INTERWEAVE

CONNECTING CARE

# Abstract of a Cookbook for Regional Interoperability in the Yorkshire & Humber Local Health and Care Record

## 2nd Edition

## Preliminary Draft

Version 1.1 –20th June 2021

# Table of Contents

## Version Control

| Version | | Release Date | Released By | Reason for Release |
|---|---|---|---|---|
| 1.0 | | 31/5/2021 | R Hickingbotham | Internal discussion draft |
| 1.1 | | 20/6/2021 | R Hickingbotham | Incorporate early review comments |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

## Reviewers

| Initials | Name | Role | Organisation |
|---|---|---|---|
| TD | Tim Davey | YHCR Solution Architect | PA Consulting |
| | | | |
| | | | |
| | | | |
| | | | |

| | | | |
|---|---|---|---|
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

# 0 Preface to this Edition

A little more than 2 ½ years separates the first and second editions of this document. During this period, the Yorkshire and Humber Care Record has evolved from being an aspiration into being an operational product. It has transitioned from being a programme of work funded and overseen by NHS England/ NHSx to being a sustainable technical asset owned by its participant organisations with dedicated management and operational resourcing.

The first edition of the Abstract of a Cookbook was written at the outset of Yorkshire & Humber (Y&H) Local Health and Care Record Exemplar (LHCRE) programme. Its purpose was to establish the intentions of the programme and to set out a vision for a future interoperability capability that would deliver long-term clinical benefit. It was a document which would be used to coalesce consensus among the senior executives of the 74 major health and care organisations in the region that the considerable work and financial investment required to realise the aspiration would be well directed. It offered an explanation to NHS England/NHSx of the goals of the Y&H LHCRE and would act as a roadmap against which the programme's progress could be measured.

Much of the first edition was written in the future tense. It referenced detailed designs which had yet to be written. It described itself as an abstract being released in advance of the forthcoming cookbook which itself would be a collection of detailed design papers which would elaborate on the concepts introduced by the abstract.

This edition has been written retrospectively. Detailed designs for the componentry described here have been published, reviewed, and ratified by the Y&H Technical Design Authority (TDA) and indeed, in the main, have been realised as fully functional, operational software products. The delta between the two editions therefore represents the difference between what was originally planned and what has now been delivered.

Propitiously, the fundamentals of the original intention have withstood the test of time. The Yorkshire and Humber Healthcare Record is a health & care information exchange which uses a hybrid model for data persistence and is founded on the open standard of HL7 FHIR. The platform is largely bespoke, and the intellectual property is held by its participant organisations. It is vendor neutral and is available for use by any data provider or consumer that complies with the open standards.

Deviations from the original intentions have been relatively minor:

- OpenEHR doesn't currently feature as a platform component, solely the FHIR standards are used for data persistence;
- there is no reliance on HSCN, security has been designed for the open internet;
- the YHCR does not offer authentication services – authentication is solely the responsibility of data consumers;
- a hierarchical relationship between a regional platform for clinical research which acquires data from local platforms for population health management has not been found to be necessary.

Most differences between early intentions and the software which is actually operating are due to extensions and refinements designed with the benefit of time and with the contribution of skilled technical collaborators.

This edition highlights deviations and extensions to the original abstract using footnotes to content.

# 1 Introduction

## 1.1 Purpose of this Document

This document is an abstract of the detailed designs which have been authored in support of the development of the Yorkshire and Humber Care Record (YHCR). The designs are a set of papers which are in the public domain[1] and which are intended to comprehensively document the software and operating processes which constitute the infrastructure behind the YHCR.

Together, this abstract and the underpinning design documents constitute the Cookbook for Interoperability in the Yorkshire and Humber Care Record. This is a technical guide and standards reference library which could allow any interested and capable party to recreate the YHCR technology. It also supports those software products, care providers, local records and inhouse applications develop capabilities to enable their participation in the YHCR. Finally, it informs those looking to reuse the platform and software products[2] to support other local shared care records.

The abstract covers the same content as the design papers, but it does so at a much more summary level and is intended to be read by a wide audience of varying degrees of technical understanding. The nature of the subject matter necessitates references to technologies and data standards but familiarity with these should not be necessary to understand the principles discussed here.

## 1.2 Background to the Yorkshire and Humber Care Record

The YHCR is the product of the Yorkshire & Humber (Y&H) region's[3] participation in the national Local Health and Care Record Exemplar (LHCRE) Programme. The LHCRE programme provided funding to five regions with aim of establishing technical capabilities that would support the creation and operation of a longitudinal health and care record. The record and technical capabilities would be used for:

- direct care;
- citizen/patient participation in their care;
- population health management.

Y&H response to the LHCRE challenge was based on certain open principles:

- open architectures which are vendor agnostic, use substitutable components, and provide autonomy and self-determination to participants;
- open standards which remove barriers to participation and promote innovation through access to data;
- open documentation which invited scrutiny over the strength of designs;
- open access to the assets created to promote reuse and wider benefit in the NHS and local government.

[1] https://yhcr.org/resources/technical-papers/

[2] The assets created for the YHCR are the intellectual property of its participating organisations and are available for others to use under various license models.

[3] The region comprises the 3 Integrated Care Systems of Humber Coast & Vale, West Yorkshire & Harrogate, and South Yorkshire & Bassetlaw. It comprises 74 major health and care organisations and has a population of 5.8 million people.

Y&H exited the LHCRE programme in March 2020 having successfully implemented the technical capabilities required of it and onboarded enough participants to demonstrate practical use.

The YHCR has now entered business-as-usual, the platform is sustainably funded, and focus has moved from capability creation to expanding meaningful use. At the time of writing of this the second edition of the abstract, eleven major health and care organisations, four software vendors, the region's General Practises, and several data analysts are providing data and/or using the platform with varying degrees of maturity. Work is also underway to deploy the platform in other geographies.

## 1.3 Governance

The YHCR operates the following governance bodies[4]:

- a Programme Board that is accountable for strategy, finances, and risk;
- a Clinical & Technical Design Authority (CTDA) Board that is responsible for clinical adoption, benefits realisation and ratification of standards and designs approving by working groups;
- a Technical Design Authority (TDA) that is a working group reporting to the CTDA and is responsible for approving technical design decisions;
- a Data Architecture Design Authority (DADA) that is a working group reporting to the CTDA and is responsible for approving data designs;
- a Management Team who produces designs, manages the development programme, and provides assurance over the onboarding of participants to the information exchange.

## 1.4 The Yorkshire & Humber Care Record Interoperability Architecture

The architecture is vendor neutral and is based on the HL7 FHIR (Fast Healthcare Interoperability Resources) open standard for sharing data between care providers.

It is vendor neutral because there is no dependency on any one system vendor. The architecture enables existing systems to contribute to a shared care record and makes no presumption as to how the record is consumed. In fact, the architecture is designed to allow the shared care record to be displayed in any technically capable user interface and affords the possibility that different care providers will have different ways of making the record available for their care professionals.

The persistence model is hybrid[5]. The YHCR prefers to keep data where it originates and to allow data to be obtained, on demand, by those who need it from those organisations that created the data. Federation of data provision is the standard way that the YHCR operates for direct care. Data is persisted within the boundary of the organisation that controls it and governance responsibility for the data is retained by that organisation. Changes to the data are immediately reflected in the YHCR and responsibility for cleansing and normalising data is devolved to those that manage the data for operational purposes.

The hybrid model allows data to be centrally persisted when there is a case do so. The YHCR's solution for population health management is a case in point. Data used for analysis purposes is

[4] The Technical Design Authority and Data Architecture Design Authority were created during the LHCRE programme and did not exist at the time of writing of the first edition.

[5] The first edition of the abstract described the persistence architecture as federated. Whilst this very much the solution's predisposition, it didn't accurately acknowledge the possibility of centralised persistence and led to a misunderstanding for some readers.

assembled in a central location in pseudonymised or anonymised form. Data is also persisted centrally where controllership is in common among participant organisation or where it is used in the operation of the YHCR.

The following graphic[6] depicts the main features of the architecture and offers readers a useful point of orientation for the rest of the document.

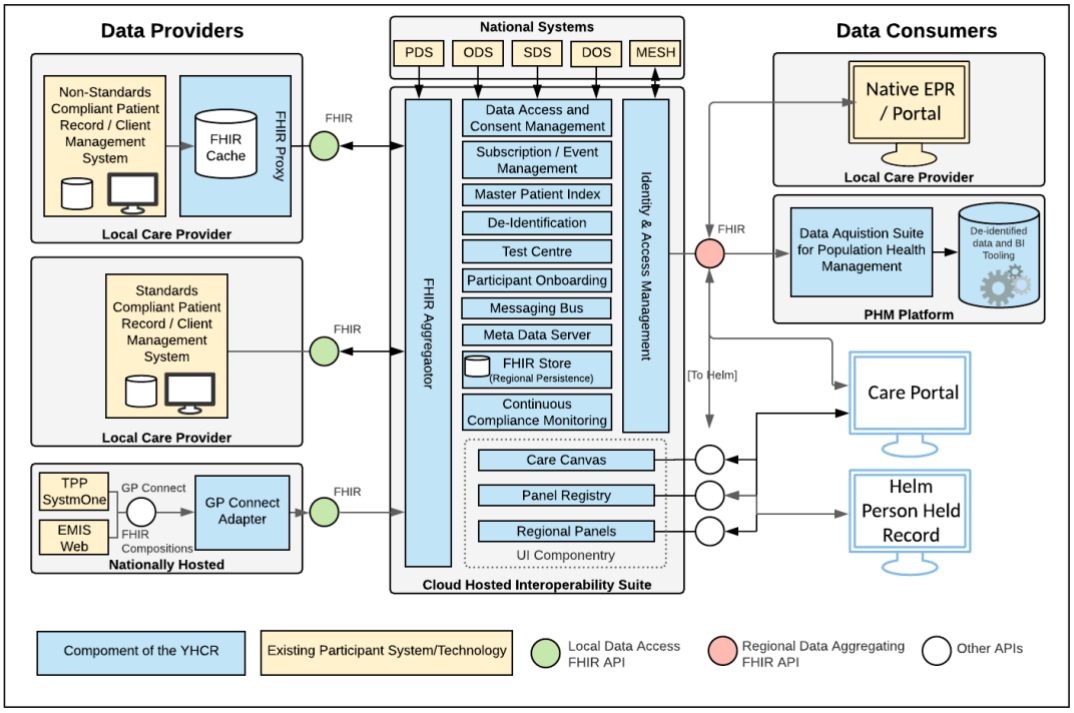


The colour coding in this diagram identifies the components that have been appropriated for the YHCR and comprise the solution. Except for some Business Intelligence tooling installed on the PHM Platform, all these components have been built for purpose under the LHCRE programme and the intellectual property is held by participants in the YHCR.

This document focuses on explaining and drilling down into the functionality of the various parts of this diagram. The detail will not be pre-empted by further explanation here other than to emphasise that the architecture is founded on FHIR standards. Data exchanged between data providers and data consumers complies with the FHIR standard resource model. Most data interactions use the FHIR API standard. Any system which is compatible with FHIR[7] can participate in the YHCR.

However, the YHCR is built on a philosophy of self-determination and it will not wait on its system vendors to meet high degree of standards compliance for it to mature. The YHCR builds data adapters where necessary to achieve compliance for non-standard sources of data. It offers a

[6] The first edition of the abstract used a diagram which was submitted in Y&Hs bid for LHCRE status. The key differences are i) OpenEHR is not used for regional persistence, ii) XDS, whilst compatible with the regional architecture and important for inter-regional interoperability as explained below, does not feature as a data source in the region.

[7] The FHIR standard has many elements, most of which are optional. The YHCR is demanding in its requirement for comprehensive FHIR support and bridging software may be needed for some products at the lower-end of the compliance spectrum.

standards compliant presentation layer to allow care professionals to access its data where standards compliance is not available from their systems of choice.

### 1.4.1 Data Providers and Data Consumers

Care settings are peers in the interoperability architecture. A care setting is a source of data which can be accessed by another care settings acting in the capacity as a data consumer. A cloud hosted interoperability suite acts as orchestrator of this relationship.

The goal for the YHCR is to make a longitudinal care record available to data consumers. This record is composed of all facts about a patient known to the health and care system and describes all encounters, procedures, diagnoses, tests, assessments, prescriptions, and plans undertaken or prepared by any clinician working for any care provider in the system.

It is of no concern to the YHCR how this data is made available to it[8] so long it meets quality and normalisation objectives and is presented securely and in compliance with the FHIR standards. The YHCR is open to what constitutes a data provider:

- a care setting which presents all data managed by any of the systems that it operates in an aggregated form;
- a single system which presents data under its management for one or more care settings;
- the GP Connect national interface into GP system data;
- a shared care record which aggregates data from several care settings.

Care settings acting as data providers typically use a FHIR proxy to simplify compliance. The FHIR proxy is available in two functionally identical versions: an InterSystems Ensemble/HealthShare/IRIS production and a standalone containerised Node.JS open source product. It provides plug-and-play compatibility with central infrastructure and allows a care setting to focus on mapping data to FHIR resources.

The YHCR is equally indifferent to what constitutes a data consumer. Examples include:

- a user interface which provides access to the longitudinal care record to a care professional at the point of care;
- a system which responds to notifications of events[9] occurring in the health and care system, perhaps attendance by a patient at a service or a diagnosis of a particular condition being made, and triggers an action such as auto-enrolling a patient on a care pathway or performing a safeguarding function;
- an algorithm analysing data such as blood test results and based on trends in measurements prompts for an intervention in care.

The YHCR platform for Population Health Management (PHM) is itself a data consumer and acquires data through the central interoperability suite using the same FHIR APIs as consumers categorised above.

[8] The first edition of the extract expressed a preference that a care setting acted as a data provider and that direct connections with individual systems would only be a tactical measure. A more detached position about the definition of a data provider has been helped by design paper 029 – Data Aggregators and Other Complex Data Providers

[9] The FHIR standard specifies enabling mechanisms for data subscriptions such as messaging which are implemented as regional capabilities.

THE PHM platform uses a data acquisition suite to capture an analyst's data requirements and interacts with data providers by placing requests for data on the FHIR Aggregator: a component of the central interoperability suite. Data is pseudonymised/anonymised in-flight but other than this, and subject to suitable Data Sharing Agreements and patient consent, the PHM platform has access to the same data as is used for direct care.

### 1.4.2 Data Normalisation

Standardisation of data as presented to the YHCR is highly important to enable consistent interpretation and avoid risk to clinical safety. Normalisation[10] covers concepts such as:

- ensuring that similar concepts presented from different data sources are presented in the same format;
- using common coding systems to classify concepts;
- de-duplicating multiple representations of the same event recorded in different data sources.

The YHCR is both ambitious and pragmatic in its objective to normalise data. Its goal is absolute uniformity in compliance with national standards for data models. However, there is a spectrum of maturity between its participants, and it uses a maturity model to allow care settings to participate early and refine their compliance to gold-standards over time and in line with clinical benefit. The early-day focus has been to standardise the format in which data is accessed and to ensure data models are fit for purpose. The maturity model and its relationship with meta-data is elaborated on in later chapters.

