



INTERWEAVE
CONNECTING CARE

Cookbook for Regional Interoperability
Detailed Design Paper #027

GP Connect Adapter

PRELIMINARY DRAFT

Version 1.0 – 14th July 2021

Abstract Interoperability Cookbook Anchor Points

Section	Title

Table of Contents

1	Introduction	5
1.1	Purpose of this Document	5
1.2	The Function of an Adapter	5
1.3	GP Connect Interface Variants and Version Numbers	6
1.4	Relationship of this Document with Other Standards.....	6
1.5	Intended Users of the This Document.....	6
2	GP Connect Adapter Architecture	7
2.1	Features of the GP Connect Adapter.....	7
2.2	The GP Connect Adapter as a "Complex" Data Provider.....	7
2.3	GP Connect Adapter SoS Interactions	8
2.3.1	Dependency on the ODS code of a patient's GP practice	9
2.3.2	Caching of SDS LDAP Responses.....	10
2.3.3	Caching of Query Results.....	10
2.3.4	Data Consumer ODS codes and ASIDs.....	10
2.4	Internal Structure of the GP Connect Adapter	10
2.5	Known Deficiencies of the GP Connect Adapter	12
3	Common Features	13
3.1	Validation	13
3.1.1	Interaction Type.....	13
3.1.2	Validates that Permissible Access of GP Connect	13
3.1.3	Resource Types and Interface Variants	13
3.1.4	Validate Search Parameters	15
3.2	Service Linkage Search.....	15
3.3	Establish Cached Data Content.....	15
3.3.1	Cache Compartmentalisation	16
3.3.2	Serialisation of Cache Interactions	16
3.4	Invoke the GP Connect API	16
3.4.1	Determine the Patient's GP Practice ODS Code	16
3.4.2	Determine End-User's ODS Code and System's ASID	17
3.4.3	Invoke SDS	17
3.4.4	Determine Primary Resource Types and Interface Variants	18
3.4.5	Determine Structured/Extended Structure Parameterisation	18

3.4.6	Invoke GP Connect APIs.....	18
3.5	Prepare and Return Results	18
3.5.1	Convert Data Conversion and Manage Cache	18
3.5.2	Prepare Results	19
4	HTML Views	20
4.1	Construction of a DocumentReference	20
4.2	Document Type Mappings.....	20
4.3	Serving of Cached Documents.....	21
5	Classification of Observations.....	22
5.1	Overview	22
5.2	Classification Mechanics.....	22
	Appendix 1 – Maturity Matrix	24

1 Introduction

1.1 Purpose of this Document

This document is one of a series of design papers which underpin the Abstract of a Cookbook for Regional Interoperability (the Abstract Cookbook). These papers, in their totality, describe the technical components and the standards which form the YHCR System of Systems (SoS). They are intended as a basis for developing or procuring software and so are expressed at a level of precision which aims to avoid ambiguity but consequentially, they are focussed to technical readers.

Design papers are anchored to topics which are discussed in the Abstract Cookbook. They are elaborations of the concepts which were first introduced by the abstract and new content is further detail rather than variations of previously established core principles.

This document (design paper 027 - "GP Connect Adapter") specifies requirements for a component of the SoS which enables data to be accessed from GP Systems over the YHCR.

GP Connect is a national interface into General Practice systems and is provisioned by NHS Digital. GP Connect is offered in four variants:

- HTML views (read only HTML views of the GP record).
- Appointments (read only see note below)
- Structured (read only medications, allergies, and appointments).
- Extended Structured (read only access to all coded data).

GP Connect also offers foundation services which provide query access to reference data.

This document provides a design for an adapter into each of these interface variants.

The design covers a read only interface into GP Connect Appointment Management. Creating and modifying appointments will be covered by a separate design should this capability be taken on by the YHCR. The appointment interface describe here will only be used to read appointment data.

1.2 The Function of an Adapter

An adapter mediates between the standards employed by the YHCR and the specific technologies used to interact with the data source, in the case GP Connect.

The SoS standards requires that a data source be able to respond to searches and retrieval requests for atomic FHIR resources on a FHIR STU3 API compliant endpoint. It places requirements on security and audit policies applied by the data source. Data should be packaged in FHIR STU3 format and comply with a profile published for a maturity level for the data items supplied.

The GP Connect interface does not natively comply with any of these requirements:

- most data is accessed using a FHIR operation (which is equivalent to a stored procedure call) rather than FHIR searches and results are supplied as a collection of related FHIR resources rather than search results;
- the endpoint address is not static and must be discovered for each interaction;
- data which is supplied as atomic FHIR resources is served from an API which implements only selected features of the FHIR STU3 API standard;

- GP Connect uses OAuth2 for authorization but requires a claim which differs from that used in the YHCR;
- the YHCR does not have access to audit records written by GP Connect;
- data for the HTML views interface variant is supplied in DSTU2 rather than STU3 format and the interface is not cognisant of the FHIR profiles used in the YHCR.

The adapter is designed to mediate between these differences.

1.3 GP Connect Interface Variants and Version Numbers

At the time of writing interface variants are being released in alignment with the following GP Connect version numbers.

	HTML Views	Appointment Mgmt	Structured	Extended Structured
GP Connect Version	0.7.3	1.2.7	1.2.7	1.5
FHIR Version	DSTU2	STU3	STU3	STU3
Foundation Support	/metadata	/metadata /Patient /Organisation /Practitioner /Location	/metadata /Patient /Organisation /Practitioner /Location	/metadata /Patient /Organisation /Practitioner /Location

In this table, "Foundation Support" details the services which are offered as foundation APIs for individual resource retrieval.

