



**INTERWEAVE**  
CONNECTING CARE

Cookbook for Regional Interoperability  
Detailed Design Paper #007

Subscriptions Infrastructure

PRELIMINARY DRAFT

Version 1.0 – 10<sup>th</sup> April 2019

**Abstract Interoperability Cookbook Anchor Points**

Section	Title
4.3	Subscriptions

## Table of Contents

1	Introduction .....	4
1.1	Purpose of this Document .....	4
1.2	FHIR Subscriptions in the YCHR .....	4
1.3	Relationship of this Document with Other Standards.....	5
1.4	Intended Users of the This Document .....	5
2	Outline of the YHCR Subscriptions Architecture .....	6
2.1	Behaviour of Data Consumers .....	6
2.2	Cohort Management .....	7
2.3	Subscriptions and Historic Data.....	7
3	Processing Model for Data Providers .....	9
3.1	Alignment of Data Provider Capabilities with the Maturity Model .....	9
3.2	Local Protections and Controls.....	9
3.2.1	Subscription Constraints.....	10
3.2.2	Publishing Subscription Constraints .....	10
3.3	Processing Model for a FHIR Proxy Server .....	11
3.3.1	Periodic Subscription Processing .....	12
3.3.2	Trigger On-Demand Subscription .....	12
3.3.3	On-Demand Subscription Processing .....	12
3.3.4	Package & Queue Subscription Results .....	12
3.3.5	Dispatch Messages .....	12
3.3.6	Dispatch REST Hooks .....	12
3.4	Validation of Subscription Management Requests .....	13
3.5	Processing Model Computational Costs.....	13
4	Regional Subscription Manager Technical Design.....	15
4.1	Creating Regional Subscriptions .....	15
4.1.1	Criteria which Reference Patients .....	15
4.2	Subscriber Identity Management, Security and Audit .....	16
4.3	Cohort Management .....	16
4.4	Rest Hooks and Messaging Channel Mediation .....	16
4.5	Capability Discovery.....	17
4.6	Revoking Subscriptions .....	17
5	Subscriptions and the National Event Management Service .....	18

---



## 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. They are intended as a basis for developing or procuring software and so are expressed at a level of precision which is intended to avoid ambiguity but with a consequence that 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 007 - "Subscriptions Infrastructure" describes a regionally standardised approach for participants in the YHCR to express an interest in data that matches specified criteria and for them to be notified when data is created that meets these criteria.

### 1.2 FHIR Subscriptions in the YCHR

The HL7 FHIR standard incorporates the concept of a subscription: the ability to register an interest in a data point. In FHIR terms a subscription is a resource that is created at a data provider and which describes:

- i) the data points that a data consumer has an interest in;
- ii) the mechanism by which the data provider may inform the consumer that data has been created that matches their interest.

The possible uses of subscriptions are widespread and varied. Examples of some use cases include:

- a clinician wishes to be alerted to events pertaining to the patients which they currently treat;
- an algorithm wishes to analyse blood test results for a cohort of patients and wishes for these to be delivered to it as they arise;
- an ambulance service wishes to understand the outcome of an A&E attendance in which it has been involved.

FHIR subscriptions could operate on a point-to-point basis with individual data consumers and providers forming bipartite relationships whereby a consumer registers subscriptions with individual providers and those providers informs the subscriber directly of data that matches the consumer's interests. However, this approach has disadvantages:

- a consumer must manage subscriptions with every data provider if it I to achieve full coverage of all possible data points across the YHCR;
- a subscriber must deliver subscription content to every subscriber and overlapping subscriptions may lead to the same content being delivered multiple times;
- often subscriptions relate to a cohort of patients and bipartite relationships do not lend themselves to regional definitions of patient groups;
- firewalls must allow data providers to distribute data to any data consumer.

This document proposes an alternative model whereby subscriptions are created against central regional service and this service assumes responsibility for propagating them out to individual data providers. Under this model data providers remain responsible for servicing subscriptions but the YHCR acts as an agent for distributing subscription results to data consumers.