## 1.5 The Yorkshire & Humber Care Record Care Portal Architecture

Ideally, the longitudinal care record which is available from the YHCR will be accessed from the same systems that care professionals use for day-to-day purposes. There will be no need for separate logins, acclimatisation to different presentational styles and mental consolidation of information available from different sources. The ideal is that data from the YHCR is blended with the local clinical record to offer care professionals a single narrative about their patients and the ideal has been realised in several the care settings using the YHCR.

However, there are situations where this will not be possible and for these the YHCR offers a standard care portal.

The care portal adheres to the design principles of local autonomy and federation and uses an architecture which separates the concepts of a canvas and panels to allow different organisations to contribute functionality and different users to have personalised interactions with the YHCR.

The architecture is illustrated below:

[10] A regional approach to normalisation is defined by design paper 026 – "Data Normalisation".

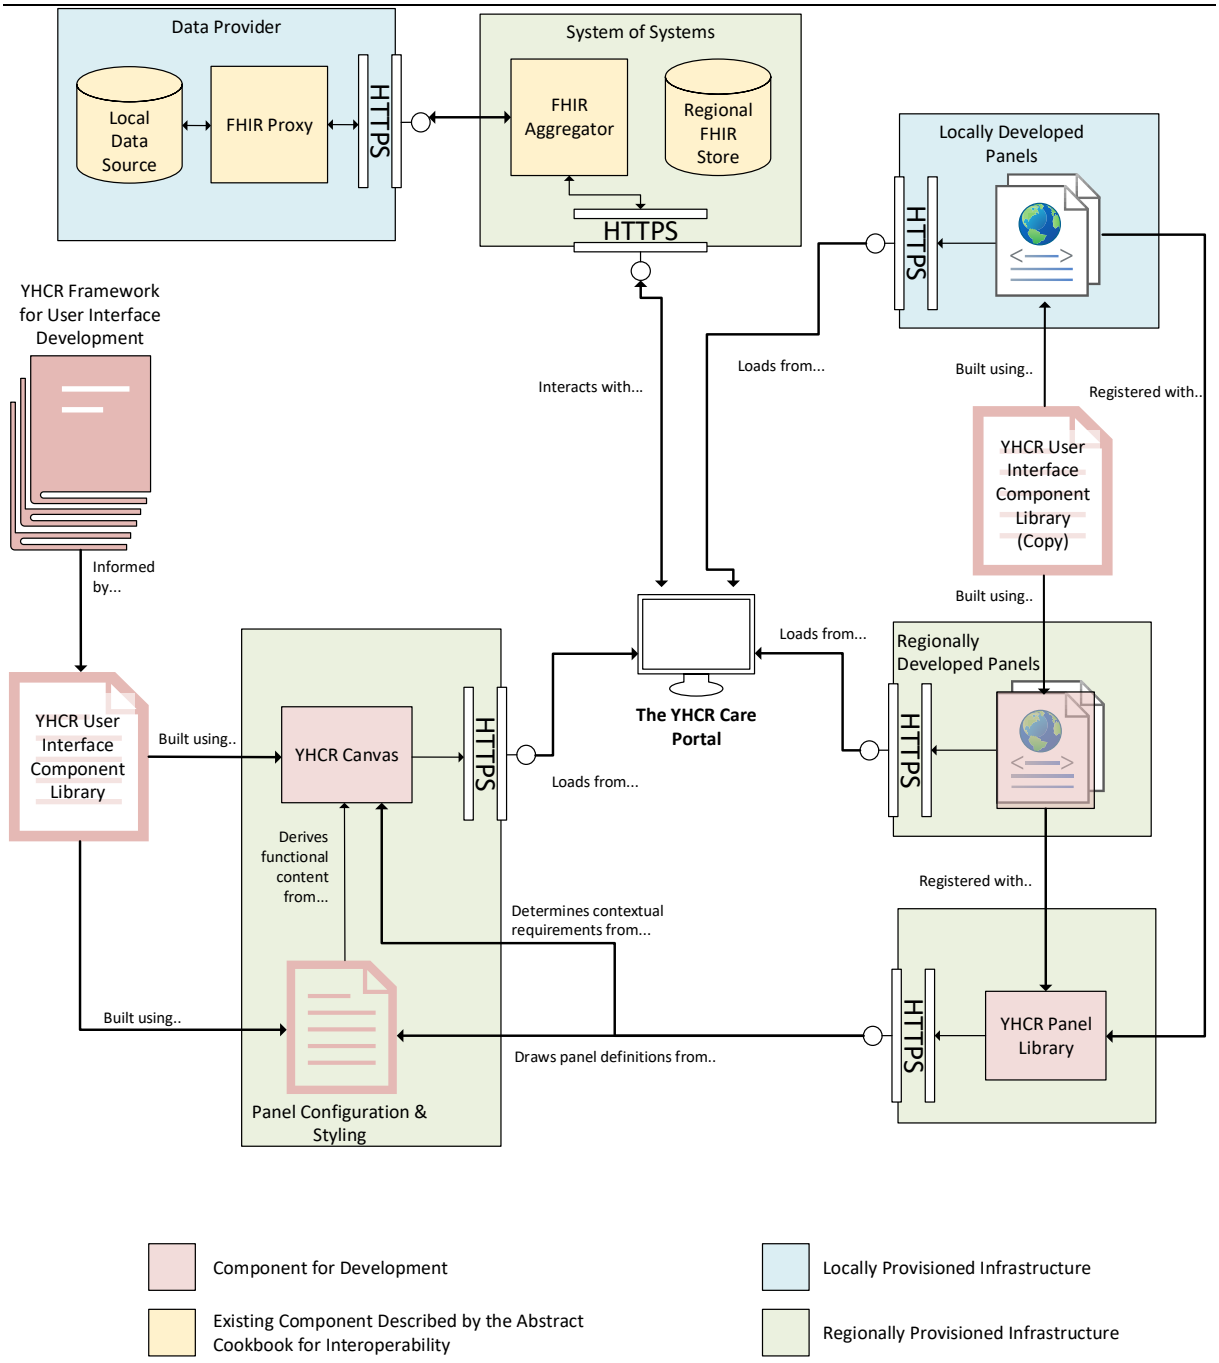


The components of the architecture are explained later but the key principles are:

- a regional canvas defines mechanisms for authentication or context launching and provides a configurable navigation system to functionality provides by panels;
- panels are regionally or locally developed (or can be offered commercially be a system vendor) and are developed using technologies of choice but comply with YHCR standards for web components;
- panels interact with YHCR data using FHIR and the regional FHIR Aggregator;
- panels must be registered with a regional panel library which provides regional control over clinical safety and alignment of user experience standards;
- a YHCR developed interface component library defines the relationships between the canvas and panels and enforces contextual constituency.

## 1.6 Architectural Design Drivers

The architectures have been driven by certain key principles and the cookbook retains these principles in a technical design:

- Reuse whatever is available: systems technical standards, coding systems and vocabularies.
- Federate data where practical: custody of data is important, don't make unnecessary copies.
- Centralise services where beneficial: there is cost-benefit to building once and using many times.
- FHIR where possible: FHIR is standard gaining international traction particularly in the NHS. It is relatively mature and provides a high coverage of use cases.
- Separate user interfaces from data: user interfaces are changed every few years, data lasts for decades.
- Normalisation where appropriate: normalisation is complex but is not essential in all use cases.
- Human interpretation of data allows for certain inconsistencies.
- Allow for local autonomy and self-determination.

## 1.7 Notes on the Persistence of Data

Whilst the architecture tends towards federation, some data at rest is persisted centrally. Several FHIR stores are deployed for persistence of:

- Citizen supplied data.
- Cached data where performance is a factor.
- Data which supports the operation of business processes across care settings.
- Audit data.
- Document locators.
- Encounter summaries.
- Data which supports the operation of regional infrastructure.
- Regionally cached copies of local resources for version control purposes.
- Pseudonymised data for data science, research and analytics.
- Person-identifiable data that is needed for execution of an algorithm to inform care intervention.
- Data that is not held in any other system, but persistence is required.

The operator of the regional infrastructure has the governance responsibilities of a data controller in this respect.

## 1.8 Notes on Infrastructure

A federated architecture relies on the availability of a relatively high-quality low-latency network. User interfaces which aggregate data in parallel from different sources will, without intentional mitigation, operate at the speed of the least performant source. The cookbook details techniques for overcoming limitations in network performance.

## 1.9 Notes on the Development and Application of Standards

Vendor neutrality or openness is wholly dependent on the adoption of standards; standards which define the way that data is interchanged, the format of the data, and its interpretation. Some of the

standards which are used by the cookbook are focussed on healthcare and it is acknowledged that social care is underrepresented.

The YHCR management team works with national and international standards bodies to improve the applicability of the standards in local government. It also hosts a regional academy which educates professionals to ensure that that the region has expertise to work with the chosen standards.

Practically applying standards to real-world situations is challenging and a uniform approach is essential for regional interoperability. A Data Architecture Design Authority has been established to facilitate discussion and the gaining of consensus on the treatment of common data topics.

## 1.10 Terminology and Presentation Devices used in this Document

The cookbook uses the term care professional to refer to someone who is involved in the care of person and includes clinicians, nurses, social workers and any other form of care provider.

FHIR resources (fragments of a clinical record) are referenced throughout the cookbook in italics i.e.: *PractitionerRole*

# 2 An Open Interoperability Platform

The Y&H approach delegates much responsibility to its participant organisations for creating and using the care record but in doing so it empowers them to pursue local objectives and retain a high degree of autonomy. The YHCR is more than a window into a virtual longitudinal care record. It is a platform that allows its participants to interact with each other in a frictionless open manner.

The YHCR's focus is on the implementation of data standards within a central cloud hosted interoperability platform. The platform is overlayed with functionality which helps participants interact in a safe and secure manner and, where necessary, with the consent of the citizens that the YHCR represents.

The central platform is supplemented with tooling which facilitates participants' compliance with the standards, rapid adoption, and re-use of developments between participants.

The central platform, the YHCR Exchange, is a functionally stable black-box API layer. The HL7 FHIR standard will evolve, and the platform will move with it. New tooling and functionality will be developed over time that improves operations and eases use by its participants. Efficiencies and performance gains will be made from ongoing enhancements. However, to its participants, the goal of the platform is to present an unchanging, reliable façade that encourages innovation at its periphery[11].

The application neutrality of the platform leads to deployment options. In the YHCR the platform is principally used to connect care setting to care setting. However, it is equally applicable to use it to connect system to system, ICS to ICS or LHCR to LHCR.

## 2.1 Interaction Patterns

The introduction alludes to the idea of a data consumer using the platform to acquire data by querying a longitudinal care record assembled on-demand from several data providers. This is an example of an interaction pattern. Whilst this is an important use of the YHCR Exchange it is only one of several which are supported.

| Interaction Pattern | Description | Use Case |
|---|---|---|
| **On-Demand Query** | **Synchronous Query**- Data is acquired on demand when it is needed – the principal mechanism employed in direct care. | • DNAR available to paramedics in an ambulance.<br>• History of local care available at a tertiary centre.<br>• Pre-natal records visible wherever a baby is born. |

[11] The YHCR develops and owns the intellectual property to a number of applications which are documented elsewhere.

| | | |
|---|---|---|
| **Data Management** | **Synchronous Update** – Participants can create and update data. Data can be maintained centrally in the exchange or the update may be targeted at data managed by another participant. | • A service referral is made by a participant directly into the EPR of another care setting.<br>• End of life care plans developed collaboratively. |
| **Bulk Acquisition** | **Asynchronous Query**- Used for requests for bulk data. Queries are serviced in the background and data returned when available.<br>Removes load from operational systems. | • Bulk data coded for direct care available for risk stratification.<br>• Correlations of condition development with treatment patterns and lifestyle factors. |
| **Notifications & Alerts** | **Subscription**- Allows an interest to be expressed in any data point in any care setting. Notifications flow when events occur. An enabler of proactive care. | • Alerting for patients at risk.<br>• Algorithmic analysis of condition development.<br>• Maintenance of frailty indices.<br>• Real-time analysis of prescription trends. |
| **Point to Point** | **Reliable Messaging**[12]- Peer organisations exchange messages when key events occur. | • Automated pre-admission of patients transported in ambulances at ED.<br>• Discharge messages and system wide patient flow management. |

## 2.2 Core Platform Features

The YHCR Exchange provides a comprehensive implementation of the HL7 FHIR STU3 API specification on a single endpoint. Behind this API is functionality that simplifies, secures, and ensures conformant relationships for participants.

### Simplify

A FHIR Aggregator provides data consumers a single endpoint from which to interact with all data providers. The FHIR Aggregator presents the data as though it came from a single source unifying data obtained from different providers around common points such as patients, organisations, and practitioners.

[12] This paper focuses on non-messaging interaction patterns. Whilst reliable messaging is an optional delivery channel for subscription notification as a pure interaction pattern messaging is only currently used in the YHCR for transfer of care and whilst an important application this papers focus is on more generic application of YHCR technology. More information on Reliable Messaging can be found in design paper 006 – "Reliable Messaging Infrastructure"

Subscriptions and Reliable Messaging are managed by central components removing the requirement for any point to point relationships between participants.

A Master Patient Index links local patient demographic data and patient identifiers to a regional patient golden record. Interactions can reference wither local patient records or the regional identity.

A Data Availability service provides a simple mechanism for a data consumer to discover the data providers which hold data about a patient/citizen.

A Cohort Manager supports enables regional maintenance of patient cohorts for local use. Cohorts can be used in queries and subscriptions to interact with data relating to any member of the cohort.

An Onboarding Suite provides participants a self-service experience when connecting to the Exchange.

Regional FHIR Stores provide central persistence capabilities and facilitate the development of collaborative applications needing to manage common datasets.

An automated test suite helps a data providers establish readiness to contribute to the Exchange.

**Secure**

The Identity and Access Management server provides an implementation of OAuth2 and enforces a YHCR managed public key infrastructure (PKI). Regional user identities are linked with local identities to ensure consistent privileges are applied regardless of the method of interaction with the YHCR.

The Onboarding Suite controls all DNS, firewall rules and configuration of the Exchange ensuring consistency in the application of security policies.

Forensic Analysis tooling highlights potential abuse of the YHCR Exchange and supports detailed analysis of usage patterns.

**Ensure Conformance**

A Metadata Server enables FHIR profiles[13] and coding systems to be regionally defined and for data transiting through the Exchange to be validated for compliance.

Continuous compliance tooling automatically validates a participants ongoing adherence to security and governance standards.

Monitoring tooling highlights performance degradations and points of failure across the whole Exchange.

Data quality monitoring tooling enables rules to be created which can be used to test data for complex integrity conditions.

