demographics source, using IHE PIXV3 and IHE PDQV3 integration profiles PIX Patient Identity Feed HL7 V3 [ITI-44] is the required transaction type. PDQ Patient Demographic Query [ITI-47] is the required transaction type.
Rationale	The *IHE Patient Identity Cross-Reference HL7 V3 (PIXV3)* integration profile enables and patient identity feeds and patient identifier cross-referencing (although NHI numbers make this unnecessary here) The *IHE Patient Demographic Query HL7 V3 (PDQV3)* integration profile lets applications query a central patient information server in order to retrieve patient demographics and encounter information.

Directive	Health Provider Source
Statement	The authoritative source of health provider identity information shall be the HPI. This rule will apply in all regions and affinity domains. It requires HIE participants to use provider identity information directly from the Health Practitioner Index (HPI) (or any replacement to the HPI) rather than from regional alternatives.
Rationale	The HPI is an attribute of the single national affinity domain policy

Directive	Provider Identity Services
Statement	Provider identity services should conform to the IHE Health Provider Directory (HPD) integration profile. HPD specifies interaction between the source and consumers of provider identity information.
Rationale	For conformance to the IHE XDS.b profile

6.4.6 Security

Directive	Authentication, Access Control and Audit
Statement	Authentication, access control and audit requirements for document sharing shall be those defined by the IHE Audit Trail and Node Authentication (ATNA) integration profile. ATNA depends on standards Transport Layer Security (TLS), WS-I Basic Security Profile and Advanced Encryption Standard (AES). The Record Audit Event [ITI-20] is a required transaction type.
Rationale	XDS.b depends upon ATNA.

Directive	Digital Signatures
Statement	Digital signatures (where used) shall accord to the IHE Document Digital Signature (DSG) integration profile. W3C XML Signature is the underlying standard.
Rationale	For conformance to the IHE XDS.b profile

6.4.7 Document and Image Management

XDS.b requires that documents and document sets have metadata in the form of a defined set of document attributes. These document attributes are populated by the document source system and repository upon submission.

Required transaction types are shown in Figure 11.

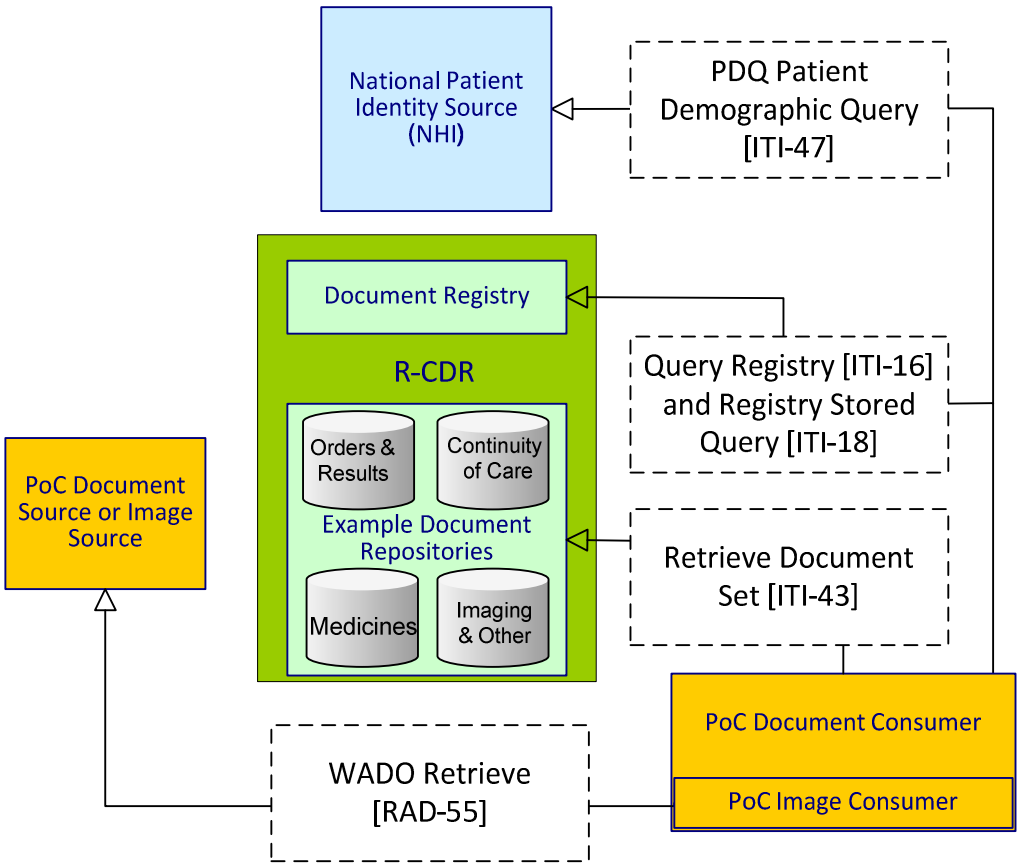

Figure 11 – Document and Image Retrieval

Directive	Medical Images
Statement	The IHE XDS-I integration profile shall be used for storage of and access to medical images *WADO Retrieve [RAD-55]* is a required transaction type XDS-I also supports medical images in non-DICOM formats. Such images can be PDF-embedded or hyperlinked. XDS-I also supports DICOM Structured Reports.
Rationale	The *IHE Cross-Enterprise Document Sharing for Imaging (XDS-I)* integration profile extends the XDS.b integration profile, enabling *DICOM Key Object Selection (KOS)* conformant image manifest files to be stored and registered in XDS.b repositories as pointers to DICOM images stored in PACS instances.

6.4.8 Network Requirements

Directive	Consistent Time Source
Statement	The communications network shall implement the IHE Consistent Time (CT) integration profile for network time synchronization. Network Time Protocol (NTP) is the underlying network time standard
Rationale	For conformance to the IHE XDS.b Profile

6.4.9 Terminology Services

Directive	Terminology Services
Statement	Terminology services shall conform to the HL7/OMG CTS2 specification. Where provided, terminology services are required to conform to HL7/OMG CTS2. CTS2 is an IHTSDO recommendation
Rationale	The use of SNOMED CT is a HISO requirement and CTS2 is the recommended service specification

6.4.10 IHE Registry-Repository Profiles

The table summarises the set of IHE integration profiles related to XDS and supporting the registry-repository model.

Functional Area	Integration Profiles
Document Sharing	*Cross Enterprise Document Sharing (XDS.b)* is for generic document storage, registration and record locator functions *Cross Enterprise Document Sharing Reliable Interchange (XDR)* and *Cross Enterprise Document Sharing Media Interchange (XDM)* are for point-to-point messaging *Cross Enterprise Document Sharing – Images (XDS-I)* enables DICOM/PACS images to be treated as documents *Cross-Enterprise Sharing of Scanned Documents (XDS-SD)* *Document Subscription (DSUB)*
Security	*Audit Trail and Node Authentication (ATNA)* describes digital certificate based authentication *Cross Enterprise User Authentication (XUA)* is for user authentication
Clinical Workstation	*Enterprise User Authentication (EUA)* and *Patient Synchronized Application (PSA)* repackage HL7 CCOW for single sign-on and patient context *Patient Demographics Query (PDQ)* describes access to identity and demo-

	graphics *Request Information for Display (RID)* provides simple (browser-based) read-only access to clinical information (e.g. allergies or lab results) located outside the user's current application *Retrieve Form for Data Capture (RFD)* enables applications to access public health reporting forms
Workflow	*Cross Enterprise Document Workflow (XDW)*
Identity	*Personnel White Pages (PWP)* provides basic directory information on human workforce members to other workforce members and applications
Network	*Consistent Time (CT*) repackages Network Time Protocol (NTP)

6.5 HIE Content Model (Bronze.2) – HISO 10040.2

This section presents the Health Information Exchange Content Model architecture building block, which frames a common shared content model to achieve semantic interoperability in information exchange.

The building block comprises architectural principles and requirements, organised under these headings:

- Semantic Interoperability
- Content Model
- Data Definitions
- Detailed Clinical Models
- Archetypes
- Terminology

6.5.1 Semantic Interoperability

Semantic interoperability requires information exchange to preserve meaning and context in a computable way. In order to achieve this, it is crucial to have a common language: the syntax, structure and semantics of information should be made explicit and shared. These definitions have to be formal and computable.

Directive	Information Exchanged
Statement	Information exchange shall be based on a definitive information model called the HIE Content Model. Clinical information needs to be exchanged in a format that complies with the HIE content Model
Rationale	The purpose of the Content Model is to enable semantic interoperability by providing fit-for-purpose, agreed and communicated data definitions

6.5.2 Content Model

This section describes the nature, scope, derivation and extensibility of the Content Model.

Figure 12 shows the standards used by the Content Model. The components of the diagram are discussed in more detail later in this document.

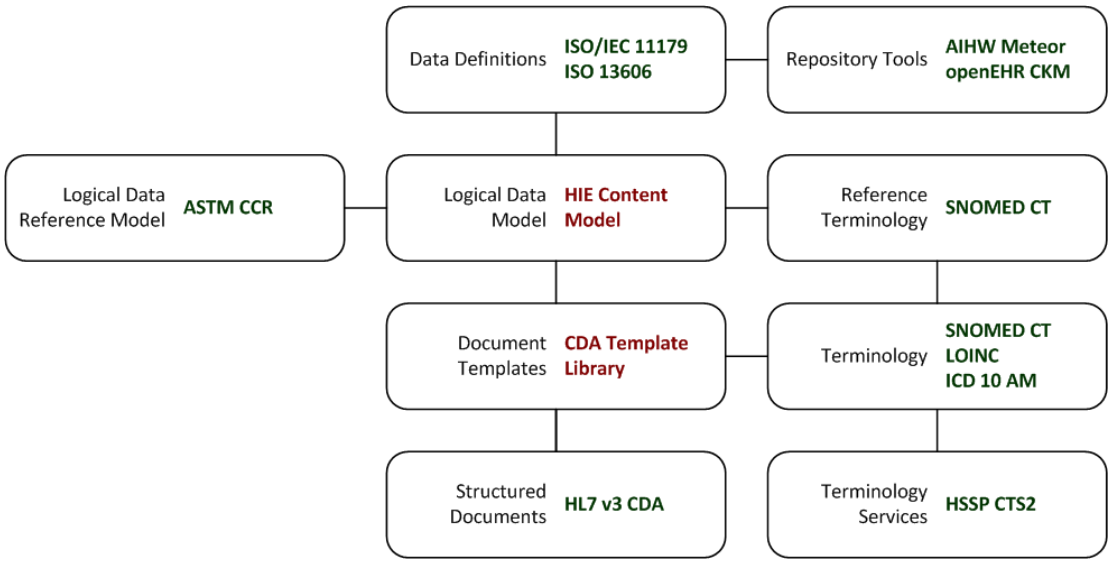

Figure 12 – Content Model Standards

Directive	Core Content Model
Statement	The Content Model shall derive from the ASTM Continuity of Care Record (CCR) specification • The Content Model will not adopt CCR in its entirety - it will be adapted using the following ways: • Subject areas – the Content Model will adopt CCR's subject area headings and scope • Conceptual data elements – the Content Model will adopt CCR's conceptual data element definitions in each subject area • Business descriptions – the Content Model will borrow CCR's business descriptions in each subject area CCR will be used as the logical data framework of the Content Model. The Content Model will tend to localise CCR and overall will be geared towards practical alignment with it rather than total conformance. The main departures are that the CCR XML schema will not be used, and that the contents of the model will be extended to suit New Zealand requirements. The population of data definitions, within the Content Model will have various sources such as from Detailed Clinical Models and IHE Content Profiles.
Rationale	CCR provides the logical data reference model this allows the data model to be aligned with an international standard, but allowing contents to suit New Zealand requirements

Figure 13 shows CCR's place in defining subject areas of the Content Model, with the 'keyhole' indicating extension into some specialty area.

This is a superset of potential health information, and work on R-CDRs and shared care records may determine that smaller subsets are used in practice, subject to information governance policy.

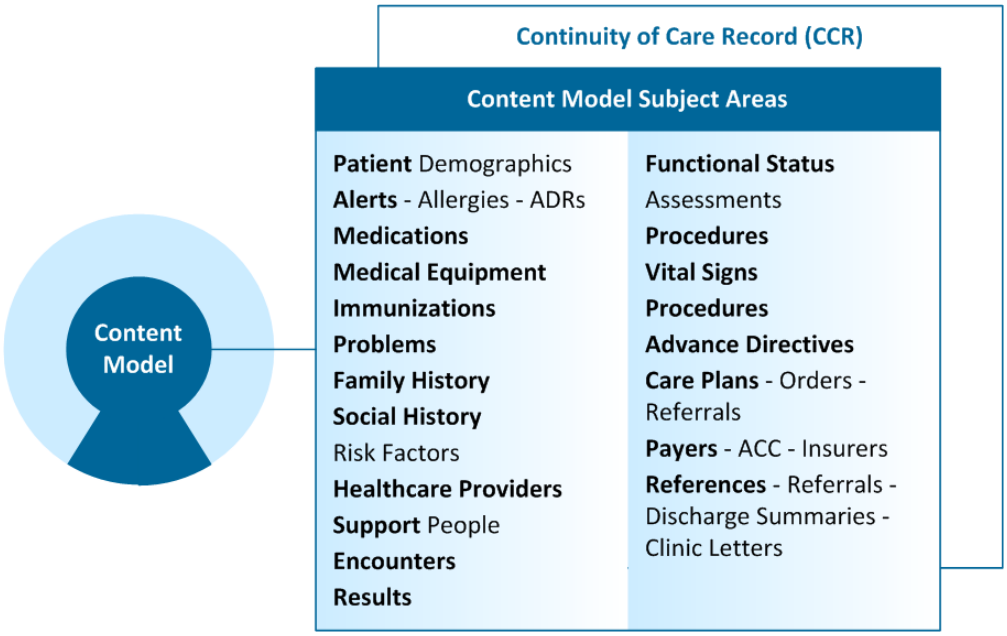

Figure 13 – Content Model Subject Areas

Directive	Extensible Content Model
Statement	The Content Model shall be extensible. The scope of the Content Model will increase over time as data requirements in specialty areas are established and documented
Rationale	The Content Model must be able to accommodate any business driven change and be resilient to its effects. Using CCR as the Data Reference Model, the process of developing the Content Model in some subject area will in technical terms involve specialising the classes adopted from CCR

6.5.3 Extending the Content Model

The methodology for extending the content model embodies the following principles:

- New items added on top of core items in extended parts for different specialties
- Items providing more detail to existing items in core can be added
- Existing core items can be constrained (e.g. provide pick lists for free text areas or value ranges can be defined for numeric fields etc.)

Figure 14 shows examples of extending the Content Model to allow three additional specialty areas.

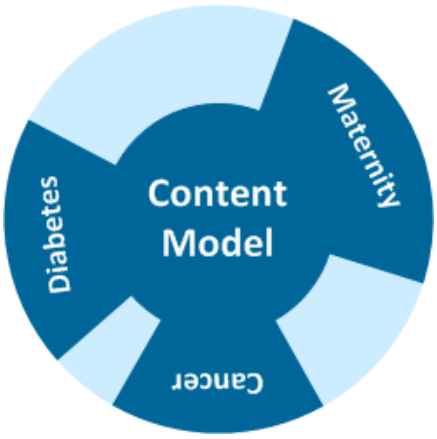

Figure 14 – Example of Content Model Extension

Reuse of definitions throughout the Content Model is essential.

Extensions will be based, wherever possible, on existing proven definitions/programmes developed by other recognised / established organisations. This will achieve not only internal consistency in the Content Model but will also promote its international alignment. There may also be additional local requirements.

6.5.4 Data Definitions

This section describes how the data definitions of the Content Model will be formally expressed.

Directive	Formulating Content Model Data Definitions
Statement	Content Model data definitions shall be formulated according to the ISO/IEC 11179 metadata standard This is the authoritative expression of the Content Model. In ISO/IEC 11179 terms, the Content Model will comprise definitions of data-sets, data elements, data element concepts, object classes, properties, value domains and classification schemes. Dataset definitions will be used to derive CDA section templates
Rationale	The national registry tool for the content model conforms to ISO/IEC 11179

Directive	Registering Content Model Data Definitions
Statement	Content Model data definitions shall be registered in accord with ISO/IEC 11179 processes and stored in a compliant registry AIHW METeOR is the chosen metadata registry tool and is ISO/IEC 11179 compliant. HISO 10014.1 Data Concept Repository Processes Standard is the local adaptation of HISO endorsed ISO/IEC 11179. It is currently under view to bring it up-to-date with the selection of METeOR and evolution of the base standard.

Rationale	The use of a single metadata registry with search facilities to hold the details of the Content Model and its datasets will make use and extension of the model easier and promote consistency and reuse

Directive	Content Model And ISO 13606 archetypes
Statement	The data definitions of the Content Model should be formulated as ISO 13606 archetypes This is an alternative expression of the Content Model. The use of archetypes is recommended rather than required
Rationale	See the sections following on DCMs and archetypes for the reasons behind this dual approach

Directive	Units Of Measure
Statement	The *Unified Code for Units of Measure* (UCUM) is the units of measure code system that should be followed for electronic data interchange
Rationale	*The Unified Code for Units of Measure* provides a single coding system for units that is complete, free of all ambiguities, and that assigns to each defined unit a concise semantics

6.5.5 Detailed Clinical Models

This section describes the development of the Content Model under the Detailed Clinical Model (DCM) approach.