### **1.3 Relationship of this Document with Other Standards**

The following standards form the basis for this document:

- FHIR Release 3 (STU) – [Messaging Using FHIR Resources](#);
- InterOpen [CareConnect FHIR profiles](#);
- National Event Management System ([NEMS](#));
- [ITK3 Messaging Distribution](#);
- Message Exchange for Health and Social Care([MESH](#))

### **1.4 Intended Users of the This Document**

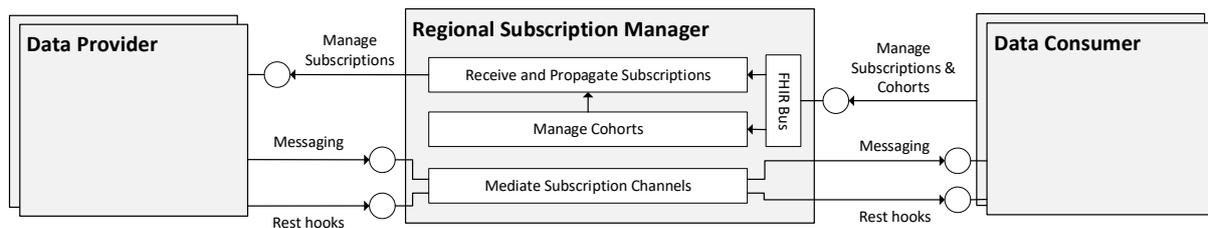
This document is a reference guide for data providers intending to send messaged within the YHCR, data consumers intending to receive messages, and developers of regional infrastructure. The document should also be read by organisations participating in National Event Management Systems pilots or those intending to publish data to this service.

## 2 Outline of the YHCR Subscriptions Architecture

The following architecture establishes a responsibility on regional infrastructure for subscription management. The Regional Subscription Manager is responsible for:

- receiving subscriptions from data consumers;
- distributing subscriptions to data providers;
- receiving subscription results from data providers;
- distributing subscription results to data consumers;
- managing cohort definitions.

The architecture is illustrated below:



Subscriptions are lodged with the YHCR Regional Subscription Manager by creating *Subscription* resources. As with all FHIR resources in the YHCR, *Subscription* resources are created, RESTfully, by issuing an HTTP POST request on the regional aggregator. The Regional Subscription Manager propagates subscriptions out to local FHIR proxy endpoints.

The FHIR standard defines a number of alternative delivery channels for subscription results. The architecture presented here offers two alternatives:

- 1) Rest hooks whereby subscription results are 'POSTed' over HTTP to a RESTful endpoint.
- 2) Messaging whereby subscription results are sent over the reliable messaging infrastructure described by design paper 006 – "Messaging Infrastructure."

A data consumer can determine the most appropriate channel for their own purposes. The Regional Subscription Manager may also operate both channels but, in order to rationalise requirements for data providers when responding to identical subscriptions made by different data consumers, it will standardize channels by search type pattern. I.e.: if multiple data consumers subscribe to the same search path then the subscription will be lodged with each data provider only once using a delivery channel that is appropriate for the search string. Note that this behaviour may result in the subscription manager mediating between rest hook and messaging channels. How the most appropriate channel is determined is an implementation detail which is not specified here.

### 2.1 Behaviour of Data Consumers

This paper doesn't make any assumption as to what constitutes a data consumer or how a data consumer behaves when it receives subscription results. Some possibilities include:

- A consumer is dashboard. Subscription result data is stored in a database and the dashboard uses the data to display summary statistics such as an illustration of healthcare services are being used at this point in time;
- A consumer is an alerting system. Subscription results are emailed to clinicians;

- A consumer is an algorithm. Subscription results are processed by to reveal trends in a data category such as platelet counts recorded in blood tests.

Data consumers may exist in care settings or centrally within the YHCR. Care settings are encouraged to manage their own relationships between clinicians and subscriptions: a subscription should be made by a care setting rather than an individual clinician and if required subscription results should be messaged to clinicians by the care setting using an appropriate internal routing mechanism.