## 2.3 A Platform for Population Health Management

The YHCR PHM Platform is independent of the YHCR Exchange and is operated as a distinct product. It acts as a peer with other data consumers and uses the same technical methods for accessing data (although its reason for use distinguishes it from consumers accessing data for the purpose of direct care and data access control policies may restrict the data that it is able to access).

[13] A mechanism for constraining a FHIR resource to apply it to a particular use case.

The PHM Platform has the following components:

- data acquisition tooling;
- data pseudonymisation/anonymisation engine;
- data persistence;
- business intelligence tooling.

The last two of these components are outside of the scope of the Interoperability Cookbook and are covered in other YHCR documentation.

### 2.3.1 Data Acquisition and De-identification

Data acquisition tooling enables a clinician or researcher to define a data set using the FHIR data model and centrally managed patient cohorts. The tooling issues asynchronous queries to the YHCR Exchange and collects results as they become available. Optionally the tooling will keep a dataset up to data by following up with subscriptions to new data created after the asynchronous queries have run.

Data acquired for use on the PHM Platform is automatically de-identified before it is persisted using the national Privitar solution. Data is either pseudonymised (in which case it can be re-identified subject to appropriate controls) or fully anonymised.

## 2.4 NHS X Definitions of the Technical Capabilities of a LHCRE

The YHCR was created under the national LHCRE programme and was developed in response to the technical capabilities which were defined as characterising a LHCR. Yorkshire and Humber exited the programme in March 2020 having been assured by NHS X as meeting all capabilities.

The following table summarises how the capabilities are supported.

| Ref | Technical Capability | YHCR Response |
|---|---|---|
| **TC1** | **Open APIs -** Enable professional and patient facing services to use open APIs to access information or interact with services across LHCRs and other points of care. | The YHCR embraces the FHIR API standard. Interactions between data consumers, data providers and the Exchange use the FHIR API. The YHCR implementation of the API standard is comprehensive both in the central infrastructure and model software provided for local implementation. |
| **TC2** | **Record Location** - Discover the location of records for a specific NHS Number. | The Data Availability service determines whether data is held on the YHCR and if so by which data providers. The service uses the NHS Number as the patient identifier. |
| **TC3** | **Events Management** - Allow events to be published into an event management service, which routes the event notification to other interested parties | The YHCR uses the FHIR Subscription mechanism for event management. |
| **TC4** | **Longitudinal Health and Care Record** - Data layer for the population enabling identifiable data sets are available for the purposes of direct care, and de-identified data sets are available for population health management. | The YHCR uses a hybrid approach to assemble a longitudinal care record. It preference is leave data where it is created and assemble data on demand. Data may be persisted centrally where there is a use case to do so. |

| | | |
|---|---|---|
| **TC5** | **Personal Health Records (PHRs)** - Allow Patients to select, from a range of PHR products, the one that best meets their needs that they can use to manage their online relationship and the personalisation of care. Capabilities include input of data by patients, access to test results, managing care plans and care preferences, and subscribing to alerts | The YHCR has a number of PHRs acting as both data consumers and data providers for patient captured data. |
| **TC6** | **Metadata Management** - Support for common metadata specification for patient records to facilitate discovery and matching. | The YHCR uses the FHIR Conformance and Terminology compartments for meta-data management and implements them un the YHCR Meta-data server. The meta-data server is compatible with the national terminology server. |
| **TC7** | **Reference Data Management** - Compilation of a core set of reference data which is the standardised view for the national data sets processing. | The YHCR holds reference data for Patients, Practitioners, Healthcare Services and Organisations. The data is drawn from national data sets and services (SDS, ODS, PDS, DOS) |
| **TC8** | **Information Standards and Terminology Classifications** - Adoption of Information Standard Notices (ISNs) to effect adoption of specific information standards across the service. Creation and management of the content for SNOMED CT, dm+d, OPCS and ICD10. Support for services around these such as the operation of a terminology server. | Adoption of ISNs is an ongoing commitment. At the time of writing the YHCR fully complies with all applicable ISNs. The YHCR Meta-data server is used for the curation of coding systems and can automatically assemble coding systems from other terminology sources where they are referenced in local FHIR profiles. |
| **TC9** | *Deleted by NHS X* | |
| **TC10** | *Deleted by NHS X* | |
| **TC11** | **Master Patient Index** - A patient index for the local population and for processing leavers and joiners. The master source of demographic data is the PDS. | The YHCR MPI indexes the local population and links substitutable local identifiers to NHS Numbers. PDS is the master source of demographic data for the YHCR. |
| **TC12** | **Data Rules Management** - A data rules engine to support the checking of data quality and other rules-based checks on the data processed nationally or locally | Data rules can be expressed using FHIR expressions and can be evaluated against all data transiting through the YHCR. |
| **TC13** | **Data Discovery Support** - Discovery of what data is held by a LHCR in terms of LHCR data schema and sources of data. | The YHCR offers 3 different methods for record discovery: a) the Data Availability service determines whether data is held on the YHCR and if so by which data providers, CapabilityStatements offered by individual providers identify the data types available from the provider, and FHIR queries can be constructed to extract precise details of data held. |
| **TC14** | **Data Transfer and Dissemination** - Transfer and dissemination of data from one place to another. This may be for the purpose of sharing an individual record for direct care to providing bulk data for other purposes. | The standard interaction patterns of synchronous query, synchronous update, asynchronous query subscriptions, and reliable messaging enable participants to transfer or disseminate data from one participant to another. |
| **TC15** | **Data Landing and integration** - Landing data sets into a local care record, such as deduplication, codification from text, data | The preference is to land data within local participating organisations and avoid a transfer of data controllership. The YHCR offers model software which helps |

| | | |
|---|---|---|
| | cleansing, conversion of data structure and codes from other standards | participating organisations and their suppliers land and normalise data.<br><br>Data can also be persisted centrally in regionally provisioned FHIR stores. The YHCR uses cloud hosted tooling to normalise data content. |
| **TC16** | **Data Processing** - Ingested data is processed to adhere to business requirements, business rules, data policies and procedures, data quality standards, this is achieved by applying series of operations/ transformations like data formatting, enrichment, cleansing, aggregation, classification, de-id/re-id etc. | Most data processing occurs locally and the YHCR provides model software to facilitate operating and transforming data. The YHCR Meta-data server offers terminology services which can be used locally or centrally for code system translation.<br><br>Where data is processed centrally cloud hosted tooling is applied. |
| **TC17** | **De-identification/Re-identification** - Ability to remove potentially patient-identifiable information from a patient's clinical record so that it can be used for secondary purposes | The YHCR use the national Privitar service for de-identification and re-identification. The service is integrated within the data acquisition tooling on the YHCR PHM Platform. |
| **TC18** | **Patient Preferences** - Patients to define their care preferences and information sharing consents, and for professionals to manage preferences and consents on their behalf | Data access control policies and patient consent preferences are recorded in et YHCR Consent Server. Patient preferences can be managed by any authorised data consumer. The Leeds Helm person record enables its uses to self-manage consent. A regionally provided user interface allows consent to be managed for section 251[14] provisions for secondary use of data. |
| **TC19** | **RBAC and Legitimate Relationships** - The services required to enable a LHCR and members of an ICS to uphold their Data Controller responsibilities. | Enforcing role based access and legitimate relationships is normally the responsibility of a data consumer and adherence to governance responsibilities is ensured during the onboarding process for data consumers. The regionally provided Care Portal is fully compliant with these requirements.<br><br>Centrally provisioned FHIR Stores use the Identity and Access Management role definitions to define scoping rules over access to data. Consistently with other functionality rules use FHIR expressions to determine the applicability of rules to data points.<br><br>Access to audit data held centrally requires users to hold appropriate privileges. |
| **TC20** | **Authorisation & Authentication** - Authorising access to APIs using the OAuth 2.0 standard to apply information sharing policies and patient consents so as to determine access permissions. Provide a capability to determine | The YHCR APIs are secured using OAuth2.0. Authentication is the responsibility of the data consumer and adherence to secure authentication standards is assured during the process for onboarding data consumers. |

[14] National Health Service Act 2006

| | | |
|---|---|---|
| | the identity of a user or service by validating their identity credentials using the OIDC and SAML identity | The regionally provided Care Portal uses OIDC for two factor authentication. |
| **TC21** | *Deleted by NHS X* | |
| **TC22** | **Care Record Access Audit** - Maintain a clear audit log for every access to a patient record within a local care record and operate a monitoring and response service. | The YHCR uses a federated approach to audit which uses FHIR resources for audit events. All interactions and data transiting through the YHCR Exchange are audited centrally. Data access is also audited by participants to the YHCR. A consolidated audit record can be assembles using standard FHIR aggregation techniques. |

## 2.5 Maturity in the YHCR

In common with most health and care economies there is significant variation in the technical maturity of the organisations of the Yorkshire & Humber region. The YHCR aims put as few barriers of entry in the way of usage as possible and this has been a design driver for the platform. The platform aims to accommodate different organisations participating at different levels of maturity but do so in a safe manner whist facilitating the most ambitious of use cases.

There are several dimensions to maturity:

- the data items offered by a data provider and the level of coverage of the medical record held by the data provider;
- compliance with industry standard FHIR profiles including adherence to coding systems;
- extent of support for the FHIR API standard and other YHCR standards;
- ability to accept inbound transactions such as service referrals;
- a data consumer's ability to interact with different types of datasets and participate in workflows;
- a data consumer's support for interaction patterns other than on-demand query.

### 2.5.1 Maturity of Data Consumers

Generally, data consumers are at liberty to mature in accordance with local needs (although as supplier of the YHCR care portal, the YHCR has a responsibility to provide the functionality demanded by its users).

Maturity of data providers is more carefully controlled.

### 2.5.2 Technical Maturity of Data Providers

Designs for technical capabilities for data providers are aligned with a maturity level in the range 1-6. Level 1 is intended to support the basic requirements of a longitudinal shared care record used for direct care. Subsequent levels support more complex requirements. Providers adopting model software provided by the YHCR have automatic support for the relevant capabilities at level 6. Provided wishing to develop their own capabilities chose the level with which they will comply. Data consumers operating above level 1 must be cognisant of compliance by the data providers with which they interact. The maturity level is available to data consumers from the Participant Registry.

### 2.5.3 Data Maturity

A similar maturity model operates for data, although the model is an evolutionary one guided by the Data Architecture Design Authority (DADA)[15]. FHIR resource content is determined by the data provider presenting it to the YHCR. The provider is encouraged to adopt:

1) FHIR profiles which are already in use in the YHCR;
2) Internationally recognised FHIR profile;
3) Develop a new FHIR profile if either of the above does not exist or the provider is unable to comply.

New FHIR profiles are endorsed by DADA (which aims to rationalise profiles in use), are added to the data maturity model, and are then available for other data providers to adopt. FHIR Profiles are registered with the YHCR Metadata server and are enforced by the YHCR Exchange. Data consumers are informed of the profile used for individual data items and can tailor functionality to specific profiles.

Using care plans as an example, this approach allows one data provider to offer basic narrative alongside another which codes goals and actions using structured data. Both explicitly reference the profile to which they are aligned, and consumers are able interpret both in a meaningful way for their users.

Working with multiple levels of maturity has challenges for data consumers. However, these are mitigated by some basic principles of profiling used in the YHCR:

- compliance with STU3 structure definitions is guaranteed and defensive coding techniques which respect the optionality and cardinality of resource attributes allows consumers to accommodate different representations of data without rigid alignment to a profile;
- profiles are associated with a maturity number and a profile at a high maturity level enhances profiles at lower levels meaning that consumers can safely operate at a low level of maturity and normally automatically accommodate higher levels;
- policies relating to providing narrative and textual descriptions of coded content provide consumers with a backstop when offering content to their users.

### 2.5.4 Data Coverage and Data Impairments

Data providers offer data that is appropriate for their care setting or system which is commensurate with clinical benefit and cost of extraction. Different models operate in the YHCR, but most data providers initially offer a core dataset and expand their offering over time.

The inevitable consequence is that all facts known about a patient may not be available to consumers at any point in time. The distinction between a medical fact being not relevant and not available is an important one for those interpreting a shared care record and the YHCR offers mechanisms which enable consumers to interpret the data that they receive.

1. FHIR *CapabilityStatements*[16] published by data providers explain the FHIR resources which are available from an endpoint.

[15] The Data Architecture Design Authority has clinical and technical representation for health and care providers, academic bodies, the Professional Standards Body (PRSB, and NHS Digital.
[16] A standard FHIR mechanism for an endpoint to publish the technical capabilities of and endpoint and the data content available for it.

2. Data providers insert Statements of Data Impairment into search results to explain any gaps in the resource types that are not available but might reasonably be expected to be provided given the *CapabilityStatement*.

Th YHCR Statements of Data Impairment[17] have been introduced because of a myriad of circumstances where, based on the content of a *CapabilityStatement,* a consumer may have a legitimate reason to expect data of a certain to be available, but for some reason the data cannot be provided meaning that there are gaps in a search result set which otherwise would be undetectable.

Statements of Data Impairment are inserted when:

- a data provider is temporarily unavailable;
- the data provider uses different systems for different healthcare service and data is available from one system but not another;
- data is only available from a certain date;
- there are known quality issues with data;
- data is being withheld due to an internal policy;
- data is being withheld due to a regionally operated consent policy and the policy authorises acknowledgement of the data's presence.

Statements of Data Impairments are implemented as *OperationOutcomes* – a FHIR resource used to inform a service user of error, warning, or information messages that result from a system action. They allow a consumer to inform users of incompleteness or quality issues and are coded using a regional coding system which facilitates standardised interpretation and machine actionability.

It should be noted that one of the key learnings from the early stages of developing the YHCR is that simplifying non-uniform data is a unique challenge for consumers of a longitudinal care record and their ability to do so distinguishes them from other systems operating in health and care.

## 2.6 Beyond the LHCRE

The LHCRE programme has established technical capabilities which lay foundations for interoperability within and outside of the region. Early day ambitions have targeted the assembly of a longitudinal shared care record. However, the potential reach and scope of the platform are very much broader.

By emphasising the mechanisms for safely sharing data, the YHCR is creating a data democracy where all systems have equal rights, within legal boundaries, to access data regardless of its provenance. There is early evidence that data democracy encourages entrepreneurship by removing barriers to new solutions to health and care problems. Systems can better operate across care settings with focus on the care pathway rather than organisational constraints.

Data democracy is also a prerequisite for the introduction of information led cognitive applications which use machine learning and advanced algorithms to provide clinical decision support, predictive diagnostics, and care automation. The Yorkshire & Humber region is well positioned to act as an innovation laboratory for this new generation of software tooling.

The open principles of the YHCR extend beyond open standards. The programme has adopted openness in its documentation and has built an open platform which is free of software licenses and where the intellectual property is held by the NHS. The cloud hosted interoperability platform is

[17] Reference design paper 017 – "Data Impairments – Interpretation and Reporting"

readily deployable outside of the YHCR and is at the time of writing of this edition, the platform is being used in two other regions. Being built for the cloud, the platform scales near linearly and can be economically applied to local integration solution, regional share care records and for connecting shared care records to provide supra-regional interoperability.