Directive	Content Model Development
Statement	Development of the Content Model shall follow the DCM approach The DCM approach is about creating discrete, modular, reusable, controlled and above all rounded specifications of information requirements in some clinical domain or sub-domain. This approach supports the above keyhole concept, by allowing the Content Model to be built up piecemeal around an established core
Rationale	The DCM approach is about creating discrete, modular, reusable, controlled and above all rounded specifications of information requirements in some clinical domain or sub-domain

Directive	DCMs Reuse
Statement	DCMs may be reused from other national programmes DCMs may be adopted and adapted from other national programmes, to save time and effort in developing the Content Model
Rationale	To save time and effort in developing the Content Model

Directive	DCMs Maximal Datasets
Statement	DCMs shall define maximal datasets. DCMs will define maximal datasets, i.e. they will include all possible data elements that may be mandatory, optional or inapplicable depending on the application or context
Rationale	Promotes the reuse and effectiveness of the DCMs

6.5.6 Archetypes

Archetypes are a robust way of describing structured health information in a way that can easily be understood and maintained. They are suited to the involvement of healthcare professionals in the development process. They combine healthcare concepts, clinical context, data elements and their organisation, terminology and metadata in a technology agnostic and computable way. Practically they specify labels, data structures, types and valid value ranges and enumerations. The premise of archetypes is that data, user interface, information exchange and integration are all based on the same specifications.

This section describes the use of ISO 13606 archetypes as another means (in addition to ISO/IEC 11179) of expressing DCMs making up the Content Model.

Directive	Expressing DCMs
Statement	ISO 13606 archetypes may be used to develop and express DCMs. Finished archetypes may be inputs to the process of formulating new ISO/IEC 11179 data definitions.
Rationale	Archetypes lend themselves to the development of new DCMs, and representing DCMs in graphical form

Directive	Archetypes Non-Native Forms
Statement	Archetypes may be transformed into other information modelling forms. Archetypes may be transformed as required from their native ADL representation into human readable (e.g. mind maps, UML) or computable (e.g. XML, CDA)

	artefacts
Rationale	While it is possible to do this transformation automatically by creating XSLT transforms per DCM, this will be progressed at a later stage

Directive	Archetype Library
Statement	There shall be a shared archetype library. The openEHR Clinical Knowledge Manager (CKM) web-based tool will be used to provide the shared archetype library
Rationale	The CKM intuitive user interface, graphical outputs and editorial process support and foster clinician engagement

6.5.7 Terminology

Directive	SNOMED CT Reference Sets
Statement	SNOMED CT Reference Sets shall be used wherever possible. SNOMED CT Reference Sets are the default choice of terminology for data elements in the Content Model. The exception of this directive is when a SNOMED CT Reference Set does not exist or has not been endorsed for use in New Zealand or when another HISO standard (such as NZPOCS) requires another terminology and has precedence.
Rationale	LOINC, ICD 10 AM and ICD 0 coding and classification systems – all mapped to SNOMED CT – have HISO endorsement and are acceptable in their respective domains. The New Zealand Medicines Terminology (NZMT) is a SNOMED CT Reference Set. CCR as the data reference model supports the use of SNOMED CT Reference Sets in most subject areas.

Directive	Terminology Bindings
Statement	The Content Model shall have explicit terminology bindings. Data elements in the Content Model should be directly associated with exactly one SNOMED CT Reference Set or another permitted coding system
Rationale	ISO/IEC 11179 data elements support explicit terminology bindings using value domains. Terminology bindings are required for semantic interoperability

6.5.8 CCR Use Case Example

This is an example use case for deriving data definitions from CCR. This example concerns the problem list. The following rules can be derived from the specification list.

- CCR supports a problem list, which may be any length, of the patient's current and resolved problems
- A problem may be classified as either a condition, diagnosis, symptom, finding, complaint or functional limitation
- Problems can be described using SNOMED CT and/or narratively
- Problems have a status of either active, inactive, chronic, intermittent, recurrent, or resolved
- Problem episodes are recorded
- The problem list can be ranked or filtered by date of onset or order of importance, e.g. for a referral
- The source of problem information may be recorded, including who and when
- Whether the subject is aware of the problem – and if not, why not – can be recorded
- There's a link to medications – when a listed problem is an indication for certain medication
- Details of functional limitation may be recorded against a problem
- Clinical documents may be associated with problems
- Problems may be recorded as the cause of allergies or adverse reactions
- The existence of a problem may be flagged as an alert
- Orders and results may be linked to problems
- A problem may be an indication for a procedure
- A problem may be a reason for an encounter
- Family history may be expressed in terms of problems (not necessarily associated with an individual)

6.6 HIE Structured Documents (Bronze.3) – HISO 10040.3

This section presents the Health Information Exchange Structured Documents architecture building block, a set of architectural requirements for the use of structured documents as the common currency of information exchange.

The building block comprises architectural principles and requirements, organised under these headings:

- Structured Documents
- Incremental Interoperability
- Data Extraction
- Document Metadata
- Document Security
- Document Attachments
- Standard Document Definitions
- Template Library

6.6.1 Structured Documents

The essence of this building block is the use of structured documents as the form in which information is exchanged between HIE participants. Structured documents have characteristics of wholeness, persistence, transport independence that differentiates this approach from the earlier messaging approach. Documents by themselves do not embody workflow; instead we rely on services to provide the necessary context and transactional capability.

Directive	Structured Documents
Statement	Structured documents shall be the common currency of information exchange. Documents may include clinical documents in a conventional sense, authored by a clinician, but also documents generated as part of the workings of interoperable systems, including clinical data repositories and shared care records
Rationale	This is the concept of document orientation. Documents have characteristics of context, purpose, wholeness, persistence and attestation

Directive	Document Types
Statement	Document types shall conform to the HL7 Version 3 Clinical Document Architecture, Release 2. HL7 Version 3 Clinical Document Architecture, Release 2 is the selected document type for all clinical information exchanged on the HIE
Rationale	Internationally, CDA has emerged as the format of choice for representing and packaging structured clinical information for exchange. It has significant tooling, documentation and deployments. This requirement for CDA builds upon the HISO endorsement of HL7 as a whole CDA is based on the HL7 v3 Reference Information Model (RIM). CDA documents are encoded using XML 1.0. There is a single defined XML schema to which all CDA documents must conform This requirement for CDA builds upon the HISO endorsement of HL7 as a whole

Directive	Persistent Documents
Statement	Documents shall be persistent. Documents will exist over time and may be reused in many contexts Document lifespan will be subject to retention rules Once created, a document can be replaced by an updated version, but the original should always be retained by the custodian, and referenced by the updated version
Rationale	Required for Information Security and Governance compliance

Directive	Documents Custodianship
Statement	Documents shall have defined custodianship. Each clinical document instance is maintained (managed, shared) by an organisation entrusted with its care
Rationale	Required for Information Security and Governance compliance

Directive	Documents Authentication
Statement	Documents shall have the potential for authentication. A clinical document is an assemblage of information that may be intended as medico-legal documentation, and contains elements that may be used to identify individuals who assert the correctness of the information within. The actual authentication method is outside the scope of this standard
Rationale	Required for Information Security and Governance compliance

Directive	Documents Context
Statement	Documents shall establish the context for their contents. A clinical document should have information establishing the default context for its contents. There is a requirement to include relevant contextual information surrounding the event documented – including details of the patient, the practitioner, the setting and the purpose
Rationale	Required for Information Security and Governance compliance

Directive	Documents Wholeness
Statement	Documents shall demonstrate wholeness. Wholeness means complete and self-contained, wholeness also covers integrity and data quality. Authentication of a clinical document applies to the whole and does not apply to portions of the document without the full context of the document
Rationale	Required for Information Security and Governance compliance

Directive	Documents Transport
Statement	Documents shall be transport independent. Documents should not be dependent on any particular transport protocol. The author of the document is responsible for the accuracy of the document content. In practice, IHE XDS.b as the transport protocol of the HIE will be that required

	most often. (See the HIE CDR Utility Services architecture building block for details.)
Rationale	Limits document use and does not future proof the documents if they are locked into a particular transport protocol

Directive	Documents Workflow
Statement	Documents shall not embody workflow. Preferred workflow mechanisms include the IHE XDS/XDW profiles, or services built from these, or carriage within an HL7 v2 or v3 message
Rationale	CDA documents in themselves are content only, they are not meant to trigger action in the way that HL7 v2 messages may do

Directive	CDA and HL7 V2
Statement	HL7 Version 2 use should be contained to solutions where CDA is not a viable alternative. This is generally where there is an existing HL7 v2 implementation, and there is no business driver to change. In particular, a number of HISO standards specify HL7 v2 and they have continuing effect irrespective of the CDA requirement specified here. New standards should be CDA-based from the outset, while existing standards will reviewed. The existing standards affected are: HISO 10011: Referrals, Status and Discharges (RSD) HISO 10030: EPharmaceutical HISO 10008 Pathology and Radiology
Rationale	HL7 V2 does not deliver the services based semantic interoperability required

Directive	CDA and HL7 V3
Statement	HL7 Version 3 messaging use (i.e. not CDA) should be contained to solutions where CDA is not a viable alternative. Certain IHE integration profiles require some limited HL7 v3 messaging. PIXV3 and PDQV3 are examples of the requirement to use HL7 v3 web services based messaging
Rationale	HL7 V3 messaging requires an expertise the New Zealand Health Sector does not currently have and there is no business requirement at this stage. Using existing IHE integration profiles negate the learning curve

6.6.2 Incremental Interoperability

Some information systems will be more capable than others in producing or consuming document content, especially when it comes to form and level of detail. Incremental interoperability is about setting minimum requirements based on human readable content, but at the same time permitting equivalent machine-readable content.

The principles of incremental interoperability are:

- Human readable content is required
- Discrete elements are optional
- Sources provide details as per their capabilities
- Consumers show human readable content
- Consumers handle discrete data as per their capabilities and its presence

This allows systems to transmit documents that have varying degrees of structure, ranging from simple text that can be read by a human, through text that is structured (i.e. has defined sections) to documents that can be parsed by an automated process.

- CDA enables incremental interoperability as follows.
- A CDA document has two parts: a header and a body:
- The header contains structured metadata about the information in the document. Examples are the author of the document, the patient who it is about, other people involved in the episode being described, etc.
  - The body contains the clinical information. This information can be represented in a number of ways, commonly described as being at one of three levels:
  - Level 1 documents have the structured header, but the content is unstructured or encapsulated, e.g. PDF, GIF or PNG with a MIME type that defines the encoding used. Each CDA document can contain only a single body.
  - Level 2 documents contain any number of sections, with each section representing a particular type of information such as medications, problem list, discharge diagnoses, etc. The information is presented as text, which means that it is human readable but not machine-readable. Each section has an identifying code.
  - Level 3 documents are also comprised of sections, but each section can also present information in a computable form. Computable data must be present in an equivalent textual form, for human readability. Each section has an associated template identifier.

Directive	Human Readable Documents
Statement	Documents shall be human readable. A clinical document should be human readable – i.e. information is (at least) in textual format. This does not pertain to the underlying XML being directly readable as to the inclusion of narrative elements at the CDA level containing the clinical information, attested by the author as correct
Rationale	Human readable documents make the documents use more flexible

Directive	CDA Level 3
Statement	Documents should be at CDA level 3. The preferred level is level 3. CDA level 3 is computable and provides the best support for semantic interoperability. This requirement is subject to the ability of the source to provide information at this level of detail. CDA level 1 use should be contained to solutions where level 2/3 is not a viable alternative CDA level 2 use should be contained to solutions where level 3 is not a viable alternative CDA level 1 and 2 use is still permitted, but contained as above
Rationale	Enables full services based semantic interoperability; because CDA level 3 is computable it provides the best support for semantic interoperability

6.6.3 Data Extraction

Directive	Document Information Extract
Statement	Information may be extracted from the document. A recipient system may extract data from the document to update its own internal stores. Care must be taken to preserve contextual information (e.g. about the patient and the source of information). Preferably, there should be a link back to the original document for audit purposes
Rationale	The extracted data can be used to populate other information stores

6.6.4 Document Metadata

Directive	Document Metadata
Statement	Document type (template) metadata shall be standardised. Metadata in CDA documents is contained as a set of attributes in the document header. Each document type defines which header elements must be present and how they are derived. These will not necessarily be the same for every document type. For example the author of a discharge summary is a person, while the author of a GP2GP message is a process
Rationale	The metadata is Important as it forms document attribute information for repository storage and retrieval

The required document attributes are as follows:

Attribute	Description	Domain
Document Type	Recognised document type	OID
Title	A human readable title	Text
Document Identifier	Persistent unique identifier for the document within the HIE and R-CDRs	GUID
Version Number	Allows a document to indicate that replaces a previous document	Positive integer
Creation Time	When the document was created (not the time of the event being documented)	Timestamp
Confidentiality	Allows a document to be recognised as normal or restricted. Normal maps to medical in confidence and restricted maps to medical sensitive	N = Normal R = Restricted
Patient	Who the document is about	NHI number
Author	Who wrote the document (person or system actor)	HPI number
Legal Authenticator	Who asserts that the information in the document is correct	HPI number
Custodian	The organisation responsible for keeping a copy of the document	HPI number

Figure 15 – Attributes for Document Header

Directive	Document Unique Identifier
Statement	Every document (instance) shall have a unique identifier. Usually a GUID, the document identifier identifies a particular document instance. An update to a document will have its own (different) identifier, and will contain a link back to the original
Rationale	Persistent unique identifier for the document within the HIE and R-CDRs

6.6.5 Document Security

Security is managed by the transport system. The CDA document does not currently support digital signatures.

Directive	Document Confidentiality
Statement	Document content shall be secured to ensure confidential, integrity, authenticity and non-repudiation to sender and recipient systems. Encryption is required along the transport system. Content must not be viewed during transport
Rationale	Required for information security and governance compliance

6.6.6 Document Attachments

Directive	Documents Attachments
Statement	Documents may have attachments. A CDA document can reference external attachments (such as images or PDF files) but cannot contain them inline within the document (except as the body of a level 1 document). The attachments are generally transported with the CDA document, most commonly as a MIME package, although options such as MTOM are permissible
Rationale	Attachments may be needed to accommodate additional information

6.6.7 Standard Document Definitions

A CDA document consists of a header (containing the contextual information) and a set of sections that holds the clinical data. A particular document type (e.g. for discharge summaries) uses the header information in a consistent way, and has a selection of sections that may be present (the definition of the document can define sections that must (or may) be present).

Each defined document type can be considered a template with an identifying template identifier, and should reference the use cases that it fulfils.

Priority business use cases under the National Health IT Plan are as follows. Each will require specific document types for information exchange.

- Core Health Information
- eReferral
- eDischarge
- Community ePrescribing
- Long Term Conditions Shared Care Record
- Maternity Shared Care Record

Directive	Template Identifier
Statement	Every approved template shall have an identifier derived from the HL7 New Zealand OID root. The HL7 New Zealand OID root is 2.16.840.1.113883.2.18 (as recorded in the HL7 International OID Registry)
Rationale	It is important that every template can be uniquely identified

Directive	Template Data Model
Statement	Template definitions shall conform to the HIE Content Model. All template definitions – header and body – are required to be consistent with the datasets and data elements of the HIE Content Model
Rationale	Templates need to conform to the HIE Content Model to natively support the functions of the HIE.

Figure 16 shows the template design process. It is driven by the data definitions of the Content Model and functional specifications.