NHS Digital accredit connections against a GP Connect version. The YHCR GP Adapter will initially be accredited for version 0.7.3 providing access to HTML Views. Subsequent accreditations for versions 1.2.7 and 1.5 will provide access to structured FHIR resources as follows:

Version 1.2.7	Version 1.5
Appointment AllergyIntolerance MedicationRequest MedicationStatement Medication	Condition DiagnosticReport Encounter Immunization Observation ProcedureRequest ReferralRequest Specimen

1.4 Relationship of this Document with Other Standards

This document relies on the following standards:

- [HL7 FHIR STU3.](#)
- [GP Connect.](#)

1.5 Intended Users of the This Document

Developers of the GP Connect Adapter and data consumers.

2 GP Connect Adapter Architecture

2.1 Features of the GP Connect Adapter

The GP Connect Adapter presents itself as a data provider to the SoS. The adapter responds to searches and request for resources by interacting with national SDS and GP Connect services. Data retrieved from these services is transitorily cached by the adapter to facilitate processing of FHIR searches and efficient retrieval of resources. Note that while the cache improves the efficiency of searches within a user session, it is used primarily to ensure conformance with the FHIR API standards.

The API exposed by the adapter aims to comply as comprehensively as possible to the STU3 standard. However, scope is limited by the capabilities of the national GP Connect Service. When an API call is received which cannot be serviced then the adapter will return an `OperationOutcome` which will be relayed back by the SoS to the data consumer. Specific limitations in API feature support include:

- only patient centric searches are supported;
- direct resource retrieval is only support for resources previously cached from a previous search request;
- data cannot be created or modified over the API;
- only the synchronous query interaction pattern is supported and not the subscription or asynchronous query pattern.

It should be noted that STU3 search embellishments such as pagination, includes and reverse includes are supported by the adapter. Support for the latter is restricted to searches for a patient which reverse includes FHIR resources available over the GP Connect interface variants.

Direct resource retrieval is supported for cached resources. If a version id is specified, then the resource is always served from cache. If the current version is implied, then the request may result in a cache refresh.

The adapter acts as a document source for "HTML views" of the GP record in compliance with design paper 019 – "Document Management".

2.2 The GP Connect Adapter as a "Complex" Data Provider

The GP Connect Adapter is a Complex Data Provider as defined by design paper 029 – "Data Aggregators and Other Complex Providers". Specifically:

- NHS Digital is the *Operating Organisation* whereas General Practices are the *Governing Organisations*;
- GP Connect only supports explicit routing meaning that either a Source or a Provenance tag must be specified in queries by data consumers.

As a complex provider the GP Connect Adapter implements the following behaviours;

- it is registered in the Participant Registry as a *Dynamic Linkage Provider* and a *Data Aggregator*;

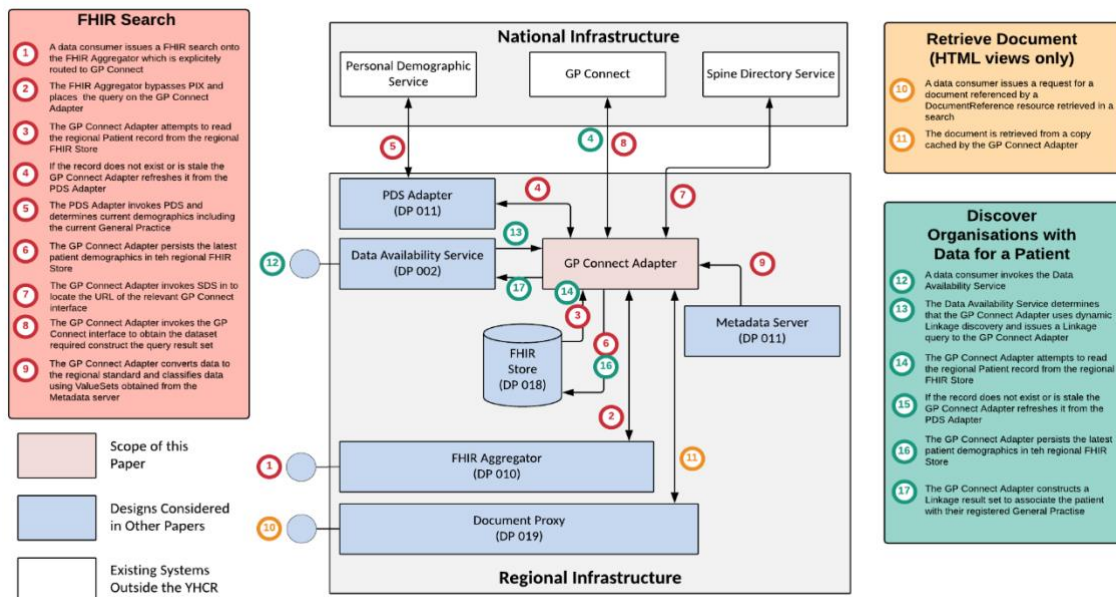
- *Governing Organisation Lists* are automatically maintained (meaning that the FHIR Aggregator is responsible for registering governing organisations in the Participant Registry as data flows from those organisations);
- It implements a *Linkage Query* which is used by the Data Discovery service;
- data returned by the adapter is unified around the regional *Patient* resource with unification performed by the adapter.