# 3 Regional and Model Component Specifications

Regional infrastructure has been designed to allow data consumers and providers to interact securely, with the consent of the person to whom the data relates (where applicable, and in accordance with local governance policies). Regional infrastructure also centralises services which provide functions that are commonly used by data providers and consumers.

Model components are software which is intended to be developed regionally and made available to YHCR contributors for local tailoring and deployment. Model components are intended to help local organisations fast-track to compliance with regional interoperability standards.

The YHCR Exchange includes a clinical data repository. The YHCR is a hybrid architecture which leans towards federation for data persistence but supports the centralisation of data where data is not naturally controlled by any single contributing organisation or where there is a valid use case for holding data centrally, including data which is necessary for the operation of the YHCR.

Exchange components make use of the clinical data repository to manage FHIR resources such as expressions of consent and a master patient index. Exchange components which manage these data items are, in many cases, no more than a FHIR service endpoint over a FHIR repository and could be specified simply in terms of resource definitions. However, in the interest of clarity of operation, these components have been separately identified and are described as functionally independent entities.

## 3.1 Components of the YHCR Exchange

### 3.1.1 Identity and Access Management Server (IAM)

IAM[18] is the key-holder for regional data. It authorises users and applications to access services on the YHCR and issues signed tokens (JSON Web Token) which can be trusted by data providers as proof of authorisation.

**Authentication v. Authorization**

A principle of the architecture is that authentication is federated: users are authenticated locally, the authentication is 'trusted' by the regional infrastructure and by data providers. Direct authentication with the regional infrastructure is only required where there is not the technical capability to exchange the data necessary to establish trust or if users are interacting with regional services through a regionally provided application.

The regional Care Portal, as a data consumer, offers several authentication mechanisms. It is deployed in tenancies which allow the most appropriate authentication method to be used for the group of users that the tenancy supports. Care Portal authentication options include:

- Multifactor Authentication – a simple username and password with a code sent over SMS.
- OpenID Connect.
- Authentication for embedded portal in EMIS – a proprietary SAML based approach.
- Authentication for embedded portal in TPP through user details passed in the launch URL.

The following methods are planned and will be deployed in accordance with use case:

[18] Reference design paper 005 – "Identity and Access Management"

- Active Directory – a Kerberos token over HTTPS validated with a Microsoft Domain Controller local to the connecting organisation.
- SAML - Digitally Signed XML over SOAP.

**Authorization Claims and Bearer Tokens**

An application authenticating a user gains access to regional infrastructure and local data providers by making an OAuth2 claim against IAM. The claim includes details of:

- the user authenticated by the local application;
- the organisation which the user represents;
- the type of user;
- their reason for accessing regional services;
- the patient that is in context in the host application (if appropriate).

If the claim is successful, then IAM returns a bearer token: a signed JSON Web Token (JWT). The signatory is the YHCR and is trusted within the YHCR Public Key Infrastructure. Authorisation attempts are logged using the regional Audit Service.

The JWT carries all information which was presented in the original claim and this information can be used by regional infrastructure and data providers to limit the scope of access (the data to which a user has access) or to tailor the presentation of data to particular purpose.

**Access Scope Control**

Fundamental to these operations is the reason for access and type of user. The YHCR defines the following options:

| Code | Reason for Access |
|---|---|
| 1.1 | Direct care (Emergency). Access is in the context of a patient; |
| 1.2 | Direct care (Non-emergency). Access is in the context of a patient. |
| 2 | Indirect care with the consent of the patient. Access is in the context of the patient. |
| 3 | Indirect care not in the context of a patient. (Not patient-centric). |
| 4 | Analytics with access restricted to pseudonymised data. (Not patient-centric). |
| 5 | Administration (Not patient-centric). |

| Code | Type of User |
|---|---|
| 1 | Clinical Professional. |
| 2 | Social Care Professional, |
| 3 | Citizen. |
| 4 | System or Robot. |
| 5 | Administrator (of YHCR systems). |
| 6 | Auditor. |
| 7 | Authorised Carer. |

Scope rules imposed by central infrastructure include:

- direct care must be in the context of the patient and searches for patient-centric resources must include a patient subject in search terms;
- an auditor only has access to *AuditEvent* resources;
- access for administration does not permit access to patient-centric resources.

Regional infrastructure includes a Consent and Data Access Control Server which uses contextual data provided in the access claim to implement policies which restrict or enable access to data.

Note that data providers may alter the presentation of data depending on information in the access claim. For instance, a response to a data request made by a Citizen may redact clinical detail that would be returned had the request been made by a Clinical Professional.

**Token Expiration**

JWTs expire. The expiry time is a property of the signed JSON object. It is the data provider's responsibility to ensure that the JWT is active. Data consumers should refresh a JWT with IAM that is near its expiry time.

Continuous compliance testing software monitors data provider adherence to security policies.

### 3.1.2 FHIR Aggregator

The FHIR Aggregator[19] presents a single endpoint from which data consumers can interact with any data provider (including regional FHIR Stores). Theoretically, data consumers can have a direct relationship with data providers and bypass the regional aggregation capability: once a JWT has been obtained from IAM then this could be used to address an individual data provider and subject to security arrangements interact directly with data managed by that data provider. Practically, however, most interactions with data are through the regional FHIR Aggregator.

**Simplification of Relations**

The FHIR Aggregator simplifies relationships between data consumers and data providers:

- it provides a single point of access and avoids consumers 'hunting' for data from a variety of providers;
- it provides a comprehensive implementation of the FHIR STU3 API standard and abstracts consumers from mediations with less capable data providers;
- it filters search results for contraventions of consent or other data access policies;
- it unifies data around certain common concepts and presents data as though it had originated from a single source.

**A Unified Data model**

Data unification[20] is the virtualisation of a longitudinal care record. Data consumers access a data model that is relationally integral around concepts which are common to all data providers. These currently comprise:

- patients;
- practitioners;
- organisations.

[19] Reference design paper 010 – "FHIR Aggregator Service"
[20] Reference design paper 001 – "A Unified Distributed Data Model for FHIR"

The FHIR Aggregator replaces references in search results to local versions of these concepts with references to regionally maintained master data. A data consumer requesting, say, blood pressure measurements, receives a result set where every blood pressure measurement references the same patient. Had they been obtained separately then each data provider would have referenced their own version of the patient resource.

There are further opportunities to expand on the unification concepts. Work is taking place on regional definitions of Locations which are commonly used between different care providers. A regional dictionary of Medications will probably be pursued in the near term.

**Synchronisation of Regional Master Data with National Services**

Regional master data is synchronised with the national Personal Demographic Service[21] (PDS),the Spine Directory Service[22] (SDS) and the Organisational Data Service[23] (ODS). The YHCR spotlights patients where there is current activity in the region and for these synchronises patient demographics with PDS with the aim of maintaining a record which is no more than 24 hours out-of-date.

PDS synchronisation will be simplified when the National Event Management Service (NEMS) fulfils its goal of offering event notifications for patient records which have changed on PDS.

**Implementation of the FHIR API Standard**

the FHIR Aggregator offers a comprehensive implementation of the FHIR API standard allowing data to be read, searched, created, and modified. The YHCR has avoided extending the FHIR standard where extensions would impact compatibility with other systems' implementation of FHIR. Any system which is fully compatible with FHIR STU3 will be able to interact with the YHCR out-of-the-box.

Extensions are limited to:

- modelling of platform configurations as resources defined using the FHIR Conformance model thus allowing the platform to be managed via the FHIR API;
- modelling of data access policies as FHIR type resources;
- extending FHIR search syntax to enable searches to be executed against patient cohorts.

The YHCR uses FHIR resources in its implementation as follows:

- *OperationOutcome* resources are used to communicate data impairments and other status information to data consumers;
- *AuditEvent* resources are used to record an audit record of interactions with the YHCR;
- *Linkage* resources are used in the Master Patient Index to cross-reference a regional master patient record to local patient records held by data providers;
- Meta data tags in the base *Resource* class are used to identify the provenance of data.

The YHCR publishes profiles for its implementation of these resource types.

[21] Reference design paper 011 – "Interfaces with the Personal Demographic Service"
[22] Reference design paper 012 – "Interfaces with the Spine Directory Service"
[23] Reference design paper 013 – "Interfaces with the Organisational Data Service"

**Data Discovery**

The FHIR Aggregator determines how to route interactions from data consumers to data providers. Patient centric searches and subscriptions are routed to those data providers who have had contact with the patient. For most providers this is determined through an encounter being recorded with the PIX service (section 3.1.6). There are a class of data providers, which are typically aggregators themselves, such as other Shared Care Records (ShCR), where patient relationships are discovered dynamically.

Data consumers can also explicitly route interactions to specific providers either be specifying the providers or governing organisations in search terms or by implying a provider by using local patient, organisation or practitioner references.

The FHIR Aggregator rebases references to regional resources in search terms received from data consumers to equivalent local references. This abstracts data providers from the regional unified data model and allows them to operate solely with local data identifiers.

**Interaction Patterns**

The FHIR Aggregator implements the interaction patterns detailed in section 2.1.

For synchronous patterns the Aggregator echo's the interaction from the data consumer onto selected providers and returns results in real-time. The Aggregator manages page caches, deduplicates, unifies regional concepts and packages results in a single result set.

For asynchronous queries, subscriptions and reliable messaging, the Aggregator acts as an agent for participants, placing queries and subscriptions, collecting results, and queuing notifications and messages.

**Performance Guarantees**

A federated architecture performs at the speed of its slowest participant. There is a risk that one poorly performing data provider could render the service unusable for all participants. The FHIR Aggregator mitigates this by offering performance guarantees to data consumers and enforcing performance contracts with data providers.

Connections to providers which do not respond within a contracted period are terminated and a potentially incomplete result set returned to the data consumer alongside details of the suspended connections.

Consumers choose whether to provide a partial result to users or override the guarantee and wait for a full response. A sophisticated user interface might display the partial results and target poorly performing providers with separate queries augmenting the display as additional results arrive.

### 3.1.3 Consent and Data Access Control Server[24]

Consent is not required to share data for the purposes of direct care. Consent is required to share data for other purposes such as data provided by the person, data not captured in the course of direct care, or data not being used for the purpose of direct care.

The YHCR models consent using standard FHIR *Consent* resources. FHIR allows consent to be applied to a policy but is not specific as to what constitutes a policy. The YHCR has extended FHIR with the

[24] Reference design paper 008 – "Data Access and Consent Management"

*Policy* resource. A FHIR *Policy* operates in a context (why the record is being accessed, who is accessing the record and where from) and specifies the data items covered by the policy. The YHCR *Policy* resource uses FHIR expressions to define data coverage and so very granular policies can be constructed.

Consent policies are defined at a national and regional level with the possibility of local variation. Policies with are defined regionally or nationally are enforced by the regional infrastructure. Policies can also be used to control the flow of data for other reasons (such as cessation of a care relationship with a patient) and so more accurately termed Consent and Data Access Control Policies.

*Consent* and *Policy* resources are managed using the same approach as other regional FHIR resources using the regional aggregator's FHIR API.

A Consent and Data Access Control server enforces policies. It tests all data items flowing through the FHIR Aggregator against relevant policies and data items withheld or qualifying statements inserted into search results which inform a data consumer of presence of sensitive data.

### 3.1.4 Audit Service

The audit service[25] is used by regional services to log all interactions with data consumers. The audit service may optionally be used by data consumers to log interactions with users and by data providers to log events pertaining to the supply of data. If a data consumer or provider chooses not to use the regional audit service for event logging, then they must log equivalent information locally and make the data available for regional interrogation via a data provider FHIR endpoint.

Audit events are recorded as FHIR *AuditEvent* resources.

A regional vocabulary is defined for audit events.

*AuditEvent* resources are managed similarly to other FHIR resources except modification is not permitted.

### 3.1.5 Regional Data Repository

Regional data is persisted in as FHIR resources in central FHIR stores[26]. It is acknowledged that FHIR was developed for healthcare and resources definitions do not necessarily cover all concepts which need to be persisted. In particular, concepts relating to Local Authority provided care do not map well to current FHIR resources. The DADA, works with NHS Digital and the Professional Records Standards Board to ensure that nationally recognised profiles and extensions to the FHIR standard are defined which accommodate the needs of all care providers.

Interactions with the Regional Data Repository

Capabilities include:

- Version controlled FHIR resources in XML and JSON format.
- All resources defined by STU3 (or new resource types defined by the YHCR using the FHIR conformance model) can be managed.

[25] Reference design paper 009 – "Auditing"

[26] The first edition of this document presented OpenEHR as alternative to FHIR for persistence. Work with the OpenEHR community to embed FHIR mappings in the OpenEHR archetype drafting process has not yet progressed to a stage that would allow a hybrid persistence layer to exist.

- All search parameters defined by STU3 can be referenced in search terms with full operator support.
- Search directives such as including referenced resources in bundles and pagination of results.
- Selective resource modification.
- Resource subscriptions (a user or an endpoint can be notified of a change to a resource).

Resource management activities are restricted by user, role, organisation, and team membership. Management restrictions are expressed as extended FHIR search terms. Extensions enable restrictions to be defined by, for instance, cohort of patients.

### 3.1.6 PIX/MPI Server

The regional PIX server[27] (Patient Identifier Cross Reference Manager) collates demographic and episodic information from data providers and manages regional FHIR *Patient*[28] resources. The role of PIX is to

i) Track the care settings that have had contact with a patient (so allowing requests for data to be targeted at only those data providers with a legitimate relationship with a patient).
ii) Populate a regional Master Patient Index with regional patient demographic and links to local identifiers and demographics.

PIX is initially seeded when a data provider onboards to YHCR with patient demographics for all patients registered with the provider. A batch data on-take service is available for this purpose.

PIX can accept an HL7v2 feed as a source of demographics. HL7 is well established in healthcare in contrast with Local Authorities where its use is sparse. A FHIR based interface is available for those to use who do not have HL7 capability.

PIX uses a traced NHS Number is as the single reliable indicator of patient identity[29].

If a registration is received from a data provider with a traced NHS number, then the person identification data is used to modify an existing patient resource which has the same NHS number. If there is no existing patient resource, then one is created. FHIR *Linkage* resources are used to cross references the master record for a patient with local identifiers.

If registration is received with an untraced NHS number, then PIX attempts to trace it using the national Spine Mini Services Provider. If the trace succeeds, then the data is treated as above. If not, then it is treated as a registration with no NHS number.