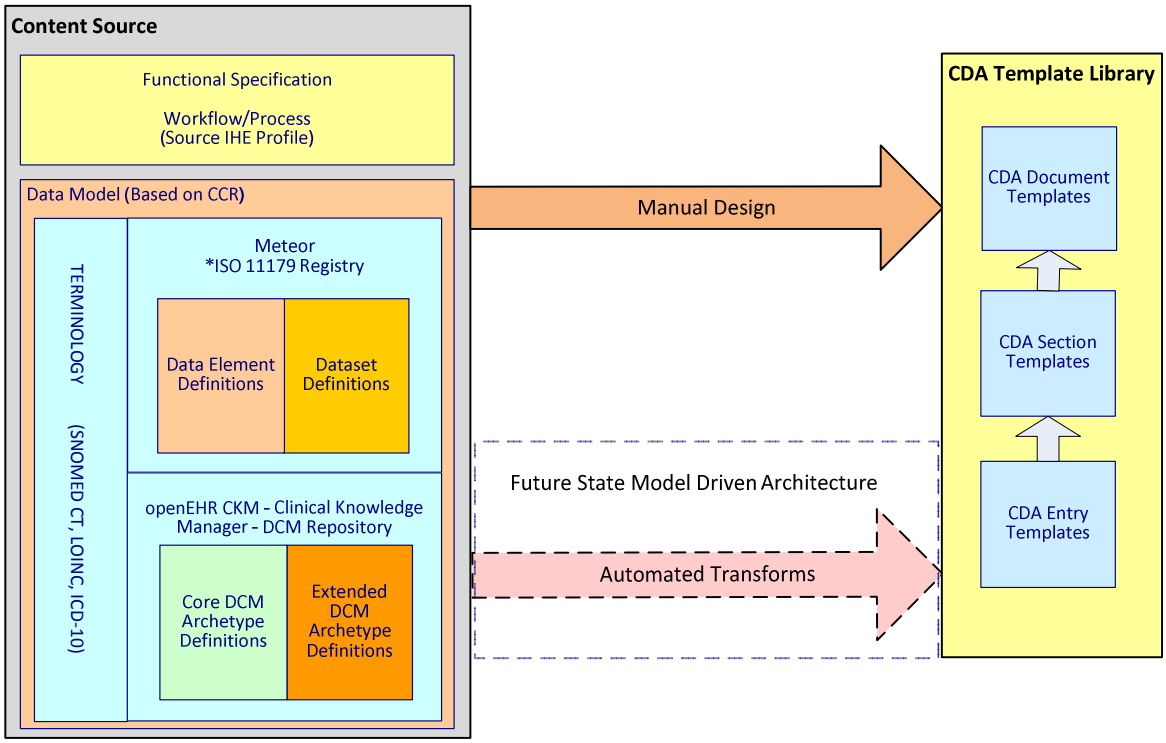

Figure 16 – Template Design Process

Directive	Document Sections
Statement	Document type (template) shall have a defined set of sections. A document type will define a set of permissible sections that may be contained. It will define whether each section is required or optional
Rationale	Sections should be re-usable across document types, thus establishing a library of standardised, re-usable sections

Directive	Templates
Statement	Templates shall be specified in either of two forms. (a) natural language definitions, within implementation guides, or (b) sets of Schematron rules. It is also required that any CDA document will be conformant to the CDA R2 XML schema
Rationale	To ensure conformity with HL7 CDA R2

Directive	Section Template Unique Identifier
Statement	When any section template is modified the process should be to create a new template identifier
Rationale	This reduces the chance of incompatibility problems

Directive	Section Template
Statement	A section template may be either CDA level 2 or level 3. Sections contain the clinical information being transferred between systems. A level 2 or level 3 CDA document can have any number of sections. The document type will define what are the allowable sections for a particular type and whether they are required or optional
Rationale	CDA sections are a requirement of level 2 and level 3 documents. Sections contain the clinical information being transferred between systems

A section can be either level 2 or level 3:

Level 2 sections have a <text> element that contains the human readable information to be transmitted

Level 3 sections also have the text element, but in addition contain any number of <entry> elements that have the same information in coded form. Any information in <entry> elements must also be present in the <text> element. However, not all information in the <text> element must be coded.

Section components are as follows.

Item	Description
Section Code	Identifies the section (e.g. discharge medications, problem list)
Title	Human readable title for the section
Text	Human readable text

Coded Entries	The same information as in the text element but in coded form Level 3 only

Figure 17 – Section Components

The <entry> element structure will vary according to the information being represented, and can become complex. The implementation guides for a particular domain should describe these, but the rules for reuse should also apply as noted above. This promotes consistency between documents.

There are other elements as well that allow an originator to describe other actors involved in the encounters described by the document plus a wide variety of other information. These need to be documented when the document type is defined.

Directive	Data Type Standard
Statement	Data types should follow the ISO 21090 standard. ISO 21090 provides a set of data types in support of information exchange. It is derived (in part) from the HL7 v3 data types specification and is compatible with CDA
Rationale	Provides a set of data type definitions for representing and exchanging basic concepts that are commonly encountered in healthcare environments in support of information exchange in the healthcare environment

6.6.8 Template Library

Directives under this heading cover:

- Template library
- Template library governance
- Template registration process

Directive	Template Library
Statement	There shall be a library of standard templates for information exchange. The library will include document templates, section templates and entry templates
Rationale	Standard templates can be reused, reducing time and cost of delivery

Directive	Template Library Governance
Statement	There shall be a governance body with responsibility for the template library. It is recommended that there be a national governance framework around the implementation of CDA documents in New Zealand and management of the template library
Rationale	Templates need to conform to the HIE Content Model and must be reusable

Directive	Template Registration Process
Statement	Templates shall be subject to a formal registration process modelled on ISO/IEC 11179 for metadata registration. The participant roles in the process are: Originator – submits new templates for registration; originators must first be accredited by the Registration Authority Registrar – receives submissions and coordinates the registration process Technical Committee – resolves, harmonises and moderates Steward – reviews and maintains as the subject matter expert Executive Committee – governs the registration process Registration Authority – owns the registration process and is custodian of the template library
Rationale	To ensure template quality and consistency is maintained

7 Silver Level Architecture Building Blocks

Silver level architecture building blocks are scoped for the second phase of development of the Reference Architecture.

7.1 Point to Point Messaging Utility Services (Silver.1)

The ABB will contain the requirements for point-to-point messaging based on the IHE XDR/XDM integration profiles. (See the point-to-point directive in the general directive section)

7.2 Terminology Utility Service (Silver.2)

The ABB will contain the requirements for terminology utility service based on the HSSP CT2 terminology service. (See the terminology service directive in the general directive section)

7.3 Shared Diagnostics Ordering and Reporting Task Services (Silver.3)

The Task Services that form this ABB will be an enabler to the Continuum of Care workstream described in the National Health IT Plan. This ABB will be based on the IHE XD-LAB and IHE LTW profiles.

7.4 Health Provider Entity Services (Silver.4)

The ABB will contain the requirements for Health Provider Index service based on the IHE HPD integration profile.

7.5 HIE Security Utility Services (Silver.5)

The ABB will contain the requirements for security over the sharing and exchange of health information, these requirements are for the HIE and participating systems and end users. The sections covered in this ABB will be Authentication, Authorization and Audit. (See the security section directives in the general directive section)

8 Gold Level Architecture Building Blocks

Gold level architecture building blocks are scoped for the third phase of development of the Reference Architecture.

8.1 Shared Care Records (Gold.1)

The ABB will specify the architectural requirements for the shared care record. The ABB will be an enabler for the Continuum of Care and Clinical Support workstreams and the Phase 2 Shared Care Programmes described in the National Health IT Plan. The ABB will use the IHE Patient Care and Coordination Profiles and infrastructure profiles such as RID.

8.2 Population Based Services (Gold.2)

The ABB will use the IHE Quality, Research and Public Health Technical Framework; the ABB will be an enabler for the Population Health workstream described in the National Health IT Plan.

8.3 Business Rules and Workflow Management (Gold.3)

The ABB will outline sector level clinical pathways and patient journeys from an architectural view point, the ABB will build on various IHE profile workflows and New Zealand specific processes. The ABB will be an enabler for the Patient Administration and Clinical Support workstreams described in the National Health IT Plan.

9 Appendix A: Glossary

The following terms are defined by or used in this document.

Term	Definition	Reference
Actor	Participant people or systems engaged in some transaction; actor is a UML term used in modelling solutions	http://www.uml-diagrams.org/use-case-diagrams.html
Additional Decryption Key (ADK)	Standard encryption technique	http://www.symantec.com/business/support/index?page=content&id=TECH149500
Advanced Encryption Standard (AES)	Standard encryption technique	http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf
Archetype	ISO 13606/openEHR archetypes are constraint-based formal models representing clinical concepts; encoded using Archetype Definition Language; templates collect and further constrain archetypes for specific needs	http://www.openehr.org/wiki/display/stds/openEHR+Archetypes+for+HL7+CDA+Documents
Archetype Definition Language 1.4 (ADL 1.4)	The normative language used to express archetypes	http://www.openehr.org/releases/1.0.2/architecture/am/adl2.pdf
Architecture Building Block (ABB)	Discrete unit of architecture specification; comprises architectural principles, constraints and requirements for some purpose, within the HealthBase framework there are architecture domains; the ABBs form the architectural contents of these domains and are collected under reference architectures	http://www.opengroup.org/togaf/
ASTM International	SDO and creator of the CCR specification	http://www.astm.org/

Audit Trail and Node Authentication (ATNA)	IHE integration profile defining security aspects of the related XDS integration profile; sets requirements for authentication, access control and audit in relation to document sharing	http://wiki.ihe.net/index.php?title=Audit_Trail_and_Node_Authentication
Australian Institute of Health and Welfare (AIHW)	Creator of the METeOR online metadata registry tool	http://www.aihw.gov.au/
CDA Release 2 (CDA R2)	Current release of CDA	http://www.hl7.org/implement/standards/cda.cfm
Clinical Data Repository (CDR)	Database that brings together clinical information from many sources for the purpose of sharing it among care teams	
Clinical Document Architecture (CDA)	HL7 RIM-based specification and XML schema for structured documents	http://www.hl7.org/implement/standards/cda.cfm
Clinical Knowledge Manager (CKM)	openEHR Clinical Knowledge Manager is a web-based tool for creating ISO 13606/openEHR archetypes	http://www.openehr.org/wiki/display/healthmod/Clinical+Knowledge+Manager
Common Terminology Services 2 (CTS2)	HL7/OMG specification for terminology services, e.g. ICD-to-SNOMED CT translation	http://hssp.wikispaces.com/cts2
Continuity of Care Record (CCR)	Widely used international specification that describes summary health status information including problems, medications, alerts, care plan, etc.	http://www.ccrstandard.com/ http://www.astm.org/Standards/E2369.htm
Containment	The term containment is used when a standard is not the current standard. The standard in containment can only be used for specified and specified circumstances; hence the standard is contained for a certain usage.	http://www.opengroup.org/togaf/
Core Health Information	A term used for important patient information, previously referred to as Patient Vitals	

Cross Community Access (XCA)	IHE integration profile that enables the interconnection of XDS affinity domains	http://wiki.ihe.net/index.php?title=Cross-Community_Access
Cross Enterprise Document Media Interchange (XDM)	IHE integration profile for information exchange using portable media such as data sticks	http://wiki.ihe.net/index.php?title=Cross-enterprise_Document_Media_Interchange
Cross Enterprise Document Sharing (XDS)	IHE integration profile for document-oriented health information exchange, based on ebXML	http://wiki.ihe.net/index.php?title=Cross-Enterprise_Document_Sharing
Cross Enterprise Document Sharing-b (XDS.b)	Latest edition of the XDS specification	http://wiki.ihe.net/index.php?title=Cross-Enterprise_Document_Sharing-b_(XDS.b)
Data Service	A service that provides interfaces to the capabilities and data of one or more data resources	
Detailed Clinical Model (DCM)	Conceptual-level specification of the structure and semantics of context-specific clinical information, e.g. adverse reactions, medications	http://www.detailedclinicalmodels.nl/dcm-en
DICOM Key Object Selection (DICOM KOS)	Part of the DICOM standard – describes an image manifest file format	http://medical.nema.org/
Digital Imaging and Communications in Medicine (DICOM)	International standard for communication of medical images	http://medical.nema.org/
Document Digital Signature (DSG)	IHE integration profile for digital signatures	http://wiki.ihe.net/index.php?title=Document_Digital_Signature
Electronic Business Extensible Markup Language (ebXML)	Standard from Oasis and the United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT) for e-business based on XML document exchange; ebXML specifies an infrastructure that allows enterprises to find services, products, business	http://www.ebxml.org/

	processes and documents in a standard way	
Extensible Markup Language (XML)	Markup language commonly used to convey structured information	http://www.w3schools.com/xml/
Extensible Stylesheet Language (XSLT)	XML-based language used to process XML documents, e.g. to create a human readable HTML version of an XML document	http://www.w3schools.com/xsl/
Functional Interoperability	An aspect of interoperability, functional interoperability exists when information exchange conforms to agreed transport protocols and message formats	
Globally Unique Identifier (GUID)	Systematically created, practically unique identifiers in computer software, based on the Universally Unique Identifiers (UUID) standard	http://www.itu.int/ITU-T/asn1/uuid.html
GP2GP	General practice patient notes transfer solution; based on CDA messages and point-to-point messaging	http://www.moh.govt.nz/gp2gp
Graphics Interchange Format (GIF)	Bitmap image file format	http://www.w3.org/Graphics/
Health Information Exchange (HIE)	Application-level communication medium with standardised content and transport, across which participants exchange health information	
Health Level 7 (HL7)	May refer to either Health Level Seven International, the organisation, or its published standards, HL7 v2 and HL7 v3	http://www.hl7.org/
Health Level Seven International	International SDO and creator of the HL7 sets of standards	http://www.hl7.org/ http://www.hl7.org.nz/

Health Practitioner Index (HPI)	National index of all health practitioners and provider organisations, services and facilities operating in New Zealand	http://www.ithealthboard.health.nz/hpi
Health Provider Directory (HPD)	IHE integration profile for the management of and access to shared health provider information	http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Suppl_HPD_Rev1-1_TI_2010-08-10.pdf
HealthBase	Based on TOGAF is the enterprise architecture for the New Zealand health and disability sector The Interoperability Reference Architecture and Health Information Exchange architecture building blocks are all part of HealthBase	http://www.infospace.health.nz/healthbase/
HIE Adapter	Software component providing interfacing support to systems that do not natively use the communication protocols of the HIE	
HL7 Version 2.x (HL7 v2.x)	Widely used health messaging standard	http://www.hl7.org/implement/standards/v2messages.cfm
HL7 Version 3 (HL7 v3)	The successor standard to HL7 v2; encompasses HL7 v3 Reference Information Model, HL7 v3 Messaging, CDA and other related specifications	http://www.hl7.org/implement/standards/v3messages.cfm
Hypertext Markup Language (HTML)	The markup language used to create web pages	http://www.w3schools.com/html/default.asp
IHE Integration Profile	Standards-based specification that describes the actors and transaction types that enable some aspect of interoperability There are foundational integration profiles that are common across clinical domains, and then there are integration profiles specific to particular clinical domains	http://wiki.ihe.net/index.php?title=Profiles

IHE IT Infrastructure Technical Framework	Set of foundational integration profiles, independent of clinical domain	http://wiki.ihe.net/index.php?title=Profiles - IHE_IT_Infrastructure_Profiles
Integrating the Healthcare Enterprise (IHE)	International organisation promoting and providing implementation guidelines for standards-based interoperability	http://www.ihe.net/
International Classification of Diseases for Oncology (ICD 0)	Used principally in tumour or cancer registries for coding the site (topography) and the histology (morphology) from a pathology report	http://www.who.int/classifications/icd/adaptations/oncology/en/
International Classification of Diseases, Australian Modification (ICD 10 AM)	Widely used international classification system for disease identification	http://www.who.int/classifications/icd/en/
International Health Terminology SDO (IHTSDO)	SDO that develops SNOMED CT	http://www.ihtsdo.org/
Interoperability	Interoperability of health information systems, or the ability to exchange information; encompasses functional, semantic and process interoperability	
Interoperability Reference Architecture	Specifically, the Interoperability Reference Architecture created as part of the HealthBase enterprise architecture, to describe interoperability standards in New Zealand	http://www.infospace.health.nz/healthbase/
ISO 13606	Five-part international standard for *Health Informatics –Electronic Health Record Communication*	http://en13606.webs.upv.es/web13606/index.php/the-ceniso-en13606-standard
ISO 13606/openEHR Reference Model (RM)	Set of technical building blocks which archetypes bind and constrain to express a particular clinical concept; the reference model consists of a set of UML classes	http://www.openehr.org/

Logical Observation Identifiers Names and Codes (LOINC)	Coding system for identifying laboratory test and clinical observations The coding system NZPOCS is based upon	http://loinc.org/
Message Transmission Optimisation Mechanism / XML-binary Optimised Packaging (MTOM/XOP)	MTOM is a method of including binary data in calls to SOAP web services XOP is a means for including binary data within XML documents MTOM/XOP is use of the two together	http://www.w3.org/TR/soap12-mtom/ http://www.w3.org/TR/xop10/
METeOR	Online metadata registration tool, conforming to ISO/IEC 11179; provided by AIHW and licenced for use in New Zealand	http://meteor.aihw.gov.au/content/index.phtml/itemId/181162
Multipurpose Internet Mail Extensions (MIME)	Internet content type family; formats for the transmission of text, images, audio, video, etc.	http://www.w3.org/Protocols/rfc1341/7_2_Multipart.html
National Health Index (NHI)	New Zealand's national master patient index; an NHI number identifies every health consumer in the country	http://www.nzhis.govt.nz/moh.nsf/pagesns/266?Open
Network Time Protocol (NTP)	Protocol used to synchronise clocks over the Internet	http://www.ntp.org/
New Zealand Medicines Terminology (NZMT)	SNOMED CT Reference Set for medicines used in New Zealand	http://www.nzulm.org.nz/
New Zealand Pathology Observation Code Sets (NZPOCS)	New Zealand specific code sets for ordering and reporting laboratory tests Based on LOINC	http://www.ithealthboard.health.nz/nzpocs
Object Identifier (OID)	Hierarchically generated persistent object identifier	http://www.oid-info.com/ - oid
Object Management Group	An IT industry SDO	http://www.omg.org/