## 2.2 Cohort Management

A very common requirement is to be notified of occurrences which pertain to a specific group of patients or a cohort. One approach is to make separate subscriptions for each patient in the group. This is an appropriate mechanism when the cohort is managed by the same care setting which is acting as a data consumer: as the cohort changes then subscriptions can be created and revoked as required.

However, situations can be envisaged where the cohort of patients is maintained regionally:

- Management of the cohort's membership is shared between different care settings such as the subjects of multi-disciplinary team meetings.
- The cohort is a segment of the regional population which is defined by health condition, treatment, or other observable attribute such as patients being treated for cancer.

In these situations, the cohort definitions are maintained regionally and persisted regionally (design paper 018 – Regional FHIR Store). Data consumers might wish to reference the cohort as a search parameter in their subscriptions. Where this is the case then the Regional Subscription Manager is responsible for translating subscriptions to cohorts into subscriptions for individual patients and lodging these with data providers. As membership of the cohort changes then the Regional Subscription Manager creates new subscriptions or retracts existing ones.

FHIR uses *List* resources to define cohorts and *List* in the regional FHIR Store can be managed externally using standard RESTful APIs. *Lists* themselves can be managed through subscriptions which are made by regional data consumers: an algorithm may subscribe to data points indicating a particular condition and manage a *List* based on subscription results.

FHIR does not provide the syntax to use a *List* in a search term and this design paper proposes an extension to the search specification to allow regional cohort management to behave as just described. As a design principle, the LHCRE should avoid extending standards. However, this extension is considered justified as it has no impact on any actors other than those with a specific interest in regional cohort management. It will only be implemented by the Regional Subscription Manager and a breed of data consumer which is yet to exist.

## 2.3 Subscriptions and Historic Data

FHIR subscriptions only apply to data which is created or updated from the point that a data provider accepts a subscription. Subscription results will only be delivered for new data which matches the subscription rule. For many use cases then this is the appropriate behaviour: a subscription to identify at-risk patients attending A&E may only be relevant for future attendances.

However, there are a significant set of use cases where a subscriber requires historic data. For example, establishing the definition of a cohort of patients being treated with Cancer requires that the cohort is created from historic data and then maintained through subscription to changes in diagnoses.

Consideration has been given to extending the *Subscription* resource definition to allow it to specify that pre-existing matches should also be returned in subscription results. The extension has been rejected in this case as uniform support would be required across all data providers for the extension to have meaningful application. This would limit providers' choices in using the native FHIR capabilities of existing applications or adopting off-the-shelf FHIR proxy software.

An alternative approach is proposed. A data consumer wishing to receive both historic and new/changed data:

1. Issues a subscription with an appropriate search string for new/updated data.
2. Runs a query using the same search string to acquire historic data.

Historic data may be voluminous. A data provider may reject a RESTful request which would consume excessive computational resources. Data consumers are recommended to use the FHIR asynchronous REST pattern support for which is required of data providers as specified in design paper 003 – Conceptual Design for a FHIR Proxy Server.

There are two recognised limitations with this approach.

1. The subscription and query may return overlapping data. The timing of the two transactions may be such that new data is created between the transactions being received by a data provider and so the data is included in both a subscription result and the query result set.
2. It is difficult to guarantee data completeness of a subscription made in response to receipt of a subscription result. A subscription may be lodged as result of subscription result being received. For instance, the Cancer cohort management algorithm may subscribe to new cancer diagnoses. Wishing to remove the patient from the cohort when the cancer clears it subscribes to a diagnosis cleared indicator. The following sequence of events would lead to data loss:

- i) Subscription received by data provider for cancer diagnoses.
- ii) Patient A diagnosed with cancer.
- iii) Subscription result issued by data provider for Patient A.
- iv) Subscription result for Patient A received by data consumer.
- v) Subscription to diagnosis cleared indicator for Patient A issued by data consumer.
- vi) Patient A cleared of cancer.
- vii) Subscription to diagnosis cleared indicator for Patient A received by data provider.

There are mitigations for both situations. These add complexity to the data consumer but not sufficient complexity to warrant an extension to the FHIR standard.