Since the first edition of this document was written, a new class of data provider has emerged[30] which have much more complex behaviour than the types of data provider originally envisaged. These more complex providers typically present data from a number of different governing organisations and offer data through proprietary APIs which need to be adapted to meet the YHCR

[27] Reference design paper 004 – "Patient Identity Exchange"
[28] The first edition of this document envisaged PIX as both a patient and encounter registration service. A use case has not yet transpired for a regional representation of an encounter although some suggestions for future models of care administrations promote regional episodes of care being managed which envelope the continuum of services provided across all care settings to a address a patient condition.
[29] Work is beginning in Wales where there is much lower coverage of the NHS Number, to introduce probabilistic matching to PIX. This functionality with provided by the national Welsh MPI.
[30] Reference design paper 029 – "Data Aggregators and Other Complex Data Providers"

standards. Like the YHCR itself, these data providers often work on a "discovery basis" for locating data and do not pre-publish relationships with patients to their consumers. PIX has developed features for interacting with these data providers and *Linkages* from a regional demographic can be created on demand by invoking discovery APIs on the complex provider.

### 3.1.7 Subscriptions Manager

The Subscriptions Manager[31] mediates between data consumers subscribing to events in the YHCR and data providers notifying subscribers that events have occurred.

A data consumer typically creates a subscription with the YHCR Exchange, and the Subscription Manager places the subscription with relevant data providers. If a subscription is for data about a specified patient, then it is only placed with those providers who have registered contact with the patient in PIX. Providers registering new contact with patients may receive subscriptions which historically been placed on the YHCR Exchange for the newly registered patients.

The same subscription registered by multiple consumers will only be placed with providers once.

Providers send notifications of events to the YHCR Exchange. The subscription manager forwards notifications to all consumers placing the subscription.

### 3.1.8 Document Sharing in the YHCR

The YHCR enables documents[32] and media to be exchanged in a similar manner to structured FHIR data. Data providers publish metadata about documents using FHIR *DocumentReference* resources.

The *DocumentReference* provides information about the document's subject, author, document type and title and can be queried through the FHIR Aggregator as other patient centric FHIR resource.

Whilst *DocumentReferences* can embed a document (and there are instances in the YHCR where this practise is adopted) it is preferred that the *DocumentReferences* references the document they describe by URL and the document is obtainable separately from its metadata. This approach allows consumers to offer a list of documents available to their users and to retrieve the bulk of data only when the user chooses to view one of the documents.

The FHIR Aggregator rebases all URLs referenced in *DocumentReferences* so that they point to a regional endpoint. The endpoint acts an HTTPS proxy to document endpoints operated by data providers so simplifying firewall management and securing access to documents: a document can only be retrieved by a user that was authorized to access the *DocumentReference* resource that describes it.

The YHCR is compatible with IHE type Health Information Exchanges which use XDS for data sharing. The FHIR *DocumentReference* is broadly equivalent to the XDS DocumentEntry. The FHIR Aggregator could act as virtualised[33] XDS affinity domain which, over an XCA bridge, would be a source of documents inter-LHCR document sharing.

[31] Reference design paper 007 – "Subscription Infrastructure"
[32] Reference design paper 019 – "Document Management"
[33] The first edition of this document proposed centralisation of an XDS document registry and direct support for the XDS SOAP action for registry management, document submission and retrieval. FHIR DocumentReferences could still be persisted centrally and XDS could still be offered as a façade to the YHCR but a use case has not yet been identified.

Similarly, an XDS exchange which supported the IHE profile MHD could act as a data provider to the YHCR.

### 3.1.9 Pseudonymization Service

The regional pseudonymisation service implemented as part of the PHM platform generates anonymous *Patient* resources which correspond to a real *Patient* resource and can be reverse engineered to the original patient in appropriately secured conditions. The service uses the national Privitar solution for pseudonymisation.

The service relies on a PKI to ensure that pseudonymized data is private to the organisation requesting pseudonymisation. The original identity of a pseudonymised patient can only be established by the holder of the private key of the party requesting a pseudonymized resource.

The regional service, with participation of the private key holder, will allow a resource to be re-pseudonymized for use by another party.

### 3.1.10 Metadata Services[34]

The metadata server[35] is a repository for FHIR profiles and coding systems that allow the YHCR to unambiguously clinical concepts using FHIR STU3 as a base. The Metadata Server also provides tooling to allow FHIR profiles to be approved and compiled (a process for collating a sequence of structure differentials into a self-contained snapshot). Resources can be validated against approved profiles using an implementation of the $validate FHIR operation.

Validation is under configuration control and be performed for all resources passing through the FHIR Aggregator or selectively for certain data providers.

By default, resources which fail validation are withheld from data consumers. Consumers can elect to receive invalid resources and, in this case, will receive a statement of non-compliance in bundles that contain them.

### 3.1.11 Connectivity and Security Services

Regional services are available over the open Internet. The YHCR has connectivity to HSCN and where a case can be provided then consideration is given to connecting with a participant over the HSCN/PSN network .

The YHCR operates the yhcr.nhs.net subdomain and associated domain name server. All data provider and data consumer endpoints have addresses on the subdomain and these are used in the organisation's IAM registration.

YHCR operates a certificating authority. All endpoints connecting to regional services must be secured by a YHCR signed certificate.

### 3.1.12 Regional User Interface Componentry

Central user interface servers are provided as follows:

[34] Reference design paper 030 – "FHIR Metadata Management"

[35] Emphasis changed during the course of the LHCRE programme from terminology services to metadata services. The YHCR recognises the value of terminology services, particularly in enabling its PHM platform to work more flexible with SNOMED-CT and will be providing an implementation of the FHIR operations for terminology services.

- canvas server;
- YHCR panel library;
- panel server.

**Canvas Server**

The canvas server is the endpoint from which the YHCR Care Portal is served. The canvas server constructs a user interface from a configuration, known in the YHCR as a tenancy, which provides authentication, navigation, and data source selection services. The canvas draws functionality from panels registered in the panel library.

**YHCR Panel Library**

The panel library contains metadata about panels used in the YHCR including contextual requirements and details of server from which panels can be accessed. The canvas validates consistency of tenancy configurations with the panel library and It operates as a governance tool, allowing the region to assure panels in use in the region.

**Panel Server**

The source of panels registered in the panel library. The canvas server acts as a proxy for web browsers loading panels from panels servers and the panels' interaction with the FHIR Aggregator. This arrangement simplifies firewall configuration and provides for central auditing of all data interactions.

The open user interface architecture affords no particular privilege to the canvas server or panel server; other user interface canvases and locally developed panels can be served alongside the regional canvas and regional panels.

The YHCR panel registry is also substitutable by local alternatives but in the YHCR allows clinical safety and technical assurance process to be operated centrally and canvas adopters which use the regional library have confidence that the panels registered with it will operate harmoniously and provide a consistent user experience.

## 3.2 Model Components

Model components are offered as regionally developed software which may be adopted, extended, and deployed as local infrastructure at the option of the local organisation.

### 3.2.1 FHIR Proxy[36]

The FHIR Proxy is a reference implementation[37] of a FHIR service bus which complies with FHIR technical standards in section 7. Versions of the FHIR Proxy server are available for InterSystems Ensemble/HealthShare and Node.js.

The FHIR Proxy server facilitates the process of FHIR enabling existing data sources by offering the following features:

[36] The FHIR Proxy is increasingly being called a FHIR Appliance. The change of name reflects the diversity of usage of the component which often involves a degree of data processing that is not reflected in the original name. This document retains the original terminology for consistency with the first edition of the paper.
[37] Reference design paper 003 – "Conceptual Design for a FHIR Proxy Server"

- A FHIR compliant endpoint with full searching capabilities.
- FHIR Aggregation.
- Support for synchronous, asynchronous and subscription interaction patterns.
- Disambiguation of physical concepts to allow patients, organisations, practitioners, and locations to be represented as single resources regardless of the data source referencing them.
- Management of FHIR versions.
- A proxy layer which facilitates mapping of FHIR resource properties to data elements in source systems.

The FHIR Proxy is represented figuratively below:

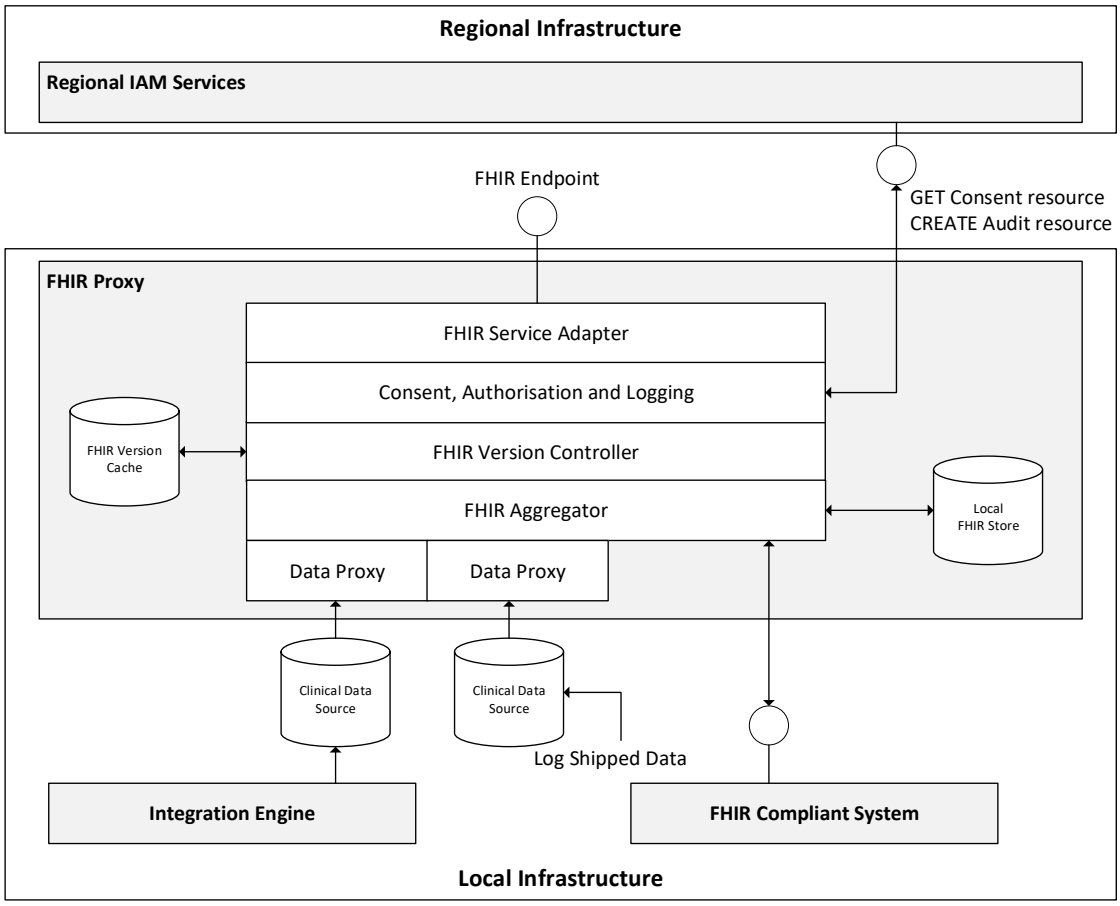


The FHIR proxy is open source and can be used without license fee.

### 3.2.2 Document Proxy

Whilst technically a capability of the FHIR Proxy, the document proxy provides useful functionality that is often independently deployed and deserves individual attention. The YHCR implements a document sharing architecture which is based on FHIR and is compatible with the IHE (Integrating the Healthcare Enterprise) MHD (Mobile Healthcare Document) standard.

The document provides out-of-the-box compliance with the standards and offers data providers a imply mechanism for offering documents to the YHCR. The proxy implements HL7v2 processing capabilities and will extract documents and prepare FHIR *DocumentReference* resources from HL7v2 messages.

### 3.2.3 Regional Canvas and Panels (aka the YHCR Clinical Portal)

The regional canvas provides a quick route to organisations becoming a participant in the YHCR as a data consumer. An organisation can configure which panels are available through the canvas for use by their care professionals. Organisations using the canvas control their own definition of roles and can tailor panels by role.

The regional canvas implements:

- methods for authenticating users including two factor authentication, federated authentication through OIDC, integration with Active Directory, and context launching from other applications;
- a menu system that allows users to navigate to panels;
- a system for filtering data providers represented in aggregated data;
- support for interpreting and users accessing statements of data impairment.

The regional canvas is tied to the YHCR panel library.

Regional panels have been designed to meet common universal requirements for interacting with the YHCR:

- a patient banner with drill down into patient demographics, care team contact and alternative contact methods;
- a locally configurable alerts banner which highlights important facts and new information available about the patient;
- a timeline of contact providing a visualisation of historical encounters with a patient;
- locally configurable summary panels[38] which blend data from different sources and provide drill down to tabular and graphical renditions of data;
- locally configurable detail panels[39] which list all instances of a FHIR resource in a searchable, sortable, and paginated table;
- a document viewer;
- an application for referring to a service or booking an appointment with a service operating withing the YHCR;
- an application for receiving a service referral and managing appointment booking slots.

The service referral, appointment booking and receiving applications are compatible with the national ITK based method for service referral and Care Connect profiles for appointment booking and so are compatible with the national Emergency Department appointment booking system EDDI and its future replacements. The service referring application has been accredited by NHS Digital the receiving application will be accredited once a use case has been determined.

Local requirements can be met through local panel development. Locally provided panels can be served from local panel servers and source code does not need to be made available for local and regional panels to be col-located within the regional canvas.

[38] Summary panels are bound to a FHIR resource type and can be used to blend data from any FHIR resource. They are intended to de-duplicate data from different sources. The definition of the resource content that defines equivalence for de-duplication purposes is in the hands of the implementor.

[39] Detail panels are bound to a FHIR resource type and can be used to list data from any FHIR resource. FHIR resource attributes displayed in a detail view are under local configuration control.

### 3.2.4 Regional Library for UI Componentisation

The standards which are used to allow panels to collocate within a canvas and to share contextual data are defined within a regional software library for user interface development. The library is written in JavaScript and can be used with most user interface development technologies to create local canvases and panels. The library is open source and can be used without license fee.

## 3.3 Participant Onboarding

### 3.3.1 Environments of the YHCR

The regional components of the YHCR are cloud hosted and are available within three publicly facing environments[40]:

- **A sandpit environment** which is easy to access, has reduced security, and provides access solely to test data is offered to facilitate participants developing compatible data provision or consumption capabilities. The YHCR offers data providers in the sandpit environment a rich set of test data which is representative of the FHIR profiles operating in the region./ Automated test tooling automatically runs against new data providers entering sandpit and provides evidence of a data provider's compliance with standards and expected levels of capability.
- **A staging environment** is a step through to live operation that all participants must pass through. Permanent connections to test instances of data providers allow data consumers to validate that they are compatible with the type of data found in the live environment. New data provision can be tested with data consumers to ensure that new data can be interpreted safely.
- **A live environment** hosts live operations.

### 3.3.2 Onboarding Suite[41]

Participants transition through the environments under software control. All participants have access to an onboarding suite that makes the onboarding process self-service under regional governance control. The suite allows participants to self-register endpoints, issue certificate signing requests and declare capabilities.

The onboarding suite also controls the security configuration of the regional infrastructure by:

- managing entries in participant registry;
- managing firewall rules;
- setting up DNS entries;
- registering certificate.