(OMG)		
Patient Administration System (PAS)	Hospital information system, responsible for booking and scheduling patients and resources	
Patient Demographic Query V3 (PDQV3)	IHE integration profile, using HL7 v3 message formats and SOAP web services transport, for distributed access to patient demographic data from an authoritative source	http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Suppl_PIX_PDQ_HL7v3_Rev2-1_TI_2010-08-10.pdf
Patient Identity Cross Reference V3 (PIXV3)	IHE integration profile, using HL7 v3 message formats and SOAP web services transport, for distributed access to patient identity feeds from an authoritative source	http://www.ihe.net/Technical_Framework/upload/IHE_ITI_Suppl_PIX_PDQ_HL7v3_Rev2-1_TI_2010-08-10.pdf
Picture Archiving and Communication System (PACS)	Class of system for storing and providing access to medical images, particularly DICOM images	
Point of Care (PoC) Point of Service (PoS)	Point-of-care or, more generally, point-of-service systems are those used proximate to the patient receiving care	
Portable Document Format (PDF)	Standard file format for unstructured document exchange	http://www.adobe.com/products/acrobat/adobepdf.html
Portable Network Graphics (PNG)	Lossless, portable bitmap image file format	http://www.w3.org/Graphics/
Practice Management System (PMS)	General practice, primary care or departmental practice or patient management system, with both administrative and clinical functions	
Process Interoperability	An aspect of interoperability, process interoperability exists when business processes are supported across multiple information systems	

Record Locator Service (RLS)	Index-based tool provided by the Health Information Exchange enabling participant applications to find patient records distributed across multiple repositories	http://healthinformatics.wikispaces.com/Record+Locator+Service+(RLS)
Reference Architecture	Collection of related architecture building blocks in some domain, a guide and template for solution architectures in that aspect	http://www.opengroup.org/togaf/
Reference Information Model (RIM)	HL7 v3 RIM is a health meta-model for the development (by constraint) of domain-specific models and message formats RIM sometimes also refers to similar non-HL7 models	http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.145.4676&rep=rep1&type=pdf
Regional CDR (R-CDR)	R-CDR are described by the National Health IT Plan as the regional collection points for objective health information, including laboratory test results, medication records, etc. – to facilitate information sharing. R-CDR will be made up of multiple repositories holding clinical information for a Region	http://www.ithealthboard.health.nz/content/national-health-it-plan
Registry-Repository Model	Information sharing architecture in which a central registry serves as an index to documents stored in multiple distributed repositories	
Representational State Transfer (REST)	Style of web services based on native use of HTTP; alternative to and simpler than SOAP web services	https://www.ibm.com/developerworks/webservices/library/ws-restful/
Schematron	XPath-based assertion language for validating XML documents	http://www.schematron.com/
Semantic Interoperability	An aspect of interoperability, semantic interoperability exists when information exchange involves commonly understood data structures and terminologies	

Service Adapters	Service adapters are software components that attach to legacy applications in order to make web services accessible to them	
Service Oriented Architecture (SOA)	Architectural philosophy or style for delivery of functionality as sets of discrete, interoperable components	
Simple Object Access Protocol (SOAP)	Method of invoking web services based on remote procedure calls and structured XML payloads	
SNOMED CT Reference Set	Collection of related SNOMED CT concepts and terms pertaining to some domain, e.g. New Zealand Medicines Terminology	http://www.ihtsdo.org/fileadmin/user_upload/Docs_01/Technical_Docs/reference_sets.pdf
Standards Development Organisation (SDO)	Usually non-profit body that exists to create standards – e.g. in health informatics, HL7, IHTSDO; in the IT industry, Oasis, OMG	
Systematised Nomenclature of Medicine Clinical Terms (SNOMED CT)	Universal medical terminology system	http://www.ihtsdo.org/
The Open Group Architecture Framework (TOGAF)	Industry standard architecture framework, see HealthBase	http://www.opengroup.org/togaf/
Transport Layer Security (TLS)	Cryptographic protocol that enables secure channels of communication over the Internet	http://www.ietf.org/rfc/rfc2246.txt
Unified Code for Units of Measure (UCUM)	Universal coding system for units of measure	http://unitsofmeasure.org/

Unified Modelling Language (UML)	Standard business process and information modelling language developed by OMG	http://www.omg.org/technology/readingroom/UML.htm
Universally Unique Identifiers (UUID)	International standard for OIDs	http://www.itu.int/ITU-T/asn1/uuid.html
Utility Services	SOA term that refers to services encapsulating common, reusable crosscutting (business process-independent) functionality	
Web Access to DICOM Persistent Objects (WADO)	HTTP-based protocol for retrieval of DICOM objects (e.g. stored in PACS), either in native form or as a rendered image	ftp://medical.nema.org/medical/dicom/2009/09_18pu.pdf
Web Services Interoperability Basic Profile (WS-I Basic Profile)	Set of consistent web services specifications, collected and profiled for use in interoperability	http://ws-i.org/Profiles/BasicProfile-2.0-2010-11-09.html
XDS Affinity Domain	Group of healthcare enterprises that have agreed to share health information under a common set of policies and with common infrastructure; a concept of the XDS integration profile	http://wiki.ihe.net/index.php?title=Cross_Enterprise_Document_Sharing
XML Schema (XSD)	XML-based language for defining XML document structure	http://www.w3schools.com/schema/default.asp
XPath	XML-based language for querying XML documents	http://www.w3schools.com/xpath/default.asp

10 Appendix B: Architecture Landscape

The architecture landscape diagram below shows the intended future state for sector services. The blue coloured section at the top of the diagram shows national services and repositories. The green coloured section in the middle shows regional services and repositories. The yellow coloured section at the bottom shows local services, this includes provider and personal services. The services can be federated, for example: national services being federated to regional level.

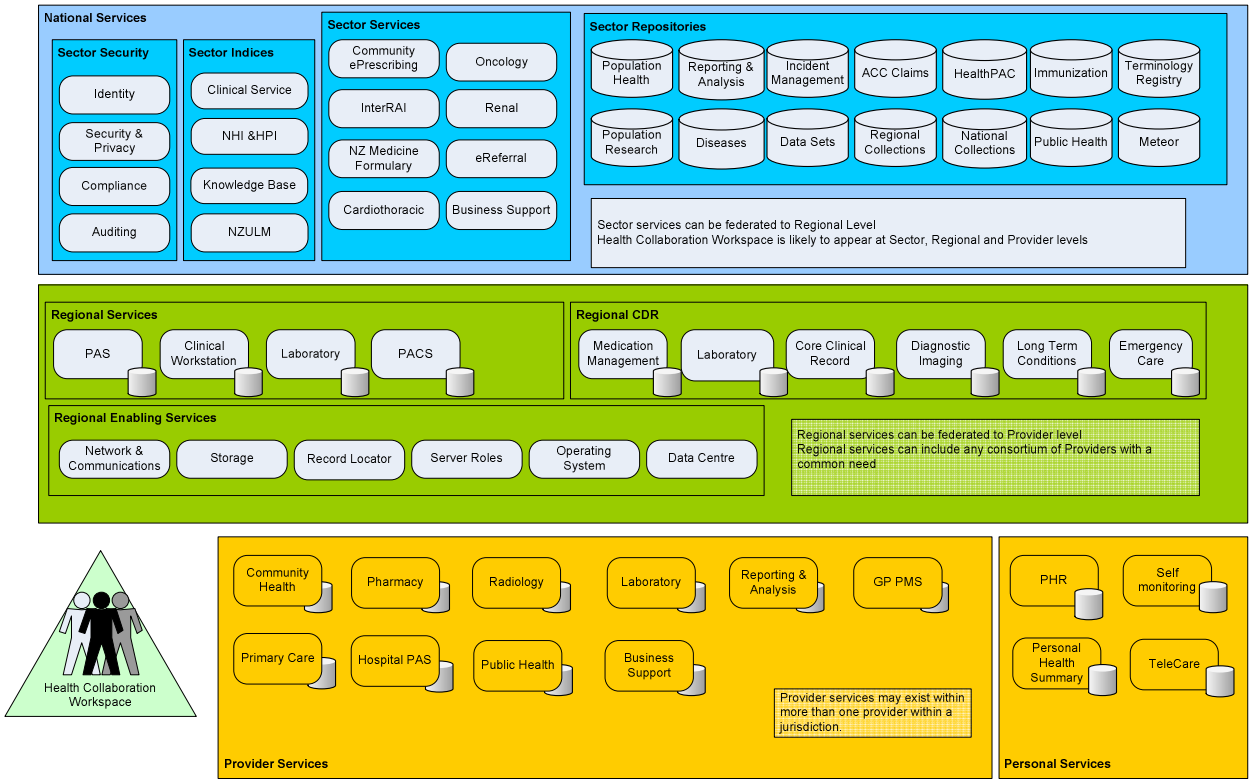

Figure 18 – Architecture landscape

10.1 National Services

National services are any services, indices or information collection endpoints that need to be presented to the sector level as one service.

10.1.1 Sector Security

Sector authentication would provide a single logon service with auditing facilities to ensure compliance and privacy. This service would use the Sector indices, HPI service for the authentication of health sector workers and NHI service for authentication of personal services.

10.1.2 Sector Indices

There are a number of sector-wide indices, such as:

- Identity service will provide HPI, NHI and postal address functions
- National Medical Warning System (NMWS)

- Clinical Service is an index that provides a list of all clinical services available from place to place - e.g. Waikato Diabetes Service, Wellington Cardiology Service. Service type and location would be provided at a minimum.
- Knowledge Base, refers to a knowledge base such as SNOMED Clinical Terms, allowing end users to look up and properly describe/code clinical conditions, e.g. determine the right coding for an instance of heart attack.
- Medicine List is the New Zealand Universal List of Medicines (NZULM). The NZULM is a single medicines list that will be universally used throughout the health system. The NZULM uses the SNOMED CT-based New Zealand Medicines Terminology (NZMT).

10.1.3 Sector Services

Business services are presented to the sector as a single, nationally available, technical service instance. These technical services provide the business with a holistic view of the particular business process being performed, e.g. interRAI provides multiple business services such as aged care assessment within a single technical service.

10.1.4 Sector Repositories

The sector repositories are used for a wide range of health related sector activities such as population base health surveillance, sector billing claims, terminology and liaison between health organisations and clinical research.

10.2 Regional Services

These are services that are deployed to a single region (rather than national deployments). The services can communicate with each other using an HIE, and can participate with registries (national or regional) following the IHE XDS model for discovery.

10.2.1 Regional Services

These are regional services that will be supported by regional systems such as laboratory results repositories, PACS systems and PAS systems.

10.2.2 Regional Clinical Data Repository

R-CDRs hold clinical information to enable its sharing between providers and the patient.

10.2.3 Regional Enabling Services

These services provide the infrastructure to enable regional business services. Enabling services are specifically described in this document, as they do not currently exist in a useable form.

Record locator services would be positioned at regional level. The federated regional record locator network could provide a virtual national EHR. Record locator services encapsulate the set of functions that enables systems to locate individual patient records across the federated system – e.g. patient X has a radiology report held in system Y in region Z.

10.3 Local Services

Local services are those accessible at district, organisation or facility levels, and include provider and personal health information services.

10.3.1 Provider Services

Provider services are those that support clinicians at the point of care.

10.3.2 Personal Services

Personal services are those used directly by patients and other consumers. Note that these services may actually be supplied regionally.

10.3.3 Health Collaboration Workspace

This workspace will be assessable by many different portal methods and locations, for example a patient's home or an Emergency department. The workspace is role based and will provide a user with applications that are appropriate for the designated user role. For example a role type could be a GP; the portal would provide the GP PMS system and any other application the GP would require for the role, including EHR repository search and access.

One of the main themes of the National Health IT Plan is the involvement of the patient in their own healthcare. This is explicit in the vision statement, and implicit in the 'Shared Care Plan', which is seen as one of the major outcomes.

It is expected that this functionality will be supplied by a number of patient portals – of which one already exists in New Zealand.

From the perspective of this Reference Architecture, a portal providing information and functionality to a patient is no different to any other consumer of health information – the 'patient portal' will use a services based approached to access functionality from other systems as required, and subject to the appropriate security and privacy constraints.

It is likely that these workspaces will follow the model being proposed by NEHTA in the Personally Controlled EHR (PCEHR) project, where the patient will generally control access to the information within the repository and is aware of all people who are accessing the data (http://www.nehta.gov.au/ehealth-implementation/pcehr-concept-of-operations).

11 Appendix C: Services Approach

Common practice today is to model business functionality in the form of re-usable services that can defined as specific functionality that be invoked using defined interfaces that are implementation agnostic – they can be used by any technology (e.g. JEE, .NET, PHP).

Examples of possible services are:

- Get a patient's demographic information
- Order a lab test
- Retrieve a patient's health status information

The advantage of a services based approach is that it avoids duplication, ensures that common functions are performed in a consistent way, and allow multiple vendors to provide implementations of a particular service allowing a consumer to choose to choose the most appropriate implementation for their needs – plug and play.

References:

- Wikipedia: http://en.wikipedia.org/wiki/Service-oriented_architecture
- HSSP practical guide: http://hssp.wikispaces.com/PracticalGuide
- Software Engineering Institute (Carnegie-Mellon) whitepaper www.sei.cmu.edu/library/assets/whitepapers/SOA_Terminology.pdf

The Reference Architecture uses a services approach to exchange information. These services will be offered and consumed at sector, regional and local levels. It is recognized that applications may not natively support the data services, and the use of HIE adapters will be used to deal with this requirement. It is expected that over time applications will natively support the standard data service interfaces, and the use of adapters will become less prevalent.

The interoperability problem that this reference architecture addresses is really about the sharing, reuse, and movement of data between business units, which can be described as 'data in motion'. In this definition, data is the content and context that is captured, manipulated and consumed by applications. This approach to integration is where data services are used to provide a solution.

The benefits of providing standardized data services are the reduced time to change applications and the cost savings that can be achieved. These savings are achieved by the reuse of the services, which leads to simplifying the current levels of point-to-point integration and reducing the time and effort savings when applications are changed or replaced. The services remove the dependencies between applications, providing certified applications with a plug and play environment.

This approach will enable the working interoperability to support the National Health IT Plan vision.

11.1 Services Taxonomy

The definition of the services required to support the different business objectives within the health sector will require significant analysis work, and will result in the creation of a services taxonomy that will group like services together. The services making up the services taxonomy will be segmented into three layers, as described later in this section. There are many references guides available to performing this work. Types of services include, for example:

- Data services (the building blocks for business services): create, retrieve, update and delete services (CRUD)

- Location services (including record locator services)
- Business services: business entity services, business process services
- Security services
- Privacy services
- Search services
- Utility services: terminology services, anonymization services
- Orchestration services
- Communication services
- Terminology services

11.2 Service Taxonomy Layers

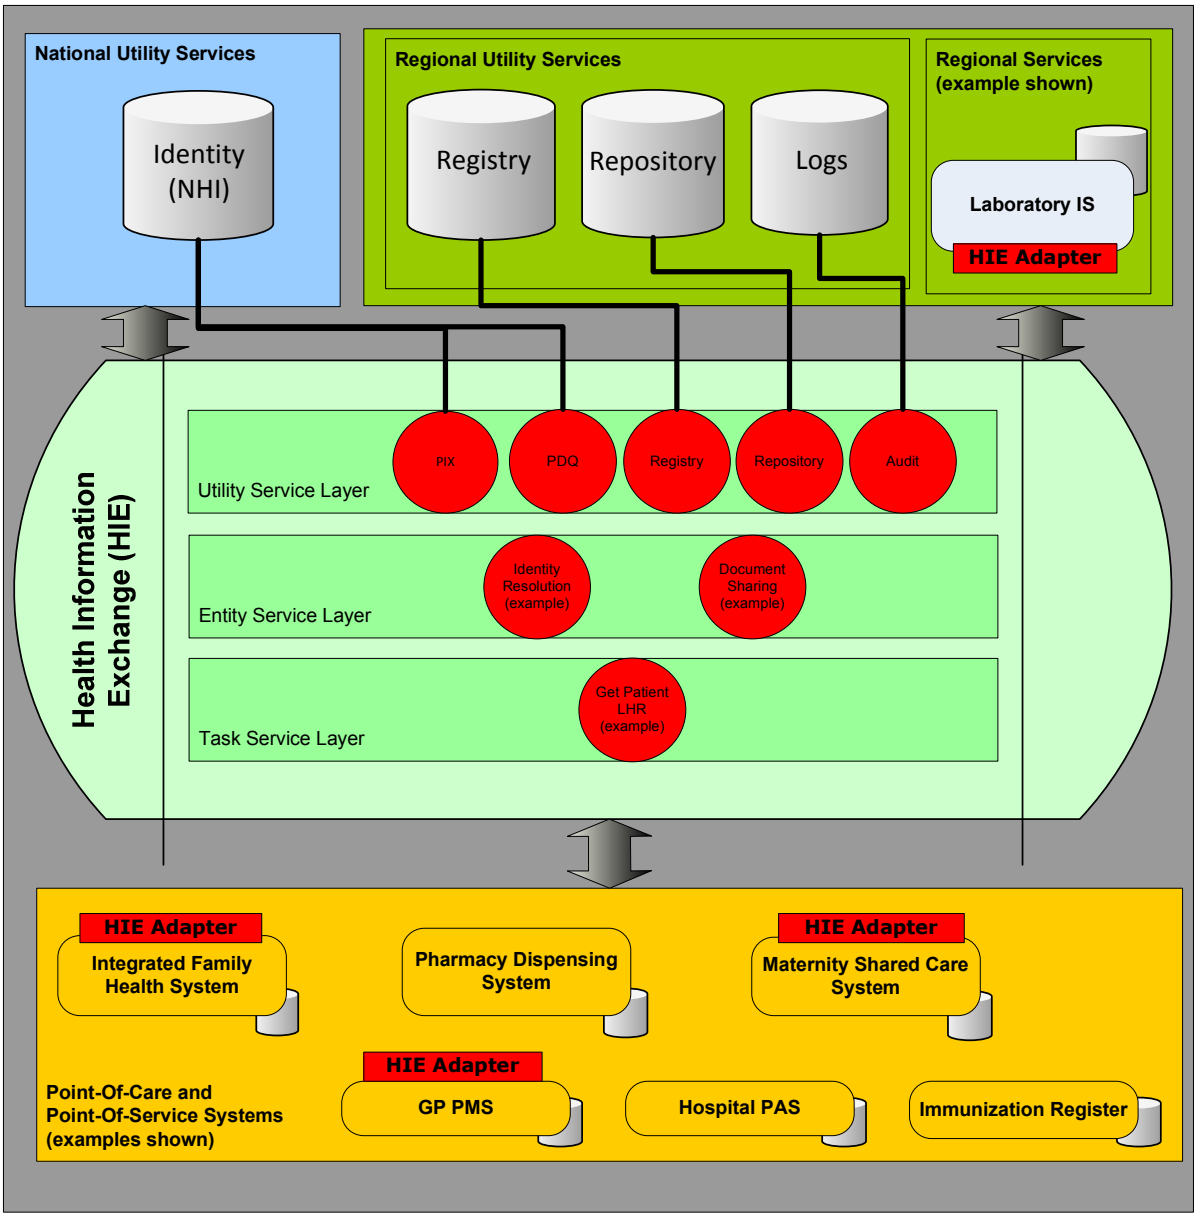

Figure 19 – Services Taxonomy Layers showing bronze utility services

11.2.1 Utility Service Layer

Utility services encapsulate common, crosscutting functionality that is useful in many contexts but is not derived from the business architecture. They are also highly reusable services due to minimal dependencies on business as well as application context. Examples of typical utility services include notification, logging, and messaging.