### 3 Processing Model for Data Providers

Design paper 003 – “A Conceptual Design for a FHIR Proxy Server” specified the technical capability which is expected of data provider. It included a processing model for a semi-productised proxy server that is intended to simplify adoption of the FHIR standard for data providers. That paper focused on a data provider’s support for generic FHIR concepts such as retrieving resources, searching for resources, and creating resources. Compliance with that paper allows a data provider to receive a *Subscription* resource and to persist it in a local FHIR repository. The paper purposefully didn’t describe process for executing subscriptions or delivering results to subscribers. These topics are the domain of this paper.

This section extends design paper 003 and should be considered by developers and implementors of FHIR Proxy Server functionality.

#### 3.1 Alignment of Data Provider Capabilities with the Maturity Model

Basic support for subscriptions is required from maturity level 1. Basic support includes the ability to:

- service subscriptions to basic facts about a patient (demographics, allergies, medications etc.) and encounters with the patient;
- support a limited set of search terms and basic comparators;
- deliver subscription results over rest hooks;
- service subscriptions within 24 hours of data being created or modified.

Support for result delivery over messaging channels is aligned with data providers’ exposing broader and more voluminous clinical content at maturity level 2.

Support for subscriptions is closely correlated with the ability to execute a FHIR search, support for search syntax, and the coverage of FHIR search terms. Capabilities in each of these areas grow with the maturity model and so subscription capability also evolves with the maturity model.

#### 3.2 Local Protections and Controls

Unrestrained support for subscriptions is potentially challenging for data providers and could, without control, lead to large computational commitments and load on operational systems.

There a number of ways that subscriptions may be implemented. Section 3.3 describes two possible models:

- asynchronous processing of a queue of resource which have been updated or modified;
- periodic querying of a data source by subscription search terms.

The two approaches have different performance profiles. The former is well suited to processing generic subscriptions with little dependency on the actual search terms used. The efficiency of the latter approach is very much tied to presence of indices which align with terms used by subscriptions.

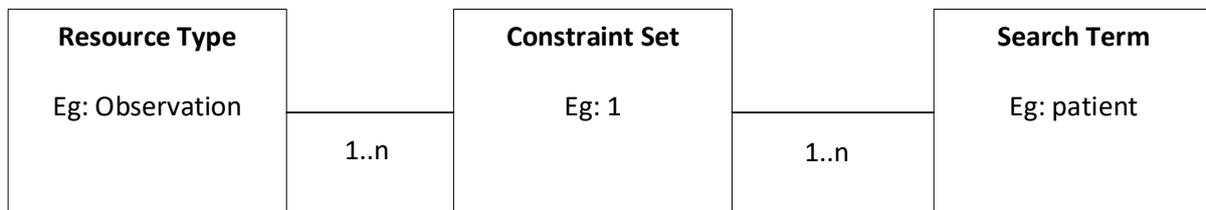
When relying on periodic querying, a data provider may choose not to process subscriptions which are too costly to process. This is consistent with real-time searching where a data provider might refuse a FHIR search request which is too computationally intensive.

Data consumers need to understand, in advance, the degree to which subscriptions will be serviced and also need to be able to target subscriptions to data points where there will be widespread support. It is also likely that data providers will wish to automate acceptance of subscriptions. These considerations point to a requirement for a formal specification of subscription capability alongside a mechanism which allows data consumers to interrogate the support offered by a data provider.

### 3.2.1 Subscription Constraints

A data provider may choose to support only a subset of the possible searches for subscription to a particular resource type. In other words, a subscription will only be accepted to a resource type if a certain combination of search terms is present in the search. For instance, a data provider may allow subscriptions to *Observation* resources if, and only if, the subject patient is present in the search. This constraint might ensure that searches for observations use an index over the patient which makes the search efficient to execute.

There may be multiple acceptable combinations of search terms for a resource type and so subscription constraints can be modelled as follows:



A search matches a Constraint Set if it includes all of the required Search Terms using an = comparator or an eq search prefix. String matching must be exact. A search will be accepted if it meets any one of the Constraint Sets for the Resource Type.