2.3 GP Connect Adapter SoS Interactions

The 3 interactions with the SoS which involve the GP Connect Adapter are:

1. **A FHIR search issued by a data consumer.** Patient centric searches which are explicitly routed are placed on the adapter. Linkages are not pre-registered in PIX by the GP Connect Adapter meaning that implicitly routed searches are not placed on the GO Connect Adapter and searches must be explicitly routed by source or governing organisation. The adapter uses the ODS code of the patient's GP practice to locate an endpoint to query. Queries which are targeted at specific resource types (eg: *DocumentReferences*, *MedicationStatements*, *AllergyIntolerances*, *Observations* or *Appointments*) are serviced by interacting with the relevant variant of the GP Connect interface. Searches for *Patient* resources which reverse include other resource types may lead to multiple interface variants being invoked. The adapter is presumptive in that it will gather more data than is requested in the query and cache it for subsequent retrieval.
2. **A request is issued a document (an HTML view) by a data consumer.** This is managed by the SoS document proxy but will offset to the GP Connect Adapter. The adapter will serve the document from its cache.
3. **A data consumer involves the Data Discovery service to ascertain the governing organisations with data for a patient.** The data discovery service involves the adapter's *Linkage Query* which uses the regional *Patient* record (potentially updating the record from the PDS Adapter) to determine the person's registered General Practise and constructing a Linkage result.

The SoS components and their role in servicing these interactions are described below:



2.3.1 Dependency on the ODS code of a patient's GP practice

The relationships with the GP Connect Adapter and other components of the System of Systems are largely determined by the interface's dependency on the ODS code of a patient's GP Practice. GP Connect consumer systems are expected to resolve the FHIR service root URL and Accredited System Identifier (ASID) for a given GP provider organisation using Spine Directory Service (SDS) Local Directory Access Protocol (LDAP) directory lookups.

These details are obtained in 2 separate calls to SDS both of which require the ODS code of the patient's GP Practice to be provided as an input parameter.

The data consumer may provide the ODS code in, or enable it to be derived from, search parameters which are included in its query. This is an optimisation that will lead to the best performance of the GP Connect Adapter. However, the behaviour of the consumer cannot be relied on and the GP Connect Adapter must be able to operate autonomously.

The patient's GP Practice is persisted in the regional FHIR Store and maintained by the regional PIX/MPI service. If the consumer's query does not allow the ODS code to be derived, then the GP Connect Adapter will attempt to use the regionally persisted data. If this data is stale, then the GP Connect Adapter will force a refresh from PDS.

NHS Digital requires that the patient record has been traced and GP practice verified within 24 hours of a call being made to GP Connect. The GP Connect Adapter will force a refresh of from PDS of Patient resources which are more than 24 hours out of date.

Th 24 hour rule also applies to data consumers and is reflected in design paper 015 – "Governance for Data Providers". A data consumer which places a query on the System of Systems which may be directed to the GP Connect Adapter must have traced with PDS the NHS Number for the subject of the query with the last 24 hours. If the consumer specifies the patient's GP Practise in the query then this must have validated this with PDS within the last 24 hours.

2.3.2 Caching of SDS LDAP Responses

NHS Digital recommends that systems connecting to GP Connect cache the results of SDS calls. However, it places restrictions on the lifetime of the cache:

"Systems **MUST NOT** cache and re-use consuming system endpoint information derived from SDS across multiple patient encounters or practitioner usage session".

In practice this means for the SoS that a cache may be held for the lifetime of a JWT (as issued by the Identity and Access Management system as described in design paper 005).

2.3.3 Caching of Query Results

NHS Digital permit query results to be cached by consumer systems for the lifetime of an end user session. This translates to a caching option for the GP Connect Adapter to hold data for the lifetime of a JWT.

2.3.4 Data Consumer ODS codes and ASIDs

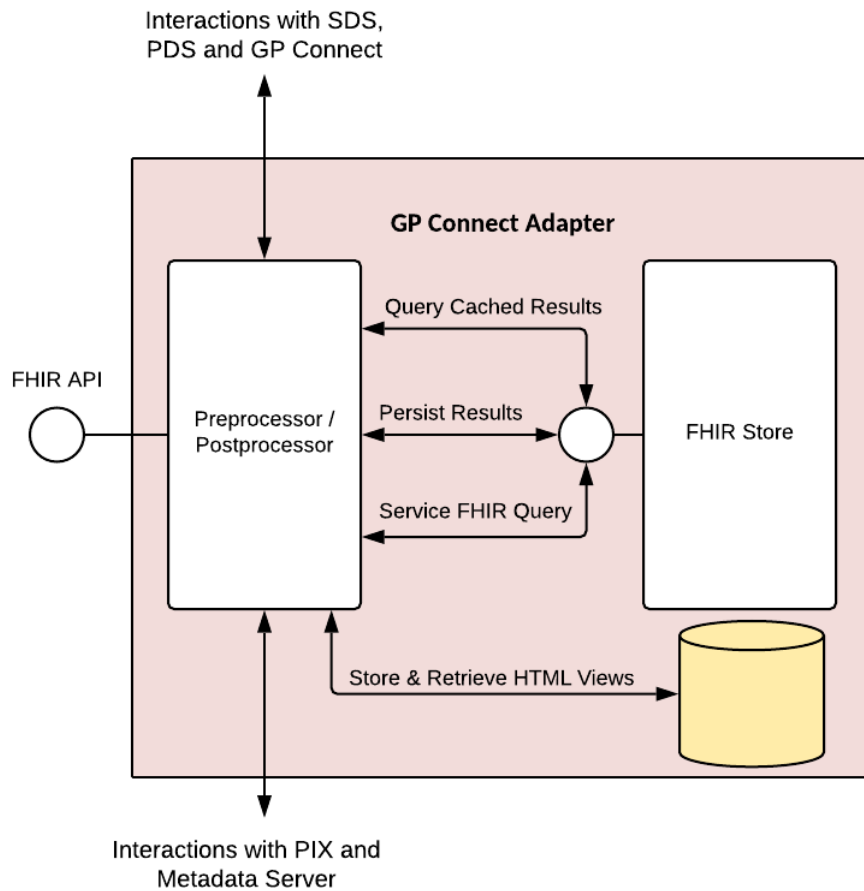
The GP Connect specification requires that the end-user, the organisation for which they work for and the system which is being used to render data are identified in API calls.