[40] Other environments are available for development and testing

[41] Reference design papers 020 –"Onboarding for Data Providers" and design paper 021 – "Onboarding for Data Consumers"

# 4 Requirements for Data Providers

## 4.1 FHIR Service Point

Every data provider must offer a FHIR service point which is compliant with the elements of the FHIR API specification identified by the YHCR as being essential for enabling data to be acquired on demand. The YHCR offers a maturity model for technical compliance, and this is outlined in section 7. At the most basic level data providers are expected to support API syntax which allows atomic FHIR resources to be queried using their identifiers, the resources' subject, and the subject identifier.

Data providers adopting the regional FHIR proxy fully comply with all levels of the technical maturity model.

The FHIR resource content available varies between data providers. Most providers offer a core dataset (*Patients* and *Encounters*) which they mature with increasingly clinical content over time.

The PRSB offers a definition of a shared care record which guides data provision in most ICSs.

| PRSB Headings and FHIR Resource Content | |
|---|---|
| **About Me**<br>*Questionnaire*, *QuestionaireResponse* | **Safeguarding & Risks**<br>*Flag, RiskAssessment* |
| **Demographics and Contacts**<br>*Patient, Person, Practitioner, CareTeam, RelatedPerson, Group* | **Formulation**<br>*DocumentReferenence, ClinicalImpression, CarePlan* |
| **Legal Information**<br>*Consent, DocumentReference* | **Medications & Allergies**<br>*AllergyIntollerance, MedicationStatement, MedicationRequest, MedicationAdministration, MedicationDispense* |
| **Social Context**<br>*Observation, ClinicalImpression* | **Plan & Requested Actions**<br>*CarePlan* |
| **Investigations, Examinations and Assessments**<br>*ClinicalImpression, Observation, Procedure, Questionnaire, QuestionnaireResponse* | **Vaccinations**<br>*Immunization* |
| **Family History**<br>*FamilyMemberHistory* | **Participation In Research**<br>*Group, ResearchStudy, ResearchSubject* |
| **Problems, Diagnoses, Conditions & Procedures**<br>*Condition, DiagnosticReport, Procedure, ClinicalImpression* | **Care Plans**<br>*CarePlan* |
| **Pregnancy Status**<br>*Observation* | **Alerts**<br>*Flag* |
| | **Documents**<br>*DocumentReference* |

The suggested FHIR mappings is that of the YHCR and is not prescriptive. The DADA evolves standards for representing clinical concepts as requirements are addressed.

Data providers must be able to serve a *CapabilityStatement*[42]

The data provider may offer secondary service points to serve documents and receive asynchronous requests for bulk data.

In accordance with section 2.5.4, when responding to FHIR searches a data provider knowingly returns an incomplete result set or data with known quality issues then the provider must insert a statement of data impairment which is coded as a FHIR *OperationOutcome* resource. There could be a number of reasons such as a data source not being available, incomplete mappings to a regional coding system, or only partial coverage of across a number of different services. The YHCR maintains a coding system which categorises data impairments.

## 4.2 Resource Management

Resource management covers the creation and modification of resources managed by a data provider using the providers FHIR service point. The capability focuses on the management of resources which allow clinical processes to be executed across care settings such as the management of referrals, appointments, and tasks. This will allow carers from different organisation to co-ordinate the care of their patients better by permitting electronic access to services regardless of where the service is provided from. For most data providers this is an advanced capability and initially data is offered only on a read-only basis.

Where resource management support is offered, the YHCR does not presume how an organisation responds to resources created or modified at its boundary or the level of integration that the organisation has between the FHIR endpoint and core patient or client administration systems. It is possible that resources created or updated from outside the organisation are manually administered. It is also possible that integrations with core record systems allow external management activities to automatically be reflected in appointment books, diaries, clinic lists etc.

Note that the YHCR is not intended as a mechanism for managing Electronic Patient Record from outside of the boundary of an organisation. The regional clinical data repository is available to create and manage clinical resources which are used between care settings.

A regional security model defines restrictions on the management of resources.

## 4.3 Subscriptions

Subscriptions allow data points to be monitored and a subscriber to be informed of a resource change. Subscriptions are expressed using FHIR search terms and specify a delivery mechanism for subscribed resources.

Whilst not an entry-level pre-requisite, the ability to accept and respond to subscriptions is a key enabler of new models of care and data providers are strongly encouraged to support this interaction pattern.

Users of the regional FHIR proxy have automatic support for subscriptions.

## 4.4 Asynchronous Query

[42] A machine interpretable of the technical capabilities and FHIR resources offered by an endpoint.

Data provider may optionally offer support for the asynchronous query pattern. This becomes mandatory if the provider is to supply data to the PHM platform.

Queries are processed in identical manner to the synchronous pattern but may be executed out of hours or against an offline copy of a database. The YHCR Exchange polls the status of queries and collects results when available.

## 4.5 Compositions and Documents

The goal of the YHCR is to assemble a comprehensive longitudinal care record that is fully structured, i.e.: where key facts are represented as attributes of FHIR resources and, where appropriate, coded to internationally recognised terminology (SNOMED-CT or DM+D). Both machine and human will be able to process the record unambiguously and it will be a rich source of data for research and service planning. However, practically speaking, much of the record currently exists in textual form as documents and the YHCR encourages data providers to offer these alongside whatever structured data they can surface.

The YHCR document architecture uses FHIR *DocumentReference* resources to provide meta-data about documents. Data providers are asked to offer the meta-data separately from the actual documents so that data consumers are able to ascertain the content that is available before downloading sizable files. This approach is supported by FHIR which allows an "attachment" to a *DocumentReference* to be provided as a URL.

It should be noted that documents also include images, video, and audio. These media types will continue be an important constituent of the longitudinal care record even when most legacy documents have been replaced with structured content.

FHIR also has a standard for structured documents which is relevant to the YHCR.

Compositions are a FHIR mechanism for packaging related atomic FHIR resources to offer a snapshot of a medical record at a point in time. Nationally, compositions are defined for transfers of care from inpatient stays (acute and mental health) and emergency departments. The YHCR has developed a regional standard for compositions for transfer of care from an ambulance to emergency departments.

Some patient record systems offer compositions in lieu of search capabilities against atomic FHIR resources. Whilst this approach can be accommodated by the YHCR, our preference is to decompose these records into searchable resources which allows them to be blended with resources from other care settings.

## 4.6 PIX Registration

Records presented by a data provider to the YHCR will only be included in resource requests for direct care purposes if the subject patient is registered by the data provider with the regional PIX service.

The YHCR offers an HL7 endpoint and for providers which process HL7 messages, a simple solution to regional patient registration is often to forward messages representing local patient registrations, demographic updates, and encounters to PIX.

An alternative mechanism is post *Patient* resources to PIX as new patients are created in local FHIR stores. The regional FHIR proxy offers this mechanism out-of-the-box.

Wherever possible patient identifier lists in HL7 messages or FHIR resources should include a traced NHS number.

## 4.7 Governance[43]

All service endpoints must validate the signature and expiry time of the regionally allocated JWT. Requests with an invalid JWT must be rejected.

All service endpoints must log all interaction either using the regional Audit Service or log equivalent data locally. Local audit logs should be searchable on the FHIR service point as *AuditEvent* resources.

Data providers must comply with the regional PKI and validate HTTPS connections are secured using regionally signed certificates.

Data provider endpoints are continuously probed for compliance with security and audit standards by regionally provisioned compliance monitoring tools.

Data providers may choose to withhold resources based on the context of a request. Full details of the claim made to the YHCR to gain access to regional records are carried in the JWT and these could be used by a data provider to control access to certain data.

Patient consent is enforced centrally[44], and filters can be setup in the FHIR Aggregator for individual providers to block access to data for any contextual reason[45] which is derivable from the OAuth2 claim. Pairings configurable in the regional participant registry can restrict access to certain data consumers. These mechanisms can be used to implement local data sharing agreements and if sufficient then no further work is required by a data provider to control scope of access.

## 4.8 Use of Adapters in the YHCR

In circumstances where a system has native support for FHIR but does not fully comply with the YHCR standards then a regional adapter is placed between central infrastructure and the data source. The nature of the adapter depends on the gap between native capability and the YHCR requirement, but examples include:

- uplifting the security of a connection to ensure that certificates are mutually authenticated;
- decomposing compositions into searchable atomic FHIR resources.

[43] Reference design paper 014 – "Governance for Data Providers"
[44] The first edition of this document placed much more responsibility on data providers for scope control and enforcement of consent.
[45] Reference design paper 031 – "Data Release Management"

# 5 Requirements for Data Consumers

Data consumption[46] is not limited to end-user applications with a user interface. For instance, a data consumer might be a software robot that automates some aspect of patient care such as auto-enrollment onto a care pathway or an algorithm that assesses the need for intervention in a patient's care.

Whilst the legal basis for sharing data might differ and consent may be required, the technical mechanics for accessing data are broadly the same whatever the nature of consumption. This section summarises requirements from the perspective of a clinical portal with exceptions noted as they apply to other types of data consumer.

## 5.1 User Authentication

End-user applications that access regional data must authenticate their users and allocate them consistent identifiers that are used across application sessions. When it is determined that a user requires access to regional data then the application must make a claim to IAM using the OAuth2 protocol.

The claim details:

- the issue time of the claim;
- the application making the claim (this must be a registered data consumer);
- the ODS code of the organisation issuing the claim;
- the patient that is the subject of the access request (for patient centric reasons for access);
- the reason for the access request;
- details of the end user.

End-user details include an identifier. Local identifiers are permitted but one of the following recognised sources are preferred:

- Electronic Staff Record number.
- Organisational Data Set code for practitioners.
- Spine Directory Service identifier.
- NHS number.
- National Insurance Number.

The user identifier is used by the regional audit service to log interactions by the user with regional data. To enable traceability, the local application must log authentication and logout and, in the log, associate the user identifier with name of the user, role being performed and the bearer token issued by the regional OAuth2 service. Logs must be retained online for a period and thereafter archived.

Non-user orientated systems accessing regional resources must authorise access with IAM using a system identifier and role which indicates that data is being acquired for system processing. If the system caches data which is subsequently accessed through an end-user application, then this application must log its authentication and record audit events which correspond to end-user

[46] A summary of governance requirements is presented in design paper 015 – "Governance for Data Consumers"

resource access either with the regional audit service or with a local audit service. An implication of this requirement is that data obtained regionally and cached locally must continue to be associated with the regional resource identifier from which it was derived.

## 5.2 Interactions with IAM

All interactions with IAM are over HTTPS. The connection is secured using a certificate signed by the regional certificating authority. The endpoint connecting to IAM must be registered with the YHCR. This may require end-user applications to interact with IAM through a proxy server.

The IAM OAuth2 service for authorisation uses a PKI which is based on certificates signed by the regional certificating authority: user identification data signed by a trusted source enables data providers to trust its provenance.

An authorization request to IAM, if successful, returns a JWT which acts as a session token and embeds data about the user, user's role, organisation, and access for direct care, a patient in context. For patient-centric reasons for access, the FHIR Aggregator will only permit data to be requested for the patient in context.

IAM offers a REST service for changing the patient in context in a JWT.

A JWT expires. A data consumer should validate the expiry date before interacting with a data provider and if necessary, refresh the JWT. IAM offers a REST service for refreshing a JWT.

Most interactions with regional services require a valid JWT. This ensures that security and consent constraints are consistently applied.

## 5.3 Interactions with the Data Availability Service

The regional PIX server implements a Data Availability[47] service (a REST interface) which can be used to determine whether records are available through the YHCR and if so from which data providers. The Data Availability services used the NHS number to identify a patient. It is preferred that all interactions with data providers use the NHS number, but local identifiers are also permitted.

The YHCR Participant Registry can also be queried to discover details of data providers' endpoints enabling potentially direct relationships between a data consumer and a provider. Whilst this is theoretically feasible, most consumers use the FHIR Aggregator to mediate relationships.

## 5.4 Interactions with the Regional FHIR Aggregator

The regional FHIR services support the FHIR resource specification of section 6 and FHIR technical standards of section 7 to the highest level of maturity required by registered data consumers.

The functionality of these services is described in section 3. In brief the regional FHIR Aggregator provides:

- Aggregated access to resources across all data providers. Regionally acquired resources have a regional identifier that can be converted into an identifier used by a local data provider (assuming the resource is not a regional resource).
- Management capabilities for resources held in the clinical data repository.

[47] Reference design paper 002 – "Data Availability Service"

Regional resources created by a data consumer, by default, are only accessible and manageable by users with a role within organisation to which the data consumer is registered. Security policies, as described in section 3, can be installed on the regional FHIR bus, to modify default rules.

## 5.5 Interactions with local FHIR Services

A consumer may choose to interact directly with the FHIR service of a data provider. The consumer should anticipate that providers may be at a lower maturity level than the consumer and tailor the user interface accordingly. Specifically, if a provider is not capable of serving a particular resource type then the user interface must differentiate the situation where no data exists from that when data is not obtainable.

A data consumer can ascertain the maturity level of a provider from its FHIR capability statement or by requesting its *Participant* resource from the YHCR Exchange,

## 5.6 Auditing

A data consumer must audit interactions with the YHCR Exchange (and any direct interactions with data providers). Audit records must be keyed by the JWT bearer token issued by the regional OAUth2 service. It is desirable that audit records are held as searchable *AuditEvent* event FHIR resources. A data consumer which is also acting as a data provider can use the regional model FHIR Proxy server for this purpose.

## 5.7 Interaction Diagrams

The following diagram illustrations the interactions arising from a user logging into a local portal and viewing a regional clinical record when the YHCR Exchange is used to aggregate resources. In the interests of clarity interactions between regional services and the regional Audit service are not shown.

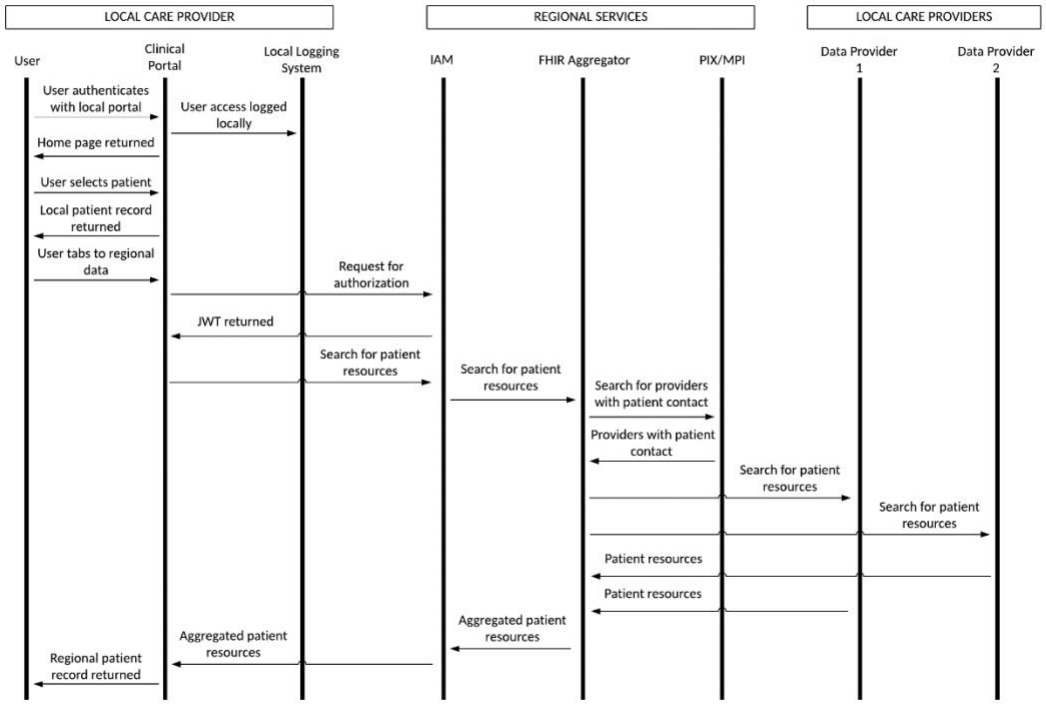


## 5.8 Subscriptions

Data consumers may register subscriptions using supported search paths either with the regional FHIR Aggregator or directly with data providers which support this capability.