11.2.2 Bronze Level Utility Services

XDS.b utility services are required to enable the R-CDR component of the shared care model described in the National IT Plan and are considered as bronze level services. The services represent the XDS.b mandatory services for R-CDRs.

11.2.3 Entity Service Layer

Entity services are derived from one or more related business entities. They are considered highly reusable because they minimize dependencies to parent business processes. Examples of healthcare-specific business entities include patient, lab order and medical summary. Example services are shown in the diagram.

11.2.4 Task Service Layer

Task services are based on a specific business process, and typically act as an entry point and controller for a service composition. As a result, task services generally have less reuse potential than the other services types. An example of a task service is a *RunAuditReport* service that retrieves, aggregates and displays audit record details for a clinical system.

11.3 Standard Services and Interfaces

Within the health sector, the HSSP project (Healthcare Services Specification Project (http://hssp.wikispaces.com/) brings together experts from the HL7 and the OMG groups to help defined standard services to be used in healthcare. These services tend to be very high level rather than granular such as:

- **RLUS** – Retrieve, Locate Update Service defining a high level interface to be used when locating and updating entities of any type (like a patient, a lab result or a medication)
- **CTS2** – Common Terminology Services for retrieving terminology codes (like SNOMED or ICD) from national/regional indices

HSSP objectives include:

- To stimulate the adoption and use of standardized plug-and-play services by healthcare software product vendors
- To facilitate the development of a set of implementable interface standards supporting agreed-upon services specifications to form the basis for provider purchasing and procurement decisions.
- To complement and not conflict with existing HL7 work products and activities, leveraging content and lessons learned from elsewhere within the organization.

HSSP provide a number of artefacts to assist with defining the required services, one of the more useful being the matrix of Business line / Service shown below (from the HSSP Practical Guide to SOA at http://hssp.wikispaces.com/PracticalGuide).

	Healthcare Unique Services							Business Services				SOA Infrastructure Services				
Business Line	Order Entry	Order Fulfillment	Patient Evaluation (DSS)	Laboratory Data Retrieval	Pharmacy Data Retrieval	EHR	Alert/Event Mgmt	Master Person Index	Terminology Service	Demographics	Billing	Scheduling	Auditing Service	Exception Mgmt	Business Process Mgmt (BPM)	Authentication	Services Directory
Pharmacy	X	X	X	·	X	X	X	X	X	X	X	·	X	X	X	X	·
Laboratory	X	X	X	X	·	X	X	X	X	X	X	X	X	X	X	X	·
Patient Administration	·	·	·	·	·	·	X	X	·	X	X	·	X	X	·	X	X
Order Entry/Mgmt	X	X	·	·	·	X	X	X	X	X	X	X	X	X	X	X	X
Scheduling	·	·	·	·	·	·	X	X	·	X	·	X	X	X	X	X	X
Registration	·	·	·	·	·	·	X	X	·	X	X	X	X	X	X	X	·
Care Management	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Referrals/Referral Mgmt	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Nursing	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
Emergency Department	X	X	X	X	X	X	X	X	X	X	X	X	X	X	·	X	X
Patient Billing	·	X	·	X	X	X	X	X		X	X	·	X	X	X	X	·
Imaging/Radiology	·	X	·	·	·	X	X	X	X	X	X	X	X	X	X	X	X
Clinical Decision Support	·	·	X	X	X	X	X		X	X	·	·	X	X	·	X	·
Facilities Management	·	·	·	·	·	·	X	·	·	·	·	X	X	X	·	X	X
Nutrition Mgmt (Dietetics)	X	X	·	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Figure 20 – HSSP matrix showing business lines versus services

12 Appendix D: Health Information Exchange

The term Health Information Exchange (HIE) describes the capability for applications and systems to share information in a manner consistent with this Reference Architecture.

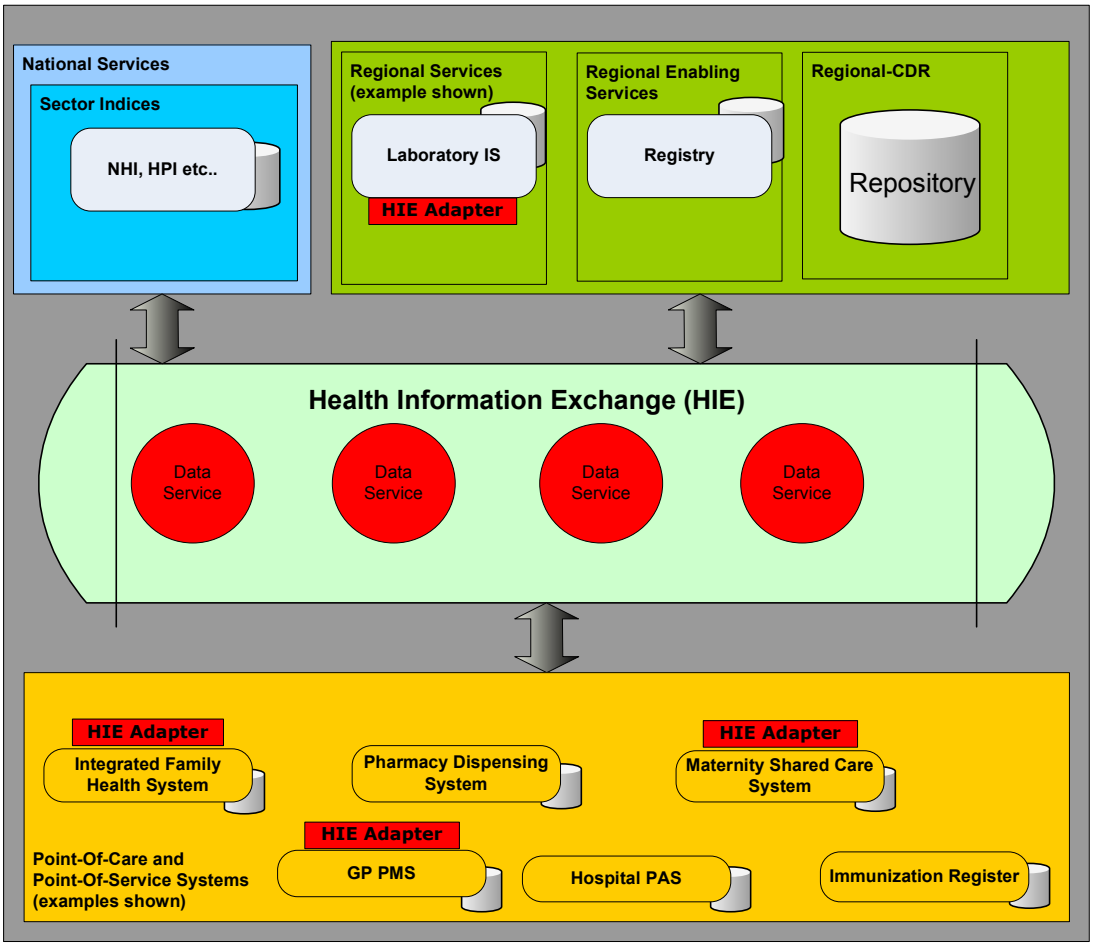

Figure 21 – Health Information Exchange

The diagram shows the interrelationships of the following key enablers of working interoperability. (Note that the systems shown are examples only.)

- Regional CDRs
- Data services
- Health Information Exchange (HIE)
- HIE adapters
- HIE transport

12.1 Regional CDRs

R-CDRs are central to working interoperability. They contain objective clinical information about the individual, such as details of problems, results, medications and encounters. Shared care applications – which may be either regional or national or specialized – will use data services to access summary health records and care plans held in R-CDRs.

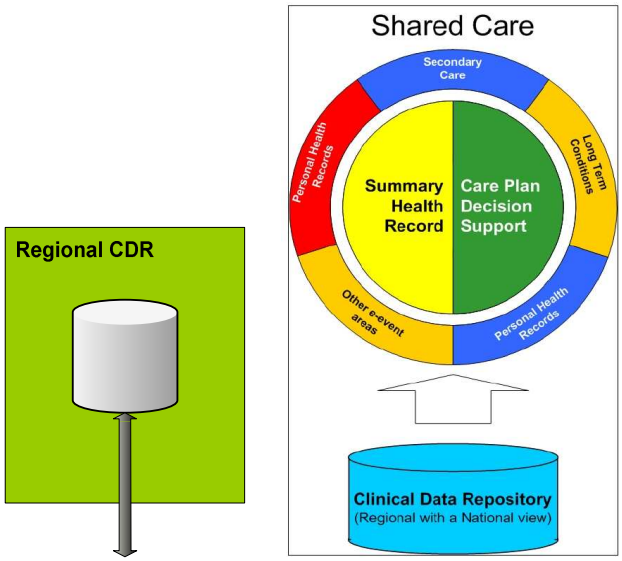

Figure 22 – R-CDR and Shared Care Model

R-CDRs will follow a registry-repository model; this will support a federated approach, allowing that national systems can be frontline repositories. This approach improves data quality by preserving the authoritative data source. The use of the XDS.b registry will ensure fast response times of patient information and provide granular security of the information.

12.2 Data Services

A data service is a service that provides interfaces to the capabilities and data of one or more data resources within a service-oriented architecture (SOA).

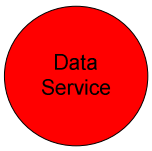

A set of standardized data services forms the interface to R-CDRs and to other systems in the regional ecosystem. Data services will be implemented by web services, although the SOA approach lends itself to other interfaces. Data services may have workflow functionality. The data services will conform to standards described in this document and will be used by certified applications or certified application adapters.

Date services can be implemented using either SOAP over HTTP, REST or DICOM for image data.

There are two actors:

- A **service provider** exposes functionality that can be consumed by other service applications. The functionality provided would be create, read, update and delete (CRUD) type access, but may include other workflow elements if required by the definition of that service.
- A **service consumer** is any application that requires information from a service provider.

12.2.1 Addressing Schema

Data services will be located by performing a lookup or discovery function using a consistent addressing schema. The following standards will be used:

Electronic Business using eXtensible Markup Language (ebXML) Registry

The primary focus of ebXML Registry extends beyond discovery into collaboration. This can be viewed on two levels: development collaboration and run-time collaboration. Due to its focus on storing and maintaining XML artefacts, an ebXML registry can enable both collaborative development of XML artefacts within an organization and run-time collaboration between trading partners. For example, users can create XML artefacts and submit them to an ebXML registry for use and potential enhancement by other users. Additionally, once trading partners have discovered each other using the discovery mechanisms defined as part of the ebXML framework (which involve CPPs and CPAs), they can collaborate in data exchange scenarios using the XML artefacts that are registered (and potentially stored) in the ebXML registry. The parties can also conduct business scenarios according to discovered business process specifications. The ebXML registry is also intended to store and manage various artefacts that support business collaboration.

The ebXML registry will accommodate the registration of business and Data Service information. The following three protocols will be used to achieve Data Service discovery. An ebXML Registry is also designed to accommodate additional types of content such as schemas, DTDs, and XML documents.

Collaboration Protocol Profile (CPP):

Describes the message-exchange capabilities of a Party involved in a business collaboration; also used for trading partner discovery purposes.

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ebxml-cppa

Collaboration Protocol Agreement (CPA)

CPA defines the capabilities that two parties need to agree upon to enable them to engage in business collaboration.

Business Process Specification Schema (BPSS)

BPSS provides a standard framework by which business systems may be configured to support the execution of business collaborations consisting of business transactions.

http://www.ebxml.eu.org/process.htm

Domain Name System (DNS)

DNS is a hierarchical naming system built on a distributed database that translates domain names to numerical identifiers.

http://en.wikipedia.org/wiki/Domain_Name_System

Web Services Description Language (WSDL)

WSDL is an XML-based standard specification for describing web services. The data service interfaces will be described in WSDL documents. WSDL defines an XML format for describing service endpoints that operate on messages that contain either document-oriented or procedure-oriented information.

WSDL service descriptions are published in an ebXML registry. WSDL documents have two main parts:

- The service interface definition describes the abstract type interface and its protocol binding, known as the WSDL binding document. A service interface is described by a WSDL document that contains the types, import, message, port type, and binding elements. A service interface contains the WSDL service definition that will be used to implement one or

more services. It is an abstract definition of a web service, and is used to describe a specific type of service. This document can reference another service interface document using an import element.

- The service implementation definition describes the service access location information, known as the WSDL service document. The service implementation document contains the service elements. A service implementation document contains a description of a service that implements a service interface. A service implementation document can contain references to more than one service interface document.

A service provider hosts a data service and makes it accessible using protocols such as SOAP over HTTP. The web service is described by the WSDL documents that are stored in the ebXML repository.

http://www.w3.org/TR/wsdl

System Directory for Document Sharing (SDDS)

Currently, the XDS.b and XCA profiles do not specify how the participants in the document exchange know about the existence and web services end points of the document repositories, the document registry, and the various gateways. An XDS Document Source needs to know the web services end points corresponding to all possible Document Repository Unique IDs that are available in the XDS affinity domain. Similarly, an XCA Initiating Gateway needs to know existence and addresses of the appropriate XCA Responding Gateways. This proposed supplement will enhance the XCA and XDS.b profiles with options to satisfy these needs.

12.3 Health Information Exchange

The HIE represents the electronic movement of health-related information among organizations according to nationally recognized standards.

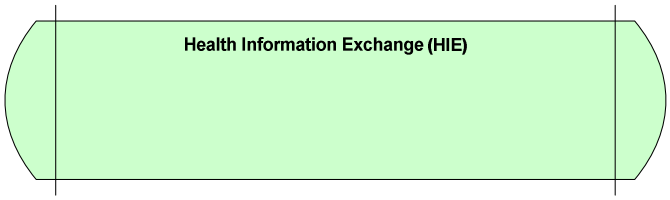

Data services are provided to their consumers via the HIE. The HIE is federated at national, regional and local levels. This provides flexibility for organizations to install and configure their own instance, or to use another's.

The HIE can be thought of as a fabric that uses standardized data services to provide the movement of data for the health system. The HIE is a logical concept and can be implemented in a number of different ways:

- Deployed in the space between systems – for example a regionally available Enterprise Service Bus (ESB) (method 2 below)
- A separate component within an organization – for example a hospital could provide a service for appointment lookup, via a local integration engine (method 4 below)
- An application exposing a data service – for example a web service that provides a laboratory result (method 1 below)

The HIE is made up of multiple middleware instances, which may include enterprise service buses, integration engines and web services depending on the nature of the data service provider and the complexity of the environment. These instances will be interconnected forming a network of middleware data services. This configuration will provide known boundaries, ensuring local flexibility

with localized security and auditing capabilities. This layer of separation between supplier and consumer of services also allows:

- Services taxonomy to be implemented at the HIE level
- Different service providers can provide the same service
- Workflow can be used within the HIE fabric

12.3.1 HIE Deployment Method 1 – Dedicated

In this example, the HIE is dedicated to a laboratory provider application and is available directly to consumer applications. A GP PMS consumer application connects directly to an external HIE, coupled to the laboratory provider application.