The GP Connect Adapter obtains end-user identifiers and their employer ODS code from the consumer's claim made to IAM when obtaining a grant to access the YHCR (reference design paper 005 – "Identity and Access Management").

The ASID is obtained from the IAM claim, if present, and otherwise from the Data Consumer entry in the Participant Registry. Sourcing the ASID in this way allows a data consumer which acts as multiple separate systems in accordance with NHS Digital's definition to legitimately identify itself to GP Connect.

2.4 Internal Structure of the GP Connect Adapter

The GP Connect Adapter presents itself as a data provider which services requests from the SoS, on-demand, through invocation of national services. Internally the adapter uses a FHIR Store to cache results and to achieve compliance with the FHIR API requirements.



The FHIR store persists the atomic FHIR resources which are contained in the bundles which form the response from the national GP Connect adapter. The persisted resources act as both a results cache and a mechanism for assembling results which match complex query parameters.

To elaborate on the last point, the adapter will be responsible for servicing any valid patient centric FHIR search. A search may trigger the download of a segment of the medical record from one or more of GP Connect interface variants. The search may impose conditions or embellishments which filter the content of the search results or cause additional resources to be included.

For instance, a search for:

`Observation?subject=1234&code=abc`

will cause the "Extended Structured" record to be retrieved for patient *1234* which will contain all Observations. These will be filtered so that only those of code *abc* are returned to the data consumer. By populating the FHIR Store with all observations and executing the query against the FHIR Store enables the Adapter to achieve full support for the FHIR search syntax without any development effort. It also leads to Observations of other codes being stored in cache and minimises the need for subsequent calls to GP Connect.

The functionality of the pre-processor and post-processor are detailed below but in summary they:

- service the Linkage query implemented as a complex data provider;
- service requests for HTML documents;
- determine whether cached data in the FHIR Store can be used to service a request;
- determine the GP Connect interface variant to use;
- determine GP Connect interface filters to apply in the API call;
- interact with PIX and the PDS Adapter to obtain the current GP Practise for a patient;
- invoke the GP Connect API;
- convert data from DSTU2 to STU3;
- persist data in the FHIR Store;
- persist documents retrieved from the HTML Document API variant;
- convert Observations to apply code classifications.

The last of the functions is intended to achieve clinical safe interpretation of coded data which may be managed differently by different GP Practise and is described in detail in section 5.

2.5 Known Deficiencies of the GP Connect Adapter

Due to the limited scope of caching permitted by NHS Digital the GP Connect Adapter does not manage FHIR resource versions. A version of a specific resource cannot be retrieved after it has been released by the adapter.

3 Common Features

The functionality of the adapter varies with the API variant being invoked and the resource type being returned. This section details the processing which is common to all API variants and resource types. Extended functionality which applies to HTML Views and *DocumentReferences* is covered in section 4. Extended Structured and *Observations* are covered in section 5.

3.1 Validation

3.1.1 Interaction Type

At this stage only the HTTP GET verb is supported by the adapter and only for synchronous access. An *OperationOutcome* being returned for any other interaction type.

3.1.2 Validates that Permissible Access of GP Connect

At this point GP Connect can only be used for the purpose of direct care. The purpose for which the YHCR is being used can be determined from the data consumer's claim made against IAM in order to gain a bearer token.

Design paper 005 – "Identity and Access Management" details the permutations of the claim. The following are required for access to GP Connect.

Claim Attribute	Acceptable Values
Reason for Access (rsn)	1.1 (Direct care, Emergency) or 1.2 (Direct care, Non-emergency).
User Role (usr.rol)	1 (Clinical Professional) or 2 (Social Care Professional)

Attempt to interact with the GP Connect Adapter without an acceptable claim results in an *OperationOutcome* being returned.

3.1.3 Resource Types and Interface Variants

For direct resource retrieval the interface variant is determined solely by the subject resource type.

Resource searches can be expressed in two ways:

- i) By searching the subject resource type directly or retrieving a subject resource by identifier.
- ii) By searching for a patient and reverse including other resource types.

In the first case the subject resource type determines the interface variant. In the latter case the resources that are reverse included determine the variant. Note that in the latter case more than one interface variant may be invoked and the results assembled from multiple sets of results.