Subscriptions monitor data points described using search terms. Support for search terms is aligned with maturity levels. Regionally applied subscriptions devolve responsibility to data providers for monitoring subscribed data. A subscription relying on search term supported at a high maturity level will not receive responses from data providers operating at a lower maturity level.

Both rest hooks and messaging channels are supported. An organisation issuing a subscription must implement an endpoint supporting one or the other[48].

[48] Reference design paper 007 – "Subscription Infrastructure"

# 6 FHIR Resource Profiles

Regional resource standards are based on FHIR STU3 (Standard for Trial Use) and are in accordance with FHIR resource standards used by national services including GP Connect and the Care Connect profiles.

A FHIR resource corresponds to a well-defined healthcare concept. The standard specifies resources in terms of their data content and, with varying degrees of stimulation, how resource properties might be used to model clinical events. Resources represent data which is both in-motion and at rest. They serve as specifications for both interfaces between systems as well as a data model which might be persisted in a database.

The YHCR aims for a high degree of normalisation whereby healthcare concepts represented consistently by healthcare providers with a high degree of coverage and no duplication and using FHIR resources.

## 6.1 Resource Catalogue

The STU3 resource catalogue (excluding resources defined for financial administration) is expressed by the HL7 organisation as follows:

| | Conformance | Terminology | Security | Documents | Other |
|---|---|---|---|---|---|
| **Foundation** | CapabilityStatement | CodeSystem | Provenance | Composition | Basic |
| | StructureDefinition | ValueSet | AuditEvent | DocumentManifest | Binary |
| | ImplementationGuide | ConceptMap | Consent | DocumentReference | Bundle |
| | SearchParameter | ExpansionProfile | | | Linkage |
| | MessageDefinition | NamingSystem | | | Media |
| | OperationDefinition | | | | MessageHeader |
| | CompartmentDefinition | | | | OperationOutcome |
| | StructureMap | | | | Parameters |
| | GraphDefinition | | | | Subscription |
| | DataElement | | | | |

| | Individuals | Entities | Workflow | Management | |
|---|---|---|---|---|---|
| **Base** | Patient | Organization | Task | Encounter | |
| | Practitioner | HealthcareService | Appointment | EpisodeOfCare | |
| | PractitionerRole | Endpoint | AppointmentResponse | Flag | |
| | RelatedPerson | Location | Schedule | List | |
| | Person | Substance | Slot | | |
| | Group | Device | ProcessRequest | | |
| | | DeviceComponent | ProcessResponse | | |
| | | DeviceMetric | | | |

| | Summary | Diagnostics | Medications | Care Provision | Request & Response |
|---|---|---|---|---|---|
| **Clinical** | AllergyIntolerance | Observation | MedicationRequest | CarePlan | Communication |
| | AdverseEvent | DiagnosticReport | Medication-Administration | CareTeam | Communication-Request |
| | Condition | Specimen | MedicationDispense | Goal | DeviceRequest |

| | | | | | |
|---|---|---|---|---|---|
| | Procedure | BodySite | Medication-Statement | ReferralRequest | DeviceUse-Statement |
| | FamilyMember-History | ImagingStudy | Medication | ProcedureRequest | SupplyRequest |
| | ClinicalImpression | ImagingManifest | Immunization | NutritionOrder | SupplyDelivery |
| | DetectedIssue | Questionnaire-Response | Immunization-Recommendation | VisionPrescription | |
| | | Sequence | | RiskAssessment | |
| | | | | RequestGroup | |

| Specialized | Public Health & Research | Definitional Artifacts | Clinical Decision Support | Quality Reporting | Testing |
|---|---|---|---|---|---|
| | ResearchStudy | Questionnaire | GuidanceResponse | Measure | TestScript |
| | ResearchSubject | ActivityDefinition | | MeasureReport | TestReport |
| | | ServiceDefinition | | | |
| | | PlanDefinition | | | |

The catalogue includes resources which are purely technical – these can facilitate the implementation of a FHIR endpoint or might be used to describe and package data or might be necessary for the operation of a particular FHIR capability.

The remaining resources are clinical or administrative and are candidates for implementation by a data provider.

The classification of resources used by the cookbook are:

| | | |
|---|---|---|
| Technical | | A technical resource that is used to describe the FHIR resource framework. These act as a data dictionary and allow FHIR endpoints to be, in part, self-generating. |
| | | A technical resource which is used by FHIR APIs to package data in transit. |
| | | A technical resource is managed by a FHIR endpoint. |
| Clinical | | A clinical resource relevant to operation of a business process, internally or across care settings |
| | | A clinical or administrative resource representing an event or fact relevant for regional consumption |
| | | A clinical or administrative resource subsidiary to a primary concept |
| | | Deprecated |

The following table suggests a classification of resources.

| Resource | Description |
|---|---|
| **ActivityDefinition** | An abstract definition of an activity to be performed maybe as part of a workflow. |
| **AdverseEvent** | An unintended consequence of a medical action. |
| **AllergyIntolerance** | The definition a patient's allergy. |
| **Appointment** | A booking of a healthcare event involving a patient, practitioner, location or device. |
| **AppointmentResponse** | Confirmation or rejection of an attempt to book an appointment |
| **AuditEvent** | A record of a security audit event. |
| **Basic** | A structural resource used to encapsulate unencoded narrative. |
| **Binary** | A structural resource used to encapsulate an image or other binary document. |
| **BodySite** | Used by ProcedureRequests and Observations to define an anatomical location for a particular patient. |

| | |
|---|---|
| **Bundle** | A structural resource used to group a number of related resources, say, in response to a search request. |
| **CapabilityStatement** | Details the capabilities of a FHIR endpoint. |
| **CarePlan** | Describes an intention of how care will be delivered to address a particular condition for a patient or group of patients. |
| **CareTeam** | An assembly or practitioners as a team. |
| **ClinicalImpression** | An assessment aimed at determining the problems affecting a patient. |
| **CodeSystem** | A framework resource defining a system from which a set of codes are drawn. |
| **Communication** | Some form of communication sent from one party to another. |
| **CommunicationRequest** | A request to receive a communication (less formal than a subscription). |
| **CompartmentDefinition** | A framework resource which defines sets of resources which are related in some way to a subject. |
| **Composition** | A structural resource used to embed the content of an immutable document. |
| **ConceptMap** | Define relationships between codable concepts in different vocabularies. |
| **Condition** | A problem, diagnosis or other issue pertaining to a patient or group of patents. |
| **Consent** | A statement of a patient's acquiescence to a consent policy. |
| **DataElement** | A framework resource to descript an item of data. |
| **DetectedIssue** | Indicates an actual or potential clinical issue with a clinical action. |
| **Device** | A medical device or other piece of equipment. |
| **DeviceComponent** | A part of a medical device or other piece of equipment. |
| **DeviceMetric** | A setting or calibration of a device. |
| **DeviceRequest** | A request for a device to be used by a patient. |
| **DeviceUseStatement** | A summary of the usage to which a device has been put to by a patient. |
| **DiagnosticReport** | The findings and interpretations of diagnostic tests applied to a subject (usually but not always a patient). |
| **DocumentManifest** | A structural resource allowing documents to be grouped. |
| **DocumentReference** | A reference to a document. |
| **Endpoint** | The technical detail of an endpoint that can be used in electronic communications. |
| **Encounter** | An encounter with a patient or group of patients. |
| **EpisodeOfCare** | A period of care during which an organisation has a responsibility to a patient. |
| **ExpansionProfile** | A technical resource relating to the use of coding systems. |
| **FamilyMemberHistory** | Health events pertaining to a person related to a patient. |
| **Flag** | Things to be aware of for a patient, medication, location etc. |
| **Goal** | An objective in a care plan. |
| **Group** | A group of patients. |
| **GuidanceResponse** | The result of issuing for guidance to a clinical decision support system. |
| **GraphDefinition** | A structural resource which allows linkages between resources to be defines. |
| **HealthcareService** | A service available at a location. |
| **ImagingStudy** | A set of series of imaging service-object pairs. |
| **ImagingManifest** | A narrative describing images available from an image store. |
| **Immunization** | The record of a vaccination being given to a patient. |
| **ImmunizationRecommendation** | The recommendation for a vaccination with supporting evidence of past immunisations. |
| **ImplementationGuide** | A framework resource representing the documentation of a use-case for FHIR. |
| **Linkage** | A structural resource which links two resources together as representing the same thing. |
| **List** | A structural resource representing a list of other resources. |
| **Location** | A physical location. |
| **Measure** | The definition of a measurement used for quality reporting. |
| **MeasureReport** | A quality measurement. |
| **Media** | A structural resource encapsulating a photo, audio or video recording. |
| **Medication** | The definition of a medication including details of packaging and batch identification. |
| **MedicationAdministration** | The act of a patient consuming a medication. |
| **MedicationDispense** | The act of dispensing a medication. |
| **MedicationRequest** | The prescription of a medication. |
| **MedicationStatement** | A report by a patient or a care professional of a past medication administration. |
| **MessageDefinition** | A framework resource defining a message exchanged between systems. |
| **MessageHeader** | A structural resource defining metadata included in a message. |

| | |
|---|---|
| **NamingSystem** | A framework resource describing a namespace within which concepts are identified. |
| **NutritionOrder** | A request to supply a nutritional composition to a patient. |
| **Observation** | A test result or assessment. |
| **OperationDefinition** | A framework resource descripting a technical operation which can be performed on a FHIR Resource. |
| **OperationOutcome** | A structural resource which describes the outcomes of an attempt to operate on FHIR resources. |
| **Organization** | An organisation. |
| **Parameters** | A framework resource describing parameters in a search string or subscription. |
| **Patient** | A patient. |
| **Person** | A person. |
| **PlanDefinition** | A pre-defined group of actions to be undertaken in a given circumstance. |
| **Practitioner** | A practitioner. |
| **PractitionerRole** | The role a practitioner undertakes in an organisation. |
| **Procedure** | A medical procedure performed on a patient. |
| **ProcedureRequest** | A record of a request for diagnostic investigations, treatments, or operations to be performed. |
| **ProcessRequest** | Deprecated. |
| **ProcessResponse** | Deprecated. |
| **Provenance** | A record of how a resource came be in its current state. |
| **Questionnaire** | A set of questions. |
| **QuestionnaireResponse** | Responses to questions by an individual. |
| **ReferralRequest** | A request to refer a patient to a healthcare service. |
| **RelatedPerson** | A link to a person who is related to another. |
| **RequestGroup** | A sequence of actions to perform for a patient or group of patients. |
| **ResearchStudy** | A basic definition of a research study. |
| **ResearchSubject** | A participation in a research study. |
| **RiskAssessment** | An assessment of the likely outcome(s) for a patient or other subject as well as the likelihood of each outcome. |
| **Schedule** | Part of the mechanism for booking appointments for a clinic/practitioner. |
| **SearchParameter** | A framework resource that describes the search options for a resource type. |
| **Sequence** | A genetic sequencing test result. |
| **ServiceDefinition** | A structural resource which defines the data required by a clinical decision support service. |
| **Slot** | A time period against which an appointment can be booked. |
| **Specimen** | A specimen for testing. |
| **StructureDefinition** | A framework resource which describes a set of properties that defines a concept . |
| **StructureMap** | A structural resource which defines a transformation between two FHIR structures. |
| **Subscription** | A structural resource representing an expression of interest in a data point. |
| **Substance** | A homogeneous material with a definite composition. |
| **SupplyDelivery** | Fulfilment of a request to supply a medication, substance or device. |
| **SupplyRequest** | A request to supply a medication, substance or device. |
| **Task** | Tracks the request and execution of a task issued to an organisation or individual. |
| **TestReport** | Results of the execution of a test script. |
| **TestScript** | A sequence of tests to execute against a FHIR endpoint. |
| **ValueSet** | A framework resource that defines a set of values or a vocabulary to be used for a resource property. |
| **VisionPrescription** | A prescription for vision aids. |

## 6.2 Profiling and the YHCR Maturity Model

Profiles are constraints or extensions to base resource definitions that make them applicable to specific concepts. Most FHIR resources are very generic, and the same resource can be applied to model many different clinical concepts. Consider for instance the Observation resource. This is used for: vital signs, laboratory results, device measurements, clinical assessment indices, personal

characteristics. Among its properties is a code which identifies what is being observed and any number of values which can be used as appropriate. A profile for a blood pressure measurement would stipulate a code to be used to identify the observation as such and a stipulation that exactly two numeric values be provided; a systolic pressure and a diastolic pressure.

Care Connect is a source of profiles with which the YHCR aims to be consistent. However, to recognise the diversity of maturity within the region, different data providers can offer data representing the same concept, but which use different profiles. Profiles are aligned with levels in a maturity model which is curated by DADA.

The maturity model for *Encounters* is used as an example of the use of profiling in the region.

An encounter represents a period of contact between a patient and a care organisation, team or professional. The FHIR definition of an encounter is loose and accommodates a number of interpretations. The following three models for a spell of care involving an emergency department attendance, connected ward stays, outpatient appointments and use of community services are compliant with the Care Connect profile. The first two are active in the YHCR and third is an aspiration for a long-term goal.