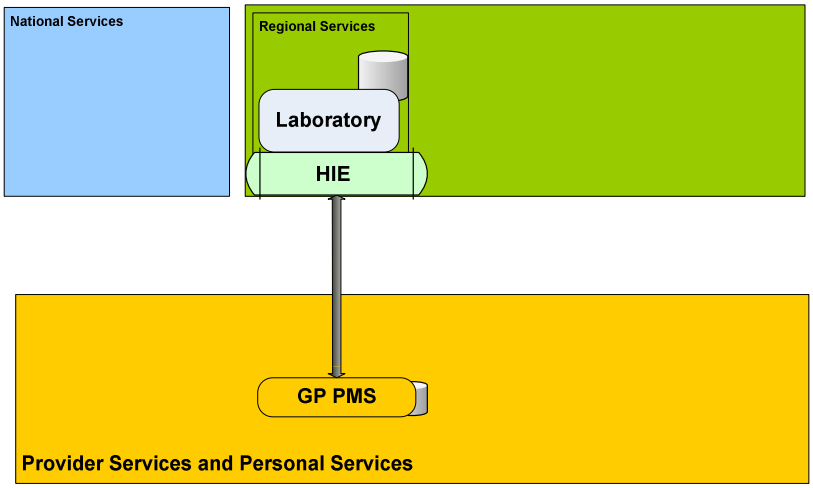

Figure 23 – HIE Deployment Method 1 (Dedicated)

12.3.2 HIE Deployment Method 2 – External

A point-of-service application connects to an externally hosted HIE.

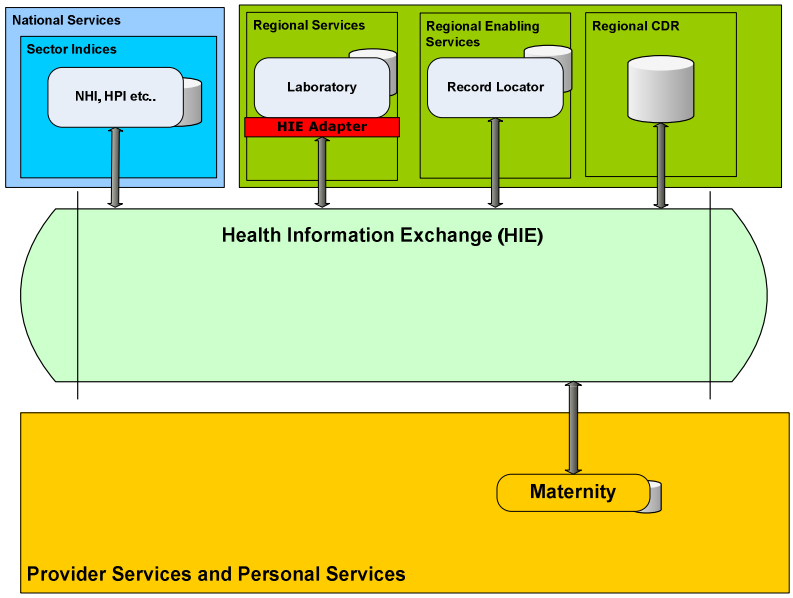

Figure 24 – HIE Deployment Method 2 (External)

12.3.3 HIE Deployment Method 3 – Internal

A hospital PAS consumer application connects directly to a local HIE.

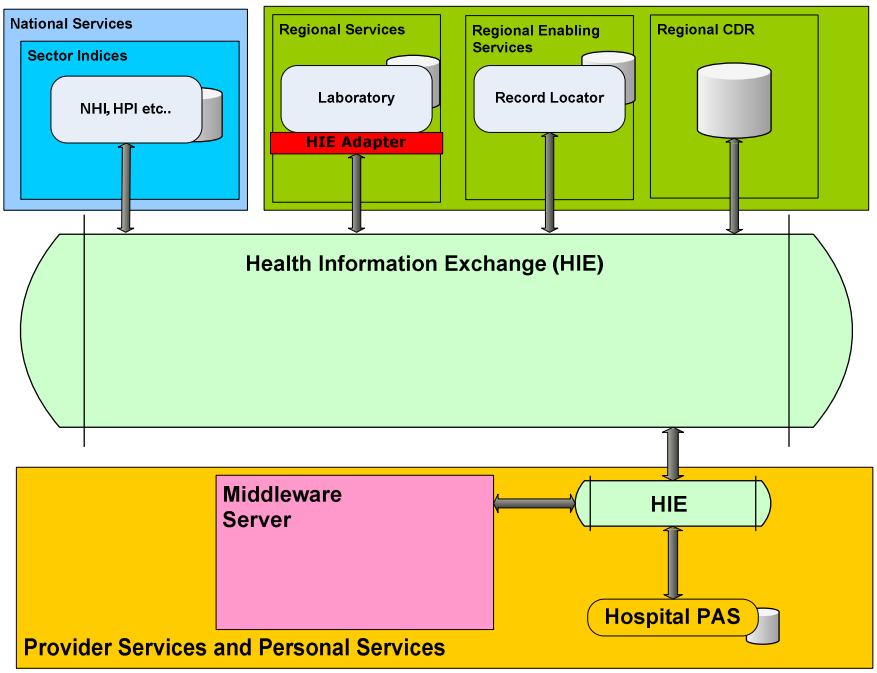

Figure 25 – HIE Deployment Method 3 (Internal)

12.3.4 HIE Deployment Method 4 – Virtualized

A public health system connects to a HIE that has been virtualized within a middleware server.

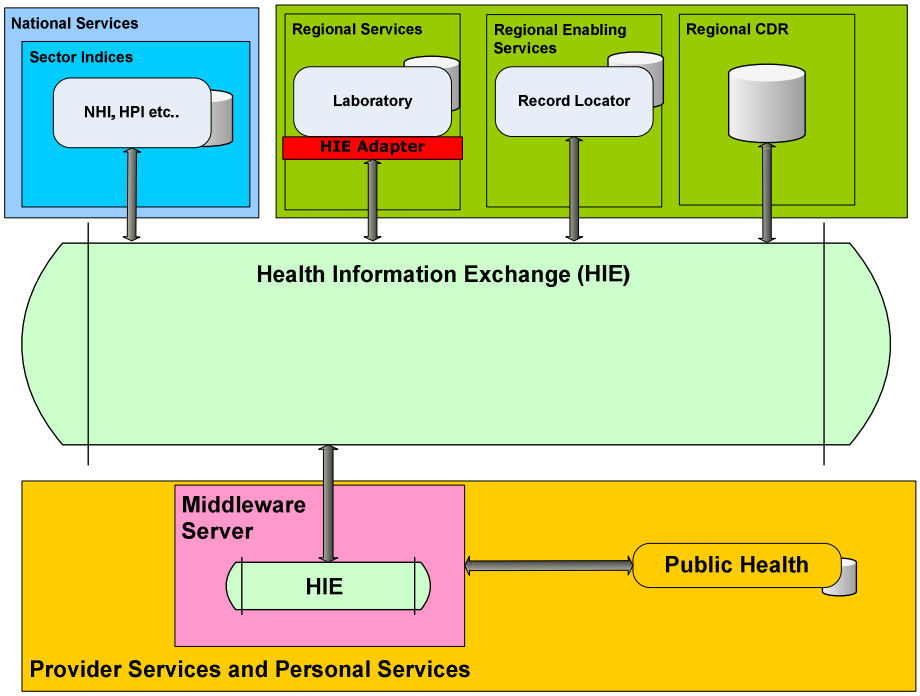

Figure 26 – HIE Deployment Method 4 (Virtualized)

12.4 HIE Adapters

HIE Adapters provide interfacing support to systems that cannot natively use the standardized data services provided by the Health Information Exchange (HIE). It is expected that as working interoperability matures the need for adapters will diminish, implying that over time all applications will natively support the HIE.

12.4.1 HIE Adapter Deployment Method 1 – Consumer Application

The consumer application adapter converts the application requests to the native HIE standard. In the example the Community Health application interface would communicate directly with the adapter and the adapter would communicate natively with the HIE, enabling the consumer to natively use the HIE.

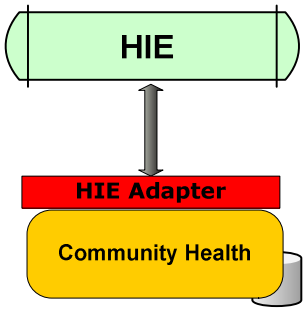

Figure 27 – HIE Adapter Deployment Method 1 (Consumer Application)

12.4.2 HIE Adapter Deployment Method 2 – Provider Application

The provider application adapter converts native HIE requests to suit the provider application or repository interface. In the example, the HIE requests are converted to meet the laboratory application's interface requirements.

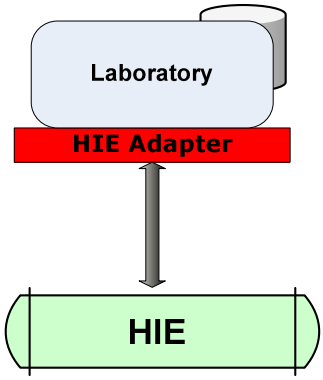

Figure 28 – HIE Adapter Deployment Method 2 (Provider Application)

12.4.3 HIE Adapter Deployment Method 3 – Middleware

There will be situations where a provider application, consumer application or repository would not be the most suitable place for the adapter. The adapter could then be deployed to a middleware server; this method will allow transparent interoperability for the application or repository.

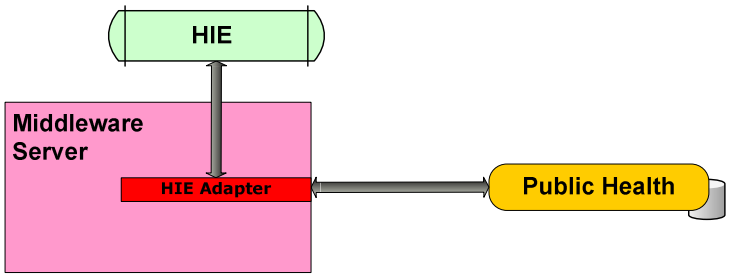

Figure 29 – HIE Adapter Deployment Method 3 (Middleware)

12.5 HIE Transport

There two approaches for interfacing with data services:

- **SOAP** (Simple Object Access Protocol – originally, at least)
- **REST** (Representational State Transfer)

Both approaches work, both have advantages and disadvantages to interfacing to web services, but it is up to the web developer to make the decision of which approach may be best for each particular case.

12.5.1 SOAP

SOAP 1.2 has fixed many of the perceived shortcomings of the technology and pushing it to new levels of both adoption and ease-of-use.

Note that using SOAP 1.2 has some additional overhead that is not found in the REST approach, but that overhead also has advantages. First, SOAP relies on XML in three ways; the envelope – that defines what is in the message and how to process it, a set of encoding rules for data types, and finally the layout of the procedure calls and responses gathered. This envelope is sent via a transport (HTTP/HTTPS), and an RPC (remote procedure call) is executed and the envelope is returned with information in a XML formatted document.

12.5.2 REST

REST embraces a stateless client-server architecture in which the web services are viewed as resources and can be identified by their URLs. Web service clients that want to use these resources access a particular representation by transferring application content using a small globally defined set of remote methods that describe the action to be performed on the resource. REST is an analytical description of the existing web architecture, and thus the interplay between the style and the underlying HTTP protocol appears seamless. The HTTP methods such as GET and POST are the verbs that the developer can use to describe the necessary create, read, update, and delete (CRUD) actions to be performed.

The REST approach, which uses a standard URI (Uniform Resource Identifier) that makes a call to a web service like http/https://www.mycompany.com/program/method?Parameters=xx.

Both technologies can be used together. REST is very easy to understand and is extremely approachable, but does lack agreed standards and is considered an architectural approach. In comparison, SOAP is an industry standard with a well-defined protocol and a set of well-established rules to be implemented, and it has been used in systems both big and small.

So this means areas that REST works really well for are:

- **Limited bandwidth and resources**; remember the return structure is really in any format (developer defined). Plus, any browser can be used because the REST approach uses the standard GET, PUT, POST, and DELETE verbs. Again, remember that REST can also use the XMLHttpRequest object that most modern browsers support today, which adds an extra bonus of AJAX.
- **Totally stateless operations;** if an operation needs to be continued, then REST is not the best approach and SOAP may fit it better. However, if you need stateless CRUD (Create, Read, Update, and Delete) operations, then REST is it.
- **Caching situations;** if the information can be cached because of the totally stateless operation of the REST approach, this is perfect.

SOAP is mature and well defined and does come with a complete specification. The REST approach is just that, an approach and is wide open for development, so if you have the following then SOAP is a great solution:

- **Asynchronous processing and invocation**; if your application needs a guaranteed level of reliability and security then SOAP 1.2 offers additional standards to ensure this type of operation. Things like WSRM – WS-Reliable Messaging.
- **Formal contracts**; if both sides (provider and consumer) have to agree on the exchange format then SOAP 1.2 gives the rigid specifications for this type of interaction.
- **Stateful operations;** if the application needs contextual information and conversational state management then SOAP 1.2 has the additional specification in the WS* structure to support those things (Security, Transactions, Coordination, etc). Comparatively, the REST approach would make the developers build this custom plumbing.

References:

- http://wiki.hl7.org/index.php?title=EHR
- http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=46087

12.6 Sample Storyboard

The following storyboard gives an example of a community pharmacy system using services exposed by the HIE to support medication dispensing. In this storyboard, there is a regionally hosted HIE that the pharmacy system can access via a local adapter.

Mr Iamsick attends a pharmacy to pick up a prescription sent electronically from an after-hours medical service. He gives the pharmacist permission to access his online records. The pharmacy system calls the record locator service via the HIE which discloses the presence of two allergy records in CDRs – one in the Northern R-CDR where the patient currently lives, and the other in the Southern R-CDR relating to a reaction to an antibiotic while on holiday in Queenstown the previous year.

The pharmacy system uses the *getPrescribedMedication* method of the dispensing service to download the prescription, and the *getClinicalData* method to retrieve the previous records.

The pharmacist notes that the patient is allergic to the antibiotic prescribed. They contact the prescribing doctor directly and arrange for the medication to be changed. The pharmacy system uses the *updatePrescription* method to amend the prescription (resulting in a notification to the prescribing doctor and the patients usual GP) and the *notifyDispensing* method that updates the local R-CDR with the dispensed medication.

All of the interactions and system accesses are subsequently visible to the patient via his on-line portal, which can use the *getAuditEvents* method to show who has accessed his record and when.

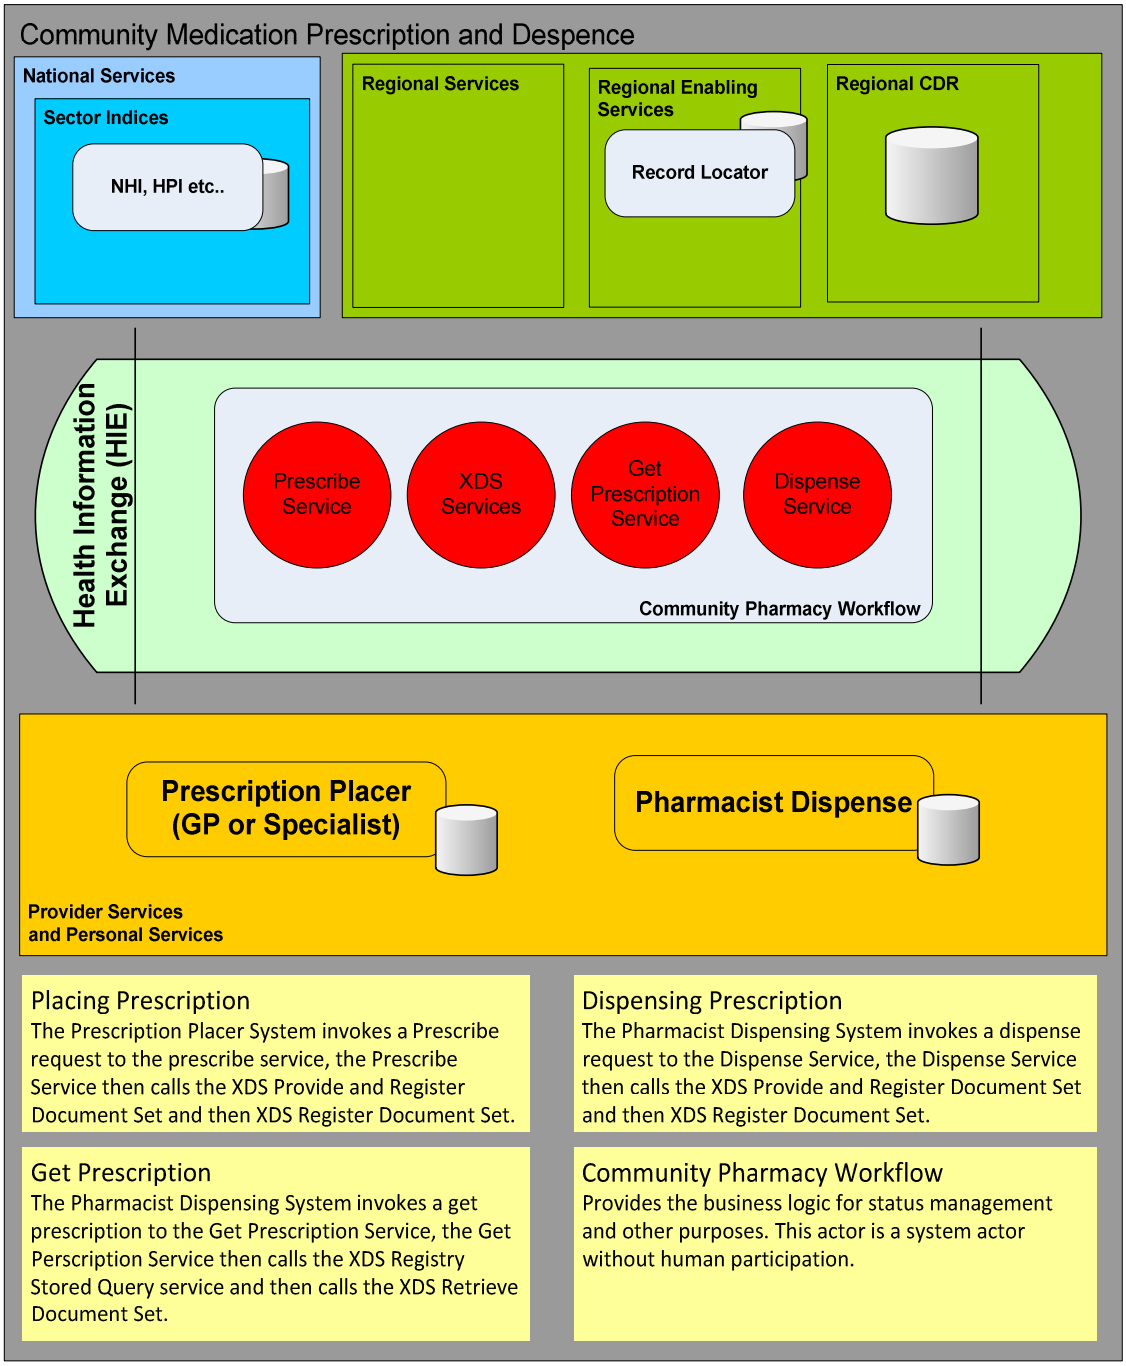

Figure 30 – Sample Storyboard

13 Appendix E: Terminologies

A clinical terminology is a structured list of concepts, their associated descriptions and relationships for use in clinical practice. These describe the care and treatment of patients and cover areas like diseases, interventions, treatments, drugs, and healthcare administration. Terminology is commonly used as an umbrella term to include coding and classification systems, nomenclature systems, controlled vocabularies, and at times biomedical ontologies.