Subject Resource	Interface Variant	Supported FHIR Interaction			
		Search Subject	Direct Retrieval	Reverse Include	Include
<i>DocumentReference</i>	HTML Views	✓	✓	✓	x
<i>Appointment</i>	Appointments	✓	✓	✓	x
<i>MedicationStatement</i>	Structured	✓	✓	✓	x
<i>MedicationRequest</i>	Structured	✓	✓	✓	x

PRELIMINARY DRAFT

AllergyIntolerance	Structured	✓	✓	✓	x
Immunization	Extended Structured	✓	✓	✓	x
Encounter	Extended Structured	✓	✓	✓	x
Condition	Extended Structured	✓	✓	✓	x
Observation	Extended Structured	✓	✓	✓	x
DiagnosticReport	Extended Structured	✓	✓	✓	x
ProcedureRequest	Extended Structured	✓	✓	✓	x
ReferralRequest	Extended Structured	✓	✓	✓	x
<i>Patient</i>	Any ⁽¹⁾ ⁽³⁾	✓	x	x	✓
<i>Organisation</i>	Foundation/Any ⁽²⁾ ⁽³⁾	x	x	x	✓
<i>Practitioner</i>	Foundation/Any ⁽²⁾ ⁽³⁾	x	x	x	✓
<i>PractitionerRole</i>	Any ⁽²⁾	x	✓	x	✓
<i>Location</i>	Foundation/Any ⁽²⁾	x	✓	x	✓
<i>Medication</i>	Structured ⁽²⁾	x	✓	x	✓
<i>Specimen</i>	Structured ⁽²⁾	x	✓	x	✓

Resource types in bold are referred to below as clinical resources.

Supported FHIR Interactions list the references to the resource type which are permitted in the FHIR interaction with the adapter and comprise:

Search Subject	The resource is the subject of a FHIR search GET Observation?subject=[patient]
Direct Retrieval	The resource is the subject of a retrieval request GET Observation/[id]
Reverse Include	The resource is reverse included on a Patient search GET Patient?identifer=[nhs]&_revinclude=Observation:subject
Include	The resource is included on a search with a different subject GET Observation?subject=[patient]&_include=Observation:patient

Notes

- (1) The patient resource, when, the subject of a search request can be targeted at any interface variant depending on the resource types which are reverse included.
- (2) *Organisation, Practitioner, PractitionerRole, Location, Medication and Specimen* resources cannot be the subject of a search request. They can only be reverse-included on a search for one of the primary resources or in the case of *PractitionerRole, Location* and *Medication* resources be the subject of a direct retrieval by identifier.
- (3) *Patient, Organisation* and *Practitioner* resources are unified by the Adapter meaning that references in search results to these resource types will be returned as references to regional resources.

An empty bundle is returned for searches which do not target a resource type which can be associated with an interface variant.

A *Patient* search which reverse-includes resource types not supported by an interface variant returns a bundle containing the *Patient* and any supported reverse included resource types.

3.1.4 Validate Search Parameters

Searches must be patient-centric. For the above, non-*Patient* resource types, the search must contain either a subject or patient search parameter and this must reference a single patient. For *Patient* searches the identifier search parameter must be included.

An *OperationOutcome* is returned is returned for non-patient-centric searches.

Patient, Organisation and Practitioner resource references in search terms must be to regional resources. The adapter unifies data around these concepts and GP Connect local references will not be disclosed to data consumers.

3.2 Service Linkage Search

Linkage searches are issued by the Data Availability Service (design paper 002). and are of the following format:

```
GET Linkage?source.identifier= http://fhr.nhs.uk/Id/nhs-number|[nhs
number]&_include=author
```

The adapter constructs a *Linkage* resource with a single 'source' item being the regional *Patient* resource that corresponds to the requested NHS number. The author is the regional *Organisation* resource that corresponds to the patient's general practice.

In servicing this request the adapter retrieves the regional *Patient* resource from the PIX FHIR Store. If the resource does not exist or the resource is more than 24 hours old then the adapter invokes the PDS Adapter, obtains a current version of the resource, and persists it in the PIX FHIR Store.

The *Linkage* resource is assigned an id which can be associated with the NHS number of the patient to which it pertains. Subsequent retrieval of the resource using:

```
GET Linkage/id
```

returns the same *Linkage* as would have been obtained by an NHS number search.

3.3 Establish Cached Data Content

Resources retrieved from GP Connect are persisted in a FHIR Store attached to the adapter. The FHIR Store acts as a cache from which results may be served for subsequent invocations of the adapter by the same user that issued searches which caused the cache to be populated.

Results will only be served from the cache only if:

- the results were created by the same user session as the one requesting results – the session being determined by the JTI of the JWT authorising access (refer to design paper 005 – "Identity and Access Management");
- the date range of cached resources encompasses the data any data range specified in the search;
- all resources requested in a search can be served from cache.

The adapter maintains cache control data which for each resource type, Patient and JTI and specifies:

- the data and time when resources were cached;
- the date range used in the original GP Connect API invocation.

If a search (or resource retrieval) can be serviced from cache, then the GP Connect Adapter is not invoked.

Note that cache control data is maintained for the JTI used to invoke the adapter. This because the contents of the cache are segmented by JTI and multiple copies of the same resource may exist in the cache.

3.3.1 Cache Compartmentalisation

The cache is a modified version of a FHIR Store which enables multiple copies of resources to be held in parallel each associated with the JTI which originally create it. Search results only contain those resources held for the JTI invoking the search. This compartmentalisation presents the persistence layer as though it were multiple separate instances of a FHIR Store.