1. Independent Encounter resources with location movements embedded in the resource


2. A hospitalization Encounter envelopes related acute contact and Episode of Care links related outpatient contact.

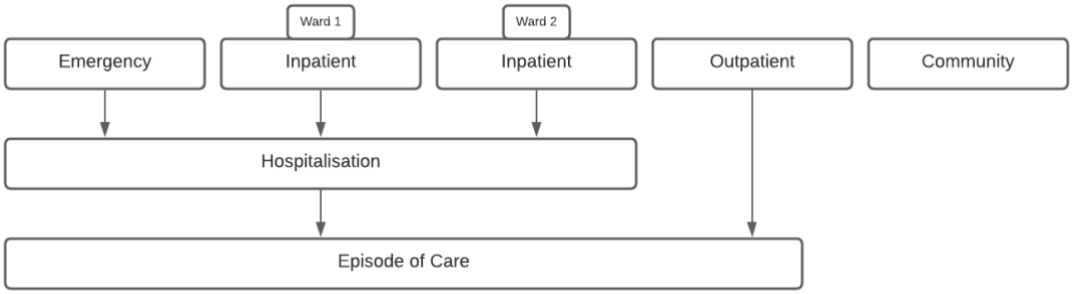


3. The Episode of Care is regionally managed and extends to related community care.

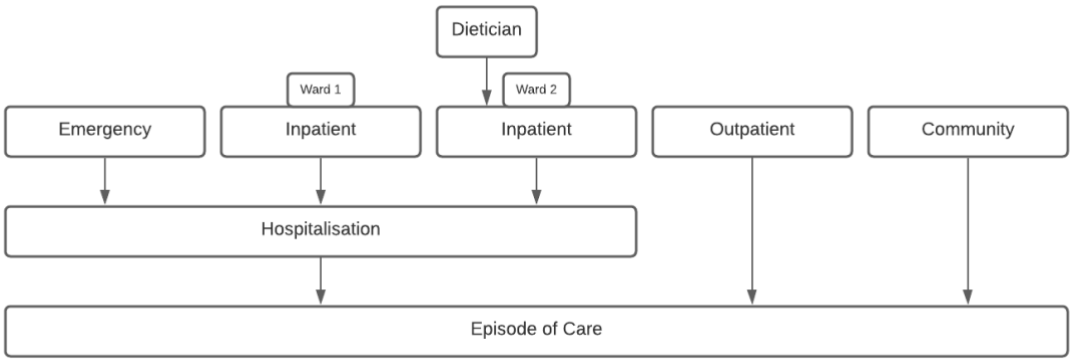


The maturity model assigns these representations to profiles which are aligned to maturity levels. The profiles provide additional constraints over the Care Connect profile.

| Profile | Summary |
|---|---|
| Level 1 | Binding provided to a YHCR standard for Encounter Type Encoding.<br>Uses structure 1. |
| Level 2 | Mandates (where exists) references to Referral Request or Appointment. |
| Level 3 | Managed status code and status history.<br>Uses structure 2. |
| Level 4 | Encounter reason coded to YHCR standard.<br>Mandates set of participants depends on encounter type |
| Level 5 | Mandates reference to causal condition or procedure. |
| Level 6 | Uses structure 3. |

## 6.3 Resource Versioning

FHIR allows resources to be version controlled. FHIR endpoints, unless specifically directed otherwise, will supply the current data when a resource is requested. Resources are stamped with a version number, if the data comprising the resource changes, then the version number is incremented. The former data is retained, and the earlier version of the resource can still be served by the endpoint if explicitly asked to.

Whilst version control is optional for FHIR STU3 it is mandated by the cookbook. Resource versioning enables:

- Audit logs can reference a resource version rather than copying data within the log. The history of data viewed by a user can be reassembled on-demand.
- Evidence can be retained of the data that informed a clinical decision. A user interface which tracks the version numbers of resources shown to a user can recreate a snapshot of the clinical record at a point in time.
- The provenance of clinical data has clinical value: a resource version history makes possible the rendering of a timeline of changes to data.
- The immutability of documents, or FHIR compositions, can be achieved through referencing FHIR versions rather than moving data from its place of origin.

Both the regional infrastructure and model components will support FHIR versioning, and it might be that data providers meet the versioning requirement by using the model FHIR proxy to persist resource versions at the service boundary.

## 6.4 Resource Identification and Disambiguation

All resources available from and endpoint must have an identifier (technically the FHIR id) which is unique within the domain available from the endpoint (i.e. not necessarily globally unique). Resources obtained from the regional FHIR Aggregator will carry a regional identifier. If the resource is a copy of a resource available from another endpoint, then the regional identifier will be constructed for this to be explicit to data consumer.

An endpoint must disambiguate the following resource types:

- Organisation
- Location
- Device
- Person
- Patient
- Practitioner
- Medication

Disambiguation means that any resource obtained from an endpoint that references one of these physical concepts will always use the same resource identifier for the concept regardless of the source of the data.

## 6.5 Vocabularies

The use of clinical vocabularies will not be mandated for early levels of maturity (although the use of Care Connect value sets will be strongly encouraged). Higher levels in the maturity mandate SNOMED-CT or DM+D coding.

## 6.6 Search Parameters

A minimal set of search parameters are mandated for each resource at level 1 maturity which enables data to be assembled by patient and by encounter. The regional FHIR Aggregator and model FHIR proxy component support searching by all search paths defined by STU3.

# 7 FHIR Technical Standards

Technical support will be for FHIR STU3 in accordance with the FHIR Rest specification at: https://www.hl7.org/fhir/STU3/http.html

The FHIR standard defines many capabilities but mandates few. The capabilities which are optional in the STU3 specification but required by the YHCR are detailed here.

"At Regional Level" implies support at the Regional FHIR Aggregator for regionally maintained data.

| Capability | Supported Options | Rational |
|---|---|---|
| **Operations** | Read, create, update | Resource management is required at the regional level at early stages of maturity. At a local level capability is introduced for a basic set of resources at level 3. |
| **Data Format** | XML, JSON | JSON is readily consumable by UIs. XML can be transformed by integration technologies. The formats are technically easy to interchange. |
| **Summary Resources** | Yes | Summary resources are useful for consumption by UIs |
| **Resource Versioning** | Yes | See section 6. |
| **Read History** | Yes | See section 6. |
| **Resource Patching** | At regional level | At local level, data will be manipulated from outside the care setting only in low volume. Capability cost does not justify benefit. |
| **Batch/Transactions** | At regional level | At local level, data will be manipulated from outside the care setting mainly as individual resources. Capability cost does not justify benefit. |
| **CORS** | Yes | Enable direct consumption of resources in a user interface. |
| **Conditional Operations** | At regional level | At local level, data will be manipulated from outside the care setting only in low volume. Capability cost does not justify benefit. |
| **Search Include and Reverse Include** | Yes | Necessary for efficient interaction between a UI and a data source. |
| **Other Search Directives** | _count, _sort | Return a count of the instances that match a search request. |

| | | |
|---|---|---|
| | | Sorts result set. |
| **Search Pagination** | Yes | Necessary for efficient interaction between a UI and a data source. |
| **Search Operators** | Eq,ne,gt,lt,ge,le,sa,eb | The search term ap (approximate searching) is excluded due to implementation complexity. |
| **Search Modifiers** | Missing,exact,text,contains, in,below,above,not-in,type | Comparatively easy to implement with high value added for UIs. |
| **Chained Search Parameters** | Yes (on subject identifier) | Essential for constructing patient centric searches |
| **Compartments** | No | Syntactical embellishment that adds little value. |
| **Messaging / Rest Hooks** | Yes | For subscriptions. |
| **Subscriptions** | Yes | |
| **Asynchronous Searching** | Yes | Required for PHM |
| **Auditing** | Yes | |
| **Linkages** | At regional level | Maintains resource equivalency information which can be used for de-duplication. |

Regional FHIR stores and the model FHIR proxy fully support these API features. Capabilities of data providers who implement their own API service point can be introduced over time in line with the requirements of the technical maturity model specified in design paper 003 – "Conceptual Design for a FHIR Proxy Server".

# 8 Security and Other Non-Functional Requirements

The YHCR is a federated architecture and security is a joint responsibility between the central management team and its participants. The central team has clear responsibility for the central infrastructure and for software that it produces for use locally. It also has a strong interest in the security of its participants: security breaches within its participants have the potential to impact the service offered by the YHCR and to cause reputational damage to the YHCR.

The measures set out here summarise the security features designed into the YHCR which are aimed both at maintain the integrity of centrally run assets and providing confidence in the security of participating organisations.

## 8.1 Securing a Federated Network Operating in the Open Internet[49]

The YHCR uses a layered approach to securing connections into the regional exchange from data consumers and from the regional exchange to data providers as follows:

1. Connections are only possible between whitelisted IP address. Firewalls protecting regional services are managed from the onboarding suite with access limited to authorised personal and all changes are logged and logs independently reviewed.
2. Connections are secured over HTTPS using TLS 1.2 or higher.
3. HTTPS connections are encrypted using certificates signed by the YHCR certificating authority. Endpoints mutually authenticate connecting parties through their certificates and connections are only accepted when certificated by the YHCR.
4. A valid bearer token (JWT) is required for connections into the regional exchange and connections into data providers. Bearer tokens can only be used for 15 minutes and must be signed by a certificate issued by the YHCR certificating authority.
5. Bearer tokens will only be issued to data consumers who have pre-registered with the YHCR presenting from a pre-registered IP address.
6. Consumers must authenticate users using a strong mechanism, establish a legitimate relationship with a patient and a valid reason for accessing the YHCR.
7. For most reasons, access must be in the context of a patient, and this is enforced by the regional FHIR Aggregator.

Whilst a number of these measures can be enforced centrally there is reliance on compliance by participants.

## 8.2 Compliance with Security Standards by YHCR Participants

Participants joining the YHCR are security assured as part of the onboarding process. The assurance checklists cover:

- boundary protection;
- penetration testing;
- accreditation against recognised security standards;
- certificate handling procedures;
- network topologies;
- authentication and password management standards;
- user onboarding and offboarding processes.

[49] Reference design paper 016 – "Securing the YHCR"

Technical compliance is validated by the onboarding suite.

Following promotion to a live environment continuous monitoring tooling validates ongoing compliance with technical standards including:

- rejection of invalid/expired bearer tokens;
- adherence to PKI requirements;
- auditing of requests.

Forensic monitoring tooling can be used to analyse end-user behaviour and pinpoint instances of potential abuse.

## 8.3 Boundary Protection

The central management team uses Google Cloud tooling to identify and mitigate potential breaches including Cloud Armour and the Security Control Centre.

Automated penetration tests run against service points and custom monitoring tools capture usage metrics and alert operators to unusual usage patterns.

The configuration of security sensitive assets (domain names, firewall rules, the participant registry, certificate signing) are under software management and all changes are logged with logs independently reviewed.

All access by the operations team is logged and operator duties segregated to minimise potential for intentional disruption.

## 8.4 Non-Functional Requirements[50]

### 8.4.1 Performance and Scalability

The YHCR offers the following performance guarantees:

Synchronous Query: response within 2.4s.

Asynchronous Query: results available for collection within 24hrs.

Subscription Placement: subscriptions placed within 3hrs.

Subscription Notification: notification delivery attempt made with 30s.

Reliable Messaging: message delivery attempt made within 1min.

### 8.4.2 High Availability and Business Continuity

Regional services are operated by an entity with formal service level agreements, 24x7 staffed service desk, formal incident tracking procedures, and defined incident response times.

Access to a regionally provisioned service desk is provided to all participating organisations.

Regional services are hosted from a high availability google cloud environment with redundancy over three data centres, synchronous data replication and service mirroring which ensures that no one point of failure results in a service outage.

[50] Reference design paper 028 – "Non-Functional Requirements for Regional Infrastructure"

Failure of local endpoints is accommodated by regional infrastructure and local data consumers. Unavailability of any one local FHIR service must not cause failure of searches to return. Potentially incomplete search results must be acknowledged in the returned data structures and user interfaces constructed to inform users of possible gaps in the longitudinal record.

### 8.4.3 Backup and Recovery

All regional data is backed up and restorable within times defined by a service level agreement. All data is maintained within the UK.

### 8.4.4 Monitoring and Alerting

The YHCR implements a monitoring suite which enables real-time analysis of performance, service availability, and usage. Monitoring is end-to-end and encompasses requests made by data consumer through to request fulfilment by data providers.

The solution uses tooling provided the by the Google Cloud Platform supplemented by custom metric reporting and custom dashboards.

The scope of monitoring includes:

- performance;
- availability;
- data quality;
- security (attacks and breaches);
- profile of usage;
- integrity of usage;
- infrastructure utilisation;
- cost.

Threshold can be set against all metrics captured which if breached cause operators to be alerted.

### 8.4.5 Release Management

YHCR software is containerised and deployed as docker images. Building of images from source code is automated and the process uses build scripts which are maintained alongside source code in a third party source code management system.

Images are a registered with the Google Clod Image Repository and  can be released to operational containers (Kubernetes Pods) without downtime.

The YHCR operates a Change Approval Board which authorises all releases and requires evidence of testing an impact assessment.

# 9 Glossary of Technical Terminology

**OAuth2** A standard which separates authentication from authorisation and which is used by the YHCR to establish a domain of trust between data providers, data consumers and regional infrastructure.

**FHIR** Fast Healthcare Interoperability Resources: an international data model and technical standard for representing healthcare concepts and communicating them between systems.

**HL7** Health Level Seven: a standards body responsible for defining the data content of messages or concepts with the intention of ensuring that they are uniformly interpretable across systems.

**HL7v2** A set of messaging standards curated by the HL7 organisation.

**HTTP** Hypertext Transfer Protocol: a language which allows client and server software to exchange data. Used by the World-Wide Web.

**HTTPS** An encrypted version of the above.

**IAM** Identity and Access Management: a computer service which centralises responsibility for authentication, access, and identification of people and organisations.

**JSON** Javascript Object Notation: a standard for structuring data and representing it in a manner that web-based applications can easily consume.

**JWT** Javascript Web Token: a sequence of digits (possibly signed) which act as a key and provide assurance to a system that the holder of the key has been authenticated by a trusted source.

**NEMS** National Event Management Service: an NHS Digital operated service for informing systems about events in the healthcare system that they have registered an interest in.

**ODS** Organisation Data Service: a set of services and a data set provided by NHS Digital, identifying organisations and practitioners known to the NHS.

**OpenEHR** A standard for developing and exchanging data models in healthcare which allows applications to be portable across databases.

**PKI** Public Key Infrastructure, a cryptographic approach for securing and signing data.

**PIX** Patient Identity Cross Reference Manager: a standard curated by Integrating the Healthcare Enterprise which facilitates the cross-referencing of patient identifiers between different organisations.

**PRSB** Professional Records Standards Body: publishers of standards for structuring data in health and social care records.

**REST** Representational State Transfer: a simple standard for defining services provided by one computer system to another.

**SAML** Security Assertion Mark-up Language: a mechanism for an application to pass details of a user's identity and authorization to another trusted application.

| | |
|---|---|
| **SOAP** | Simple Object Access Protocol: a standard for wrapping data in metadata which allows data to be exchanged between systems. |
| **STU** | Standard for Trial Use: a version control numbering system used by HL7 for the FHIR standard. |
| **XDS** | Cross-Domain Document Serving: a standard curated by Integrating the Healthcare Enterprise which allows documents to be shared in an affinity domain. |
| **XML** | eXtensible Mark-up Language: a standard for structuring data and representing it in a format that can be interpreted by different systems in comparatively human readable format. |