13.1.1 Description

The Exchange Content Model defines a structure for holding information, but individual items of information need to refer to a terminology to be meaningful. For example the code 34101-6 is meaningless, unless you know that it is a LOINC code for a general medicine outpatient note.

The Exchange Content Model defines the structure and semantics of health information related to capture and representation. For example a particular section may define whether it will hold instances of clinical observations or prescription orders. Likewise a data element may specify a diagnosis or an anatomical site and so on. But there is no real-world domain knowledge embedded into the Exchange Content Model – this is what differentiates an information model from terminology or ontology. The latter are formal representations of real-world facts or objects; such as heart is part of the cardiovascular system or malignant melanoma is a type of skin cancer. This can be as simple as hierarchies given in ICD or in the other end exhaustive semantic relationships defined in SNOMED.

One should also make clear the purpose of using terminology; there is no point in encoding each and every data element. However using standard terms for commonly queried items (e.g. diagnosis, procedures, reasons for encounter, medications) or semantically significant items (e.g. depicting whether information is about subject of care or family history) is essential.

13.1.2 Reference Sets

Most terminologies are designed for acting as ‘reference’. This means a particular terminology includes facts about almost everything in its scope. Therefore most terminologies are ‘big’ – tools and expert knowledge are required to make use of it. In reality only a small part - actually a tiny subset of terminology will be needed to encode information being exchanged. Defining a custom set of terms for a specific purpose is called terminology sub-setting and the resulting set is called as Reference Set or simply RefSet. This not always simple as selecting all items under a hierarchy or hand picking certain terms but involve complex querying using relationships and logical operations.

For example a specifying a RefSet in SNOMED for a list of infectious diseases of the urinary track that is not of viral cause involves a query against terminology using disease and body system/anatomical location axes and type and causative agent relationships. Currently there is no standard querying language for defining RefSets, although IHTSDO and HSSP (a joint OMG/HL7 initiative) is working towards this.

13.1.3 Terminology Bindings

Linking parts of Exchange Content Model (certain data elements or structures) to terminology is called terminology binding. This allows for specifying rich semantics that acts like a reality check for correct interpretation and is crucial for advanced decision support. In DCM, specifically 13606/openEHR archetypes, it is possible to link every item to one or more terminologies. This can either be done by specifying the term code, terminology name and version for any item or using a RefSet. In the UK NHS have done a significant amount of work on terminology bindings.

13.1.4 Terminology Services

With the services oriented approach in mind, provision of terminology through a standard interface as a software service is called a terminology service. This insulates systems using the service from the complexity and internal workings of the terminology. One other advantage is that changes to the terminology can be made without breaking existing applications using the terminology service. HL7 has defined the Common Terminology Services (HL7 CTS) in this space. However, its use has been limited due to the intrinsic dependencies to other HL7 v3 artefacts. The second generation of the standard (CTS2) is currently being developed by HSSP – a joint initiative of HL7 and OMG. Its purpose is to specify a universal representation model and common operations for all terminologies and then provide standard interfaces to the terminology service.

In New Zealand a number of international and national terminologies are currently being used. While a few are mandated by HISO as national standards (e.g. NZPOCS, based on LOINC) and some are either endorsed (ICD10 AM and ICD-O) or used as de-facto standards (e.g. READ codes) by the sector. Some others are in the pipeline such as the NZULM based on NZMT.

14 Appendix F: Behaviour

The purpose of this chapter is to prescribe how inter-system behaviour should be architected and specified. The main driver in the choice of approach is the need to enable the continuum of care and shared care records specified by the National Health IT Plan.

The behaviour model delineates the roles, boundaries, activities and interactions of the various actors – systems, components and users – that participate in a particular solution. It is a dynamic view that complements the static view provided by the Exchange Content Model and payload specifications (described separately in this document). The behaviour model derives from requirements analysis by a formal process.

This chapter discusses the topic and states directives under these headings:

- Service Oriented Architecture
- Analysis and Design Methodology
- Functional Model
- Behaviour Modelling
- Technical Frameworks
- Continuum of Care Domain
- Localization

14.1 Analysis and Design Methodology

Interoperability is complex, requiring consideration of a range of factors:

- The business requirements behind the exchange
- The nature of the information to be exchanged
- How the information will be represented
- How data will be formatted
- How data will be transported
- How data will be accessed
- How workflow will be effected

14.2 Specification Framework

The HL7 Services Aware Interoperability Framework (SAIF) can help to organize the analysis and design process and its outputs. SAIF is geared towards practical achievement of working interoperability in some problem space. SAIF is technology agnostic (despite the HL7 association).

Details can be found at http://wiki.hl7.org/index.php?title=SAIF_main_page and http://en.wikipedia.org/wiki/HL7_Services_Aware_Interoperability_Framework.

SAIF organizes reference architecture and solution architecture material by an adaptation of ISO Reference Model for Open Distributed Processing (RM-ODP) viewpoints, and distinguishes conceptual, platform independent (logical) and platform specific (physical or implementable) models.

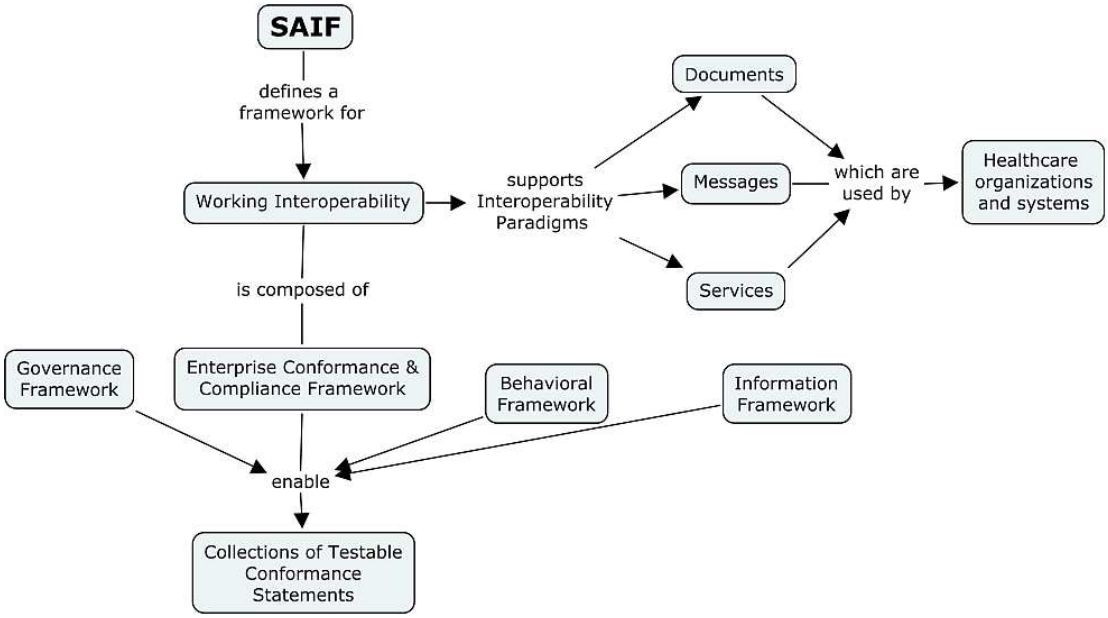

Figure 31 – SAIF composition

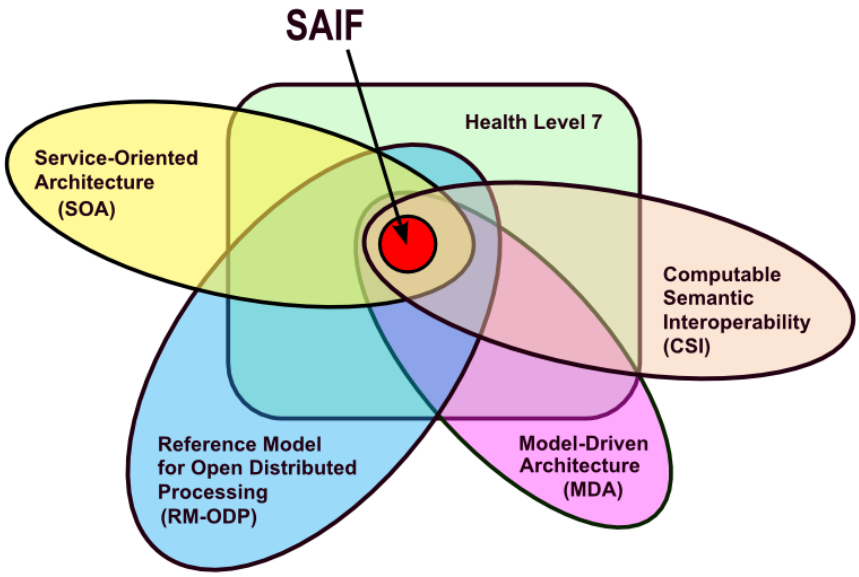

Figure 32 – SAIF as a confluence of approaches and methodologies

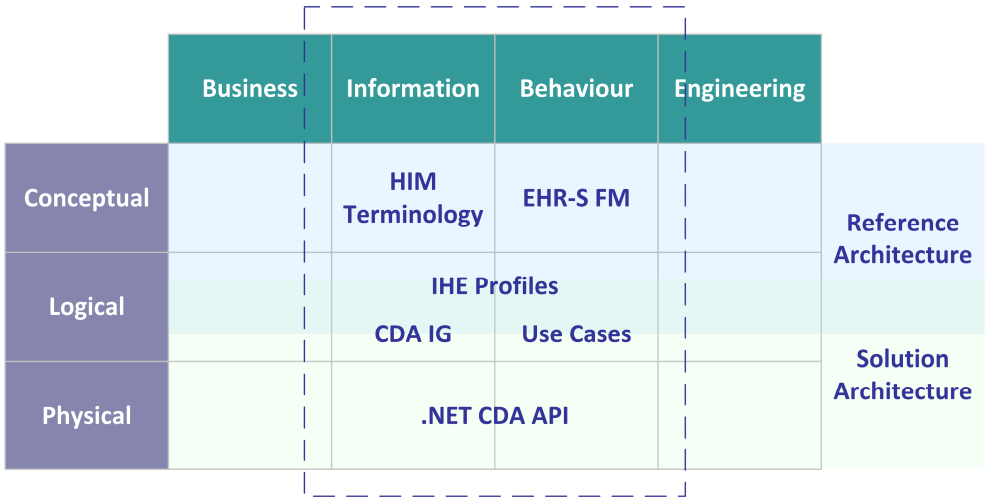

Figure 33 – SAIF specification matrix with example artefacts

14.3 Functional Model

ISO/HL7 10781:2009 defines requirements for functional interoperability of HIEs and R-CDRs. It represents the functions that an EHR should perform, but makes no statement about technology, internal design or messaging interfaces. The functions may, in turn, be a good basis for defining services to access and update them.

Direct Care	DC.1	Care Management
	DC.2	Clinical Decision Support
	DC.3	Operations Management and Communications
Supportive	S.1	Clinical Support
	S.2	Measurement, Analysis, Research and Reports
	S.3	Administrative and Financial
Information Infrastructure	IN.1	Security
	IN.2	Health Record Information and Management
	IN.3	Registry and Directory Services
	IN.4	Standard Terminologies and Terminology Services
	IN.5	Standards-Based Interoperability
	IN.6	Business Rules Management
	IN.7	Workflow Management

14.4 Behaviour Modelling

Required behaviour should be expressed in use case models and dynamic models, with Object Management Group (OMG) Unified Modelling Language (UML) the preferred formalism.

The table shows the taxonomy of Unified Modelling Language (UML) diagram types, indicating those relevant to solution design for interoperability shown in bold.

<table>
<tr><td>Structure Diagrams</td><td colspan="2">Package Diagram
Class Diagram
Component Diagram
Deployment Diagram
Profile Diagram
Composite Structure Diagram
Object Diagram</td></tr>
<tr><td rowspan="2">Behaviour Diagrams</td><td colspan="2">Use Case Diagram
Activity Diagram
State Diagram</td></tr>
<tr><td>Interaction Diagrams</td><td>Interaction Overview Diagram
Sequence Diagram
Communication Diagram
Timing Diagram</td></tr>
</table>

The behaviour model delineates the roles, boundaries, activities and interactions of the various actors – systems, components and users – that participate in a particular solution. It is a dynamic view that complements the static view provided by the structure model.

14.5 Technical Frameworks

Many organizations around the world have developed technical frameworks for the development of standards-based health informatics solutions. Integrating the Healthcare Enterprise (IHE) is the most international among these and has a substantial body of specifications that continues to grow. It has been decided that IHE technical frameworks will provide the key starting point for analysis and design for interoperability.

IHE has a technical framework for each of a number of healthcare domains including patient care coordination (continuum of care), laboratory and pharmacy (medication management). Technical frameworks are collections of what are called integration profiles, which are template solutions to business process problems.

Once developed, integration profiles are released for ‘trial implementation’; if successful they become ‘final text’; and they may be put ‘under revision’. Revision levels are numbered.

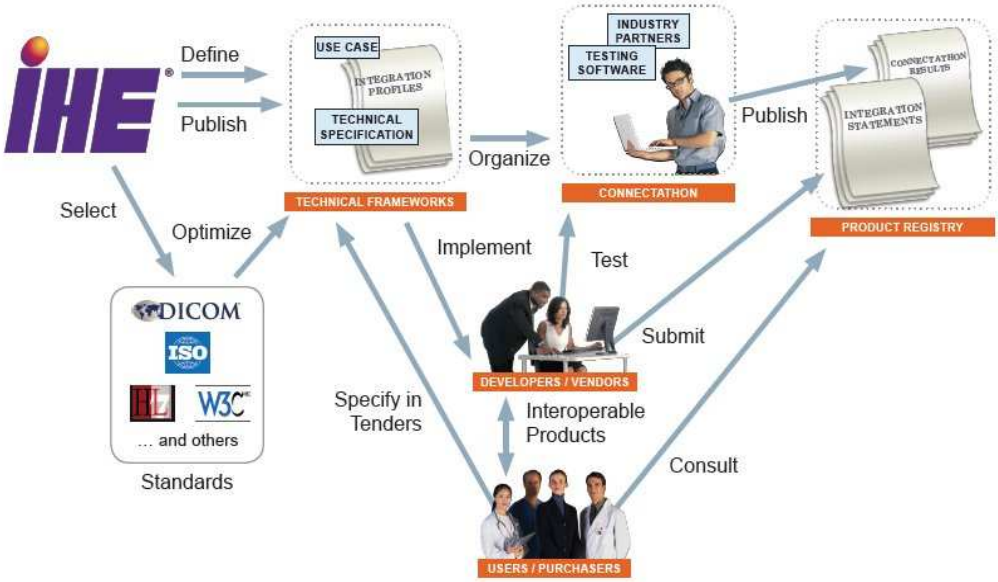

Figure 34 – IHE integration profile lifecycle

Each integration profile defines the actors, transactions and information content required to address the clinical use case by referencing appropriate standards. Content profiles ultimately define pay-loads.