Details of how this is achieved is left the implementor. One possible strategy is to use a standard FHIR Store, prefix resource id's by JTI, and include the JTI as a meta tag which is added to all searches.

3.3.2 Serialisation of Cache Interactions

Interactions with the cache and cache control data are queued and serialised by JTI. This allows multiple requests to be processed in parallel for different JTIs but maintains the integrity of the cache compartments when multiple requests are received for the same JTI.

3.4 Invoke the GP Connect API

3.4.1 Determine the Patient's GP Practice ODS Code

The most performant use of the GP Adapter is for a consumer to reference the subject patient's organisation as a search parameter in its request.

If the search targets a *Patient* resource, then this is achieved by including a general-practitioner search term that targets a single ODS code. An example expression might start:

```
Patient?identifier=https://fhir.nhs.uk/Id/nhs-number|{nhs number}&general-practitioner.identifier=https://fhir.nhs.uk/Id/ods-organization-code|{ods code}
```

The remainder of the expression would "reverse include" resource types to be included in the search results.

A search which targets a clinical resource type would achieve similar by using search parameters chained from the resource's subject or patient attribute. For example:

```
DocumentReference?subject.identifier=https://fhir.nhs.uk/Id/nhs-number|{nhs number}&general-practitioner.identifier=https://fhir.nhs.uk/Id/ods-organization-code|{ods code}
```

There are other syntaxes which would achieve the same result. There are also other options for referencing a general practitioner. These include:

- a) Referencing the NHS Spine organisation resource ie:

```
general-practitioner=https://directory.spineservices.nhs.uk/STU3/Organization/{ods  
code}
```

- b) Referencing a local *Organisation* resource.
c) Referencing the YHCR *Organisation* resource.

The last two options cause the GP Connect Adapter to interact with the regional FHIR Store and obtain the YHCR *Organisation* resource either directly or by querying *Linkages* if the reference is to a local resource. The YHCR *Organisation* resource is the source of the required ODS code.

If the GP Practice *Organisation* is not determinable from the query parameters, then the GP Connect Adapter will determine it from the YHCR *Patient* resource.

The method for determining the YHCR Patient resource depends on the type of request:

- a query for a clinical resource will reference the patient as a query parameter;
- a query for a Patient resource will include the patient identifier as a query parameter;
- a request to retrieve a clinical resource will have established the latest cached version and the patient can be determined from its subject;
- a request to retrieve a *Location* will use the patient identifier embedded in the JWT as presented by the consumer.

There are options for referencing a *Patient* resource and these will lead to different mechanisms for retrieving it from the PIX FHIR Store. The patient may be referenced as:

- a YHCR Patient reference and the resource will be obtained by direct retrieval;
- a local or GP Connect reference and the resource will be obtained by a *Linkage* query;
- an NHS number and the resource will be obtained by an identifier query.

If the Patient resource was updated within the last 24 hours (as determined by its *lastUpdated* meta data property, then the *general-practitioner* that it references will be the source to the required ODS code. If the Patient resource does not exist or is more than 24 hours old then from the PDS Adapter (design paper 011 – Interfaces With the Personal Demographic Service).

3.4.2 Determine End-User's ODS Code and System's ASID

The end-user's employer organisation ODS is extracted from the JWT presented to IAM.

The end-user's system ASID is extracted from JWT if provided and otherwise from the data consumer's registration in the Participant Registry.

If either the end-user's ODS code or ASID is not available, then an *OperationOutcome* is returned.

3.4.3 Invoke SDS

Separate invocations are required for the GP Connect service endpoint and the provider system's ASID.

Details of the LDAP call are [here](#).

In line with the guidance provided by NHS Digital, the adapter will cache, in memory, the result of calls to SDS. The cache will be keyed on the consumer's JWT and ODS code.

NHS Digital's requirement to cache per session basis is met by keying the cache on JWT.

Cache entries expire on expiry of the JWT and memory should be released.

3.4.4 Determine Primary Resource Types and Interface Variants

For most FHIR interactions processed by the adapter there will be a single primary resource type that determines the interface variant this being the subject of the FHIR interaction. The exception is a Patient search which reverse includes clinical resource types. Each clinical resource reverse includes may necessitate a separate interaction with GP Connect. Note that multiple resource types can be retrieved in a single Structured or Extended Structured GP Connect interaction

3.4.5 Determine Structured/Extended Structure Parameterisation

API parameters control which resource types are included in results returned by the structured/extended structured interface variants. More than one parameter is often required to ensure that all instances of the resource type and references to other resources are included.

API parameters are set as follows:

<i>AllergyIntolerance</i>	includeAllergies, includeConsultations
<i>Condition</i>	includeAllergies, includeMedication, includeConsultations, includeProblems, includeImmunisations, includeUncategorisedData, includeInvestigations, includeReferrals, includeDiaryEntries
<i>MedicationStatement</i>	includeMedication, includeConsultations, includeProblems,
<i>MedicationRequest</i>	includeMedication, includeConsultations, includeProblems
<i>Observation</i>	includeUncategorisedData
<i>Encounter</i>	includeConsultations
<i>Immunization</i>	includeConsultations, includeProblems, includeImmunisations
<i>DiagnosticReport</i>	includeInvestigations
<i>ProcedureRequest</i>	includeInvestigations, includeDiaryEntries
<i>ReferralRequest</i>	includeReferrals

When including allergies, resolved allergies should also be included. When including medications, prescription issues should also be included.

3.4.6 Invoke GP Connect APIs

API specifications are available at: [GP Connect Foundation interface](#), [HTML Views](#), [Structured Data](#), [Appointments interface](#) and [Extended Structured Data](#).

3.5 Prepare and Return Results

3.5.1 Convert Data Conversion and Manage Cache

Search bundles are parsed, and atomic resources are extracted or constructed. Resources retrieved from the HTML Views interface are in DSTU2 format and are converted to STU3 format in accordance with the specifications in section 4.

Patient, Organisation and Practitioner resources are replaced with regional resources and references to these resource types in other resources are replaced with references to the regional resources.

Observation resources are classified in accordance with the specification in section 5. Unclassified observations are discarded.

Cached copies of resource types extracted from API results are discarded for the current JTI and extracted resources are persisted in the FHIR Store. Cache control data is updated for the JTI and Patient for the relevant resource types.

3.5.2 Prepare Results

A modified version of the query originally issued by the consumer (or resource request) is executed against the embedded FHIR store. The modifications are as discussed in 3.1.5 and are designed to exclude cached data that may pollute the result set.

The modification are discussed for each interface variant in the section below.

The results of the query along with appropriate *OperationOutcome* describing any data impairments are returned to the data consumer. Data impairments (design paper 017) arise from non-uniform support by the different GP system vendors.

4 HTML Views

4.1 Construction of a DocumentReference

The HTML Views interface variant returns a search set containing:

- a Composition (who's section text contains the HTML View requested);
- an Organisation (representing the GP Practise);
- a Practitioner (representing the Patient's practitioner at the GP);
- a Patient.

In accordance with design paper 019 – "Document Management", the adapter presents document metadata as FHIR DocumentReference resources with attributes derived from the GP Connect composition as follows:

status	current
type	see table below
subject	reference to regional Patient resource
created	current date/time if no upper date provided in search otherwise the provided date
indexed	current date/time
author	reference to regional Practitioner resource representing the general practitioner resource as identified by the SDS code in the GP Connect composition
custodian	reference to the regional Organisation resource as identified by the ODS code in the GP Connect composition
description	GP Connect record section as per the table below
content.attachment[1].contentType	text/html
content.attachment[1].data	base64 encoded copy of
content.attachment[2].contentType	text/html
content.attachment[2].url	Url constructed by the adapter
context.practiceSetting	SNOMED CT code 394814009 "General Practise"

Note that including 2 representations of the attachment places the DocumentReference in compliance with the regional document architecture, which requires documents to be able to be served independently from meta-data from a URL offered in the DocumentReference, and, by also offering the attachment as a base64 embedded document, the design recognises that as an explicitly routed data source, data consumers will be addressing GP connect with full knowledge of the documents that are available and in almost all circumstances a request for DocumentReference indicates a desire to access the document itself.

4.2 Document Type Mappings

SNOMED-CT		GP Connect Record Section	
Code	Description	Code	Description
5491000179105	Medical Record Summary	SUM	Summary

371531000	Encounter report	ENC	Encounters
371524004	Clinical report	CLI	Clinical Items
717711000000103	Problems	PRB	Problems
886921000000105	Allergies & Reactions	ALL	Allergies & Adverse Reactions
163111000000100	Medication Record	MED	Medications
886721000000107	Referral details	REF	Referrals
1102421000000108	Observations	OBS	Observations
1102181000000102	Immunisations	IMM	Immunizations
716611000000107	Clinical Administration	ADM ⁽¹⁾	Administrative

Notes

(1) Not available from EMIS

4.3 Serving of Cached Documents

Documents extracted from the composition are persisted by the adapter on a file system and are served by the adapter in response to the HTTPS URL published in the DocumentReference being invoked.

Cached documents are deleted when the corresponding DocumentReference is removed from the adapter's FHIR Store.

5 Classification of Observations

5.1 Overview

Much of the structured data held in GP Systems is presented over GP Connect as *Observation* resources. This category of data includes any coded data that is not appropriately represented as one of the other clinical resources and is known to the GP Connect APIs as Unclassified Data. Such data includes assessments made by GPs, preferences of the Patient, and measurements such as lung capacity and blood pressure.

The structure of the Observation resource by GP connect differs depending on whether the resource represents a measurement or other types of data. For measurements the Observation code defines what is being measured and the Observation value provides the result of the measurement. So a lung capacity measurement may be represented as:

Observation code: 54715006 " Diffusion capacity of lung, function"

Observation value: 5400 mL

Assessments and preferences are only represented by the Observation value which is coded as a codable concept. A preferred place of death may be represented as:

Observation value: 108401000000102 "Preferred place of death: hospice"

Much of the unclassified data is collected and assigned codes on templates which can be configured for individual GP Practises. The 'bag-of-codes' approach to representing assessments and preferences introduces, in the opinion of the author, a clinical risk: a data consumer must recognise every code used by every practise to determine whether a concept has been coded. Using the preferred place of death as an example, a consumer may recognise 108401000000102 as signifying a preferred place of death has been registered as a hospice but unless they also recognise 931401000000103 as signifying a preferred place of death as a hospice community lodge then they might assume that no preferred place of death had been registered were this code to be presented by a GP Practise. Slight deviations in templates use by GPs make data uninterpretable when aggregated.

Template standardisation is essential and the SoS has a role in classifying data so that data consumers can ensure that they are applying a regionally standard interpretation.