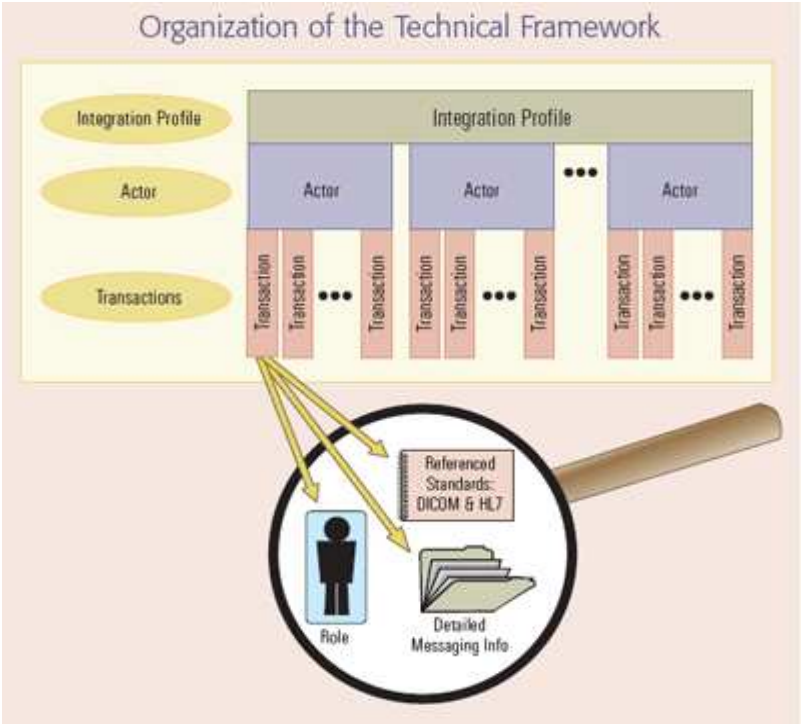

Figure 35 – Organisation of an IHE Technical Framework

An affinity domain is a group of healthcare enterprises that have agreed to common policies for in-formation sharing and have common infrastructure.

14.5.1 Healthcare Domains

IHE defines the following healthcare domains and publishes a technical framework for each. The most important to the National Health IT Plan are highlighted. IHE has a development roadmap and continues to add to the list.

Cardiology	Patient Care Coordination
Dental	Patient Care Device
Endoscopy	Pharmacy
Eye Care	Quality, Research and Public Health
Laboratory	Radiation Oncology
Anatomic Pathology	Radiology

Supporting all of the above is the infrastructure technical framework, which has record management and security integration profiles used by all other technical frameworks.

The pharmacy technical framework has integration profiles for medication management in both the hospital and the community, with these settings receiving somewhat different treatment (the hospital solution is HL7 v2 based while community uses v3). This difference reflects changing standards as profiles are created.

14.5.2 Workflow

Workflow refers to the steps that need to be performed to complete some business process. The business process is likely to involve multiple participants. An example is e-prescribing, where the participants will include (at a minimum) the ordering clinician, the patient, a pharmacist and a system that tracks the workflow as it executes.

Establishing the workflow required to support a business process is the essential first step in creating the information systems that are required to support that workflow.

14.5.3 Continuum of Care Domain

Requirements for the continuum of care are addressed for the most part by the IHE Patient Care Coordination technical framework. The diagram shows key integration profiles within that domain.

The table provides more detail on the most relevant of these profiles.

Functional Area	Integration Profiles
Referrals and Discharges	*Medical Summaries (MS)* defines the content and format of discharge summaries and referral notes *Emergency Department Referral (EDR)*
Clinical Workstation	*Query for Existing Data (QED)* gets core clinical information on problems, medications, immunizations and diagnostic results
Shared Care Record	*Care Management (CM)* is about management of long term condi-

	tions *Request for Clinical Guidance (RCG)* is for clinical decision support *Functional Status Assessments (FSA)* *Exchange of Personal Health Record Content (XPHR)* *Emergency Department Encounter Record (EDER)* *Patient Plan of Care*
Privacy	*Basic Patient Privacy Consents (BPPC)* defines a model and vocabulary for recording patients' wishes with respect to information sharing This profile will be localised as necessary in order to meet the requirements of the Health Information Privacy Code and the developing Shared Care Record consent model.
Maternity	*Antepartum Care Summary (APS)*
Medication Management	[*Pharmacy Technical Framework] Community Medication Prescription and Dispense (CMPD)* *[Pharmacy Technical Framework] Hospital Medication Workflow (HMW)*
Laboratory	*[Laboratory Technical Framework] Sharing Laboratory Reports (XD-LAB)*

The IHE Patient Care Coordination Technical Framework has a number of well-developed profiles; some of those shown appear to be well matched to current initiatives of the National Health IT Plan and warrant further investigation. Show the mapping.

14.5.4 Localization

Localization (in relation to standards) is the process of shaping an international standard for local use.

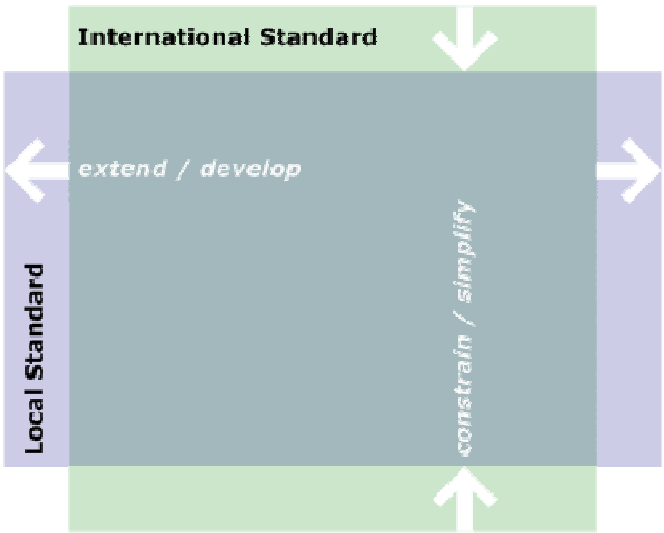

Figure 36 – Shaping an international standard for local use

In some respects, the local standard will tend to constrain, simplify and reduce choice, while remaining fully compliant with the international standard. Optionality will be either removed or made mandatory, which is a safe process. However, there is usually also the pressure to extend or develop beyond what the international standard allows. Because this is not inherently safe, the need should always be carefully examined and subject to governance.

The process of localization of international standards requires governance, particularly when extensions are proposed.

Local standards are required to be traceable to international standards (where they exist) and justified in any deviation from them. Participation in international standards development efforts will be important.

15 Appendix G: openEHR Detailed Clinical Models

The Detailed Clinical Models (DCMs) approach is a robust way of describing structured health information in a way that can easily be understood and maintained by healthcare professionals. DCMs combines healthcare concepts, clinical context, data elements and their organisation, terminology and associated metadata in a technology agnostic way. Practically, they specify labels, data structures and types, valid value ranges and enumerated values for each information item. The main premise of DCMs is that data model, user interface, messaging and document exchange and legacy system integration are all based on the same specifications.

DCM creation is commonly known as two-level modelling.

- The **first level** comprises a reference model where these common building blocks are formally defined by the standards (e.g. ISO 13606/openEHR and HL7). These are fairly stable technical artefacts that depict the generic characteristics of health information (e.g. data structures and types) and provide the means to define clinical context to meet ethical, medico-legal and provenance requirements. They also make it possible to leverage the vast amount of standardised terms and semantics by binding (linking) information items to biomedical terminologies. This is fundamentally important for automated decision support.
- At the **second level**, the clinical concepts are constructed by pulling together common technical building blocks and constraining those (e.g. defining hierarchy, optionality, repeatability, providing default values and etc.) using visual tools.

A good analogy to understand how reference model, archetypes and terminologies relate to each other is using a limited set of standard lego blocks to assemble many different structures.

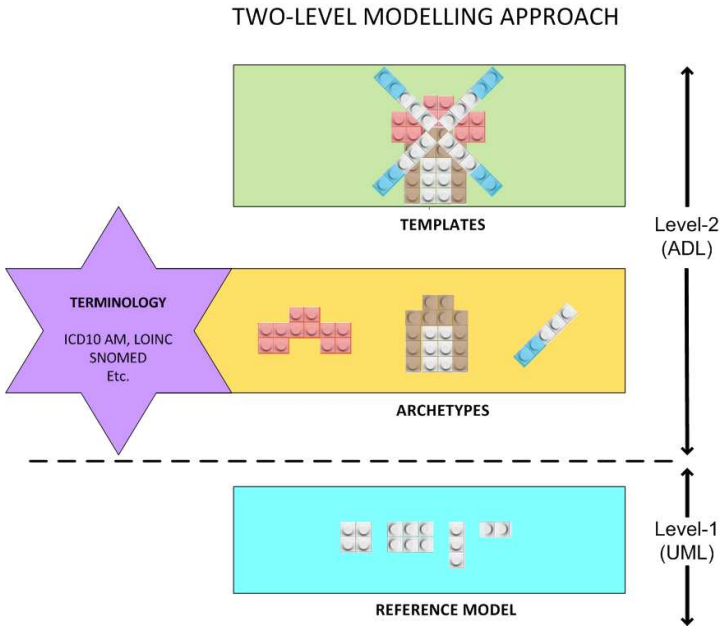

Figure 37 – Two-level modelling approach

Blood pressure measurement is a typical example of a DCM. It consists of a data part that holds the actual measurement data (e.g. systolic and diastolic blood pressure). Within the DCM other essential information required for the correct interpretation by a different system (or a clinician), such as cuff size (adult, child, etc.) and patient position (e.g. lying, sitting) are also captured. A blood pressure measurement DCM represented as a mind map is shown in figure 38 (from NEHTA).

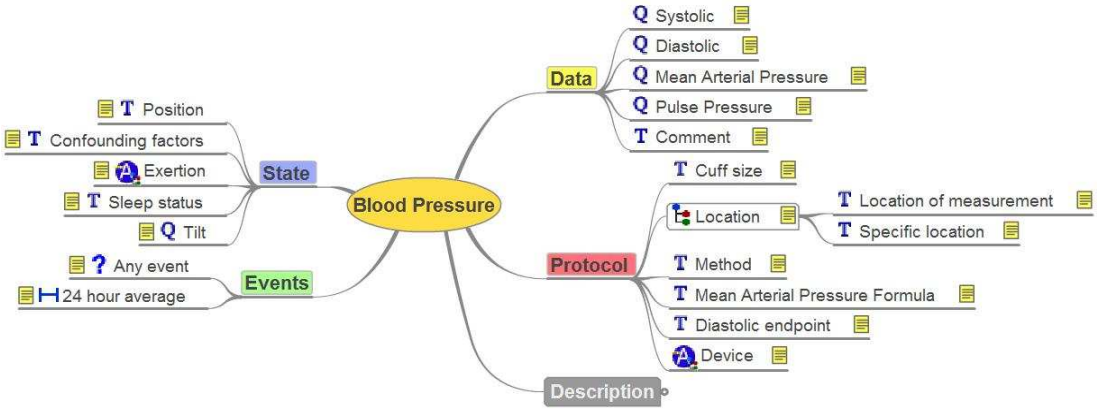

Figure 38 – Blood pressure measurement archetype

DCM is designed to contain all possible data elements for a given concept; hence it can be seen as a *maximal data set*. In most cases only a fraction of data elements defined in a DCM will be used by any one system; however using data elements from a common set will ensure consistency among implementations and when the data are aggregated from different sources they will conform to the same DCM and thus be comparable.

DCM can be automatically transformed into human readable (e.g. mind maps, UML) and computable (e.g. XML schema) artefacts and made available for technical professionals – see figure 39. This effectively removes much of the dependency between healthcare and technical professionals and separates the business and technical concerns. This capability can significantly reduce the time and effort required to build and maintain health information systems while keeping sector-wide implementations fairly consistent and interoperable.

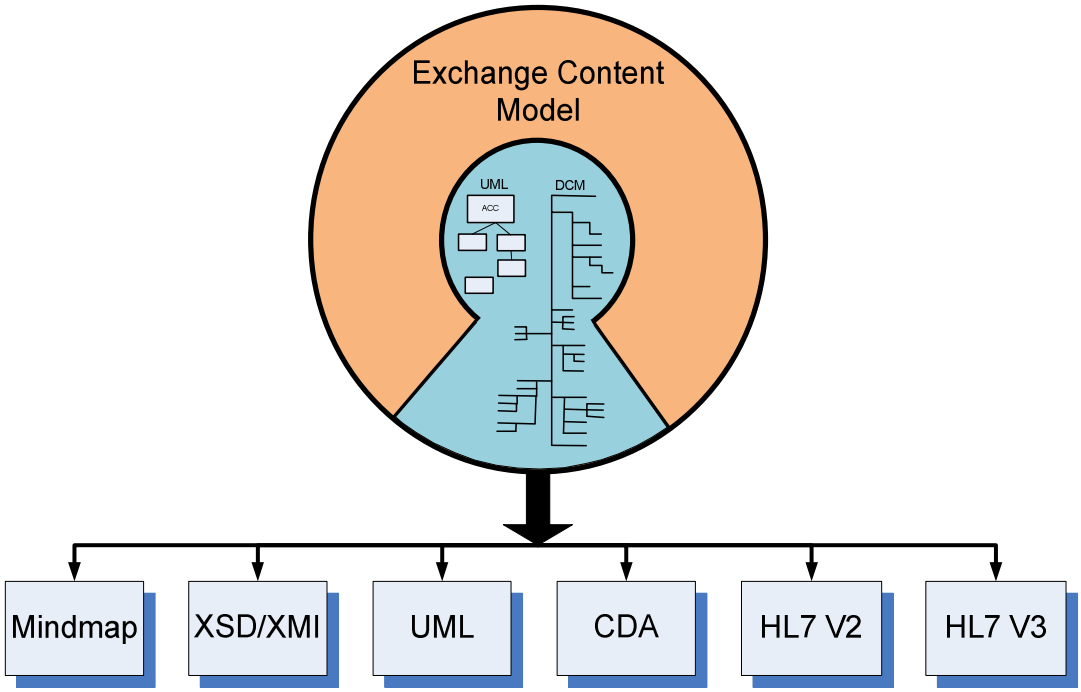

Figure 39 – Serialisation of the Exchange Content Model

Governance of the Exchange Content Model, especially keeping the core model stable and consistent over time, will be critical. DCM approach offers many advantages in that respect by means of formal archetype specialisation mechanism, versioning, and a web based collaboration and artefacts repository (such as Clinical Knowledge Manager, recently deployed by NEHTA).

Another advantage of using the DCM approach is that it is possible to adopt existing DCMs from other national programmes (e.g. NHS, NEHTA) that will significantly reduce time and effort to develop our own models. It may be necessary to alter some of these (using archetype specialisation methodology to keep international compatibility) and create New Zealand specific DCMs where there are significant differences between the health systems.

Since a DCM can be automatically transformed into other forms, keeping as much of the content as possible as DCMs will offer single source control over the Exchange Content Model. Where some non-clinical content cannot be represented as DCM then some manual work may be necessary during transformation.

It is important to realise that this model applies only to the interoperability space – it does not dictate how an individual health information system (e.g. EMR, CDR) models its data. However, in order to interoperate, they will need to be able to output data with the required semantics and granularity, and then map to the Exchange Content Model.

Archetypes should be used as the mechanism to define DCM, as these are readily understood by clinicians and have defined protocols for transformation into other artefacts for information exchange.

Archetype specialisation should be adopted to extend the core model using DCMs. Where DCMs cannot be used then these concepts should be extended using the same principles following object oriented specialisation/generalisation methodology. This will ensure backward data compatibility and also being able to perform generalised queries against highly granular specialist repositories.

15.1 Governance

The Exchange Content Model needs controls over its creation, ongoing development and use. There should be controls over the metadata registry and DCM repository.

A multi-organisation and multi-disciplinary editorial panel should be formed and have governance over the model. This panel should appoint domain experts for review and shall be responsible for the quality assurance of DCMs.

There should be controls over payload definitions. It will be important to get these definitions right first time to avoid flux and impact on mature and widely deployed software.

Web 2.0-style Exchange Content Model management tool (Clinical Knowledge Manager) should be set up for effective governance. This will help clinician engagement and enable collaborative and single source control.

16 Appendix H: Security

This section provides definitions and states requirements for security over the sharing and exchange of health information. These are requirements for HIEs, their participant systems and end users.

16.1 Authentication

Authentication is about identifying the individual – whether a practitioner, the patient or someone else wishing to access information on the patient's behalf. This requires both some identifier for the individual, and some way for them to prove to the system that they are who they say they are.

16.2 Authorisation

Authorisation is the act of controlling access to information. Commonly this is defined in terms of a role – e.g. a clinician, patient or an administrator. An individual persons role may change depending on the circumstances – e.g. a doctor accessing their own records is in the role of a patient, but when treating another is in the role of a clinician.

16.3 Audit

Audit is the process of recording accesses and changes to information – who does what, when and why. All accesses to patient identifiable information must be recorded and retained according to security principles, which should detail the right of the patients and others to access those audit records.

16.4 Privacy

As health information is one of the most sensitive categories of information available, both security and privacy issues are paramount. This section describes the specific subset that is privacy.

Privacy is all about controlling access to information and is a hotly debated topic, with legislative and other components. There is a separate work stream under the Sector Architecture Group dealing with security and privacy, so this topic is dealt with only superficially here.

This Reference Architecture states the following guidelines, but it is understood that there are national and local policies that apply.

- Patients always have access to their own information. It may be that this is given in a controlled fashion – for example a diagnosis of cancer should be given in person so that support can be given.
- The patient has overall control over who has access to their information, except in exceptional circumstances such as mental illness or where an unconscious patient is to be treated in an emergency situation.
- The patient is always able to find out who has accessed their information.
- The patient has the right to expect that information about them is held securely, and the system holding that information implements the defined security principles around storage and access.