This design proposes that the adapter filters assessment and preference type Observations to ensure that only those that use approved code are released to consumers and that the adapter should format them in a manner that makes the concept being coded explicit.

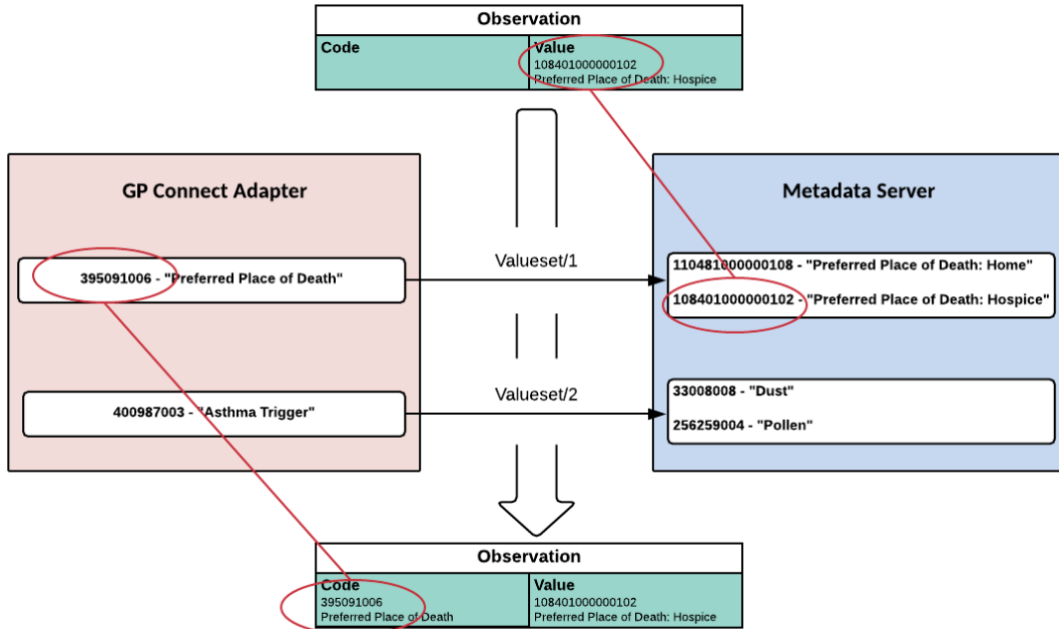
This technical solution will accompany a process for template rationalisation which will improve the consistency of coding and lead to coded concepts being released in alignment with data consumption use cases.

5.2 Classification Mechanics

The GO Connect Adapter will use *ValueSets* stored in the regional Metadata server to collate codes which are approved for release from the adapter. The adapter will associate relevant *ValueSets* with a code that identifies the clinical concept being described.

When an *Observation* resource is received from GP Connect which does not have a Code attribute but uses a *CodeableConcept* as a value then the GP Connect Adapter classifies the *Observation* by searching relevant *ValueSets* for the *CodeableConcept* code. If it is located the GP Connect Adapter assigns the *Observation* code that it associates with the ValueSet.

The process is illustrated below:



Note that uncoded *Observations* that cannot be classified u Uncoded *Observations* with multiple values which cannot be uniquely classified will be withheld by the GP Connect Adapter.

Appendix 1 – Maturity Matrix

Section	Narrative	Consultative	Draft	Normative
1 Introduction	X			
1.1 Purpose of this Document				
1.2 The Function of an Adapter	X			
1.3 GP Connect Interface Variants and Version Numbers				X
1.4 Relationship of this Document with Other Standards	X			
1.5 Intended Users of the This Document	X			
2 GP Connect Adapter Architecture		X		
2.1 Features of the GP Connect Adapter				
2.2 The GP Connect Adapter as a "Complex" Data Provider		X		
2.3 GP Connect Adapter SoS Interactions		X		
2.3.1 Dependency on the ODS code of a Patient's GP Practice				
2.3.2 Caching of SDS LDAP Responses		X		
2.3.3 Caching of Query Results		X		
2.3.4 Data Consumer ODS codes and ASIDs		X		
2.4 Internal Structure of the GP Connect Adapter		X		
2.5 Known Deficiencies of the GP Connect Adapter		X		
3 Common Features		X		
3.1 Validation				
3.1.1 Interaction Type				
3.1.2 Service Linkage Search		X		
3.1.3 Resource Types and Interface Variants		X		
3.1.4 Validate Search Parameters		X		
3.2 Service Linkage Search		X		
3.3 Establish Cached Data Content		X		
3.3.1 Cache Compartmentalisation				
3.3.2 Serialisation of Cache Interactions		X		
3.4 Invoke the GP Connect API		X		

PRELIMINARY DRAFT

3.4.1 Determine the Patient's GP Practice ODS Code				
3.4.2 Determine End-User's ODS Code and System's ASID		X		
3.4.3 Invoke SDS		X		
3.4.4 Determine Primary Resource Types and Interface Variants		X		
3.4.5 Determine Structured/Extended Structure Parameterisation		X		
3.4.6 Invoke GP Connect APIs		X		
3.5 Prepare and Return Results		X		
3.5.1 Convert Data Conversion and Manage Cache				
3.5.2 Prepare Results		X		
4 HTML Views		X		
4.1 Construction of a DocumentReference				
4.2 Document Type Mappings		X		
4.3 Serving of Cached Documents		X		
5 Classification of Observations		X		
5.1 Overview				
5.2 Classification Mechanics		X		