### 3.2.2 Publishing Subscription Constraints

A data consumer can interrogate the constraints which are imposed by a data provider using a FHIR [Operation](#) which specifies a resource name as an input parameter:

```
GET /Subscription/$constraints?resourceName=Observation
```

This returns all constraints that are in force for the Observation resource.

The output parameter is a *Bundle of Lists of SearchParameters*. Each list corresponds to a constraint set in the above data model. The output parameter is returned in the HTTP response body.

An empty *Bundle* infers that there are no constraints in force for the resource type and all searches will be accepted.

A *Bundle* which contains an empty *List* infers that searches are not supported for the resource type.

Data providers which apply subscription constraints must implement this operation and describe it their FHIR server's *CapabilityStatement*.

If the operation is not supported, then a data consumer can safely assume that there are no constraints on subscription search terms.

### 3.3 Processing Model for a FHIR Proxy Server

Design paper 003 – “Conceptual Design for a FHIR Proxy Server” provided an outline processing model for adding FHIR capabilities to existing data sources. The model covered retrieval, search, creation, and modification of FHIR resources for two categories of data source:

- i) data held in SQL databases;
- ii) data held as FHIR resources in a document database.

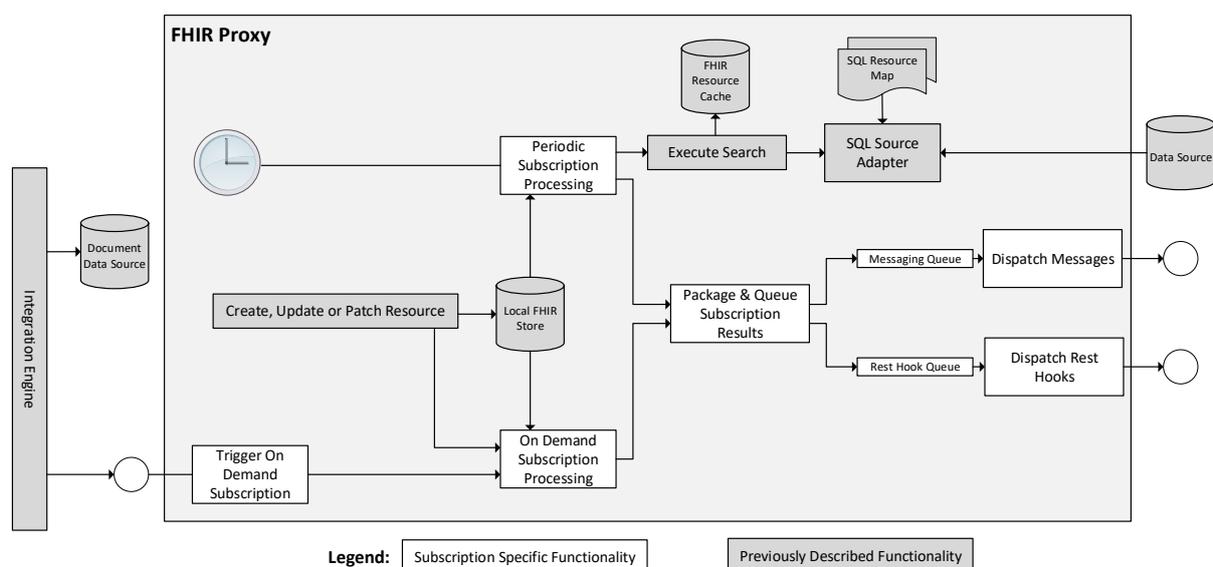
The first category is targeted at FHIR enabling live systems or an offline copy of a live system. The second assumes that data is extracted from a live system and is transformed into resources structures as part of an Extract-Transform-Load (ETL) process.

Subscriptions can operate against both types of data, each lends itself to a different processing model.

SQL databases are good candidates for a periodic processing model: subscriptions are processed according to a schedule. On each invocation of the schedule, queries are constructed which interrogate the data source for data which both matches the subscription and has been created or modified after the last time the subscription was processed.

Document database are well suited to real-time subscription processing: a hook is inserted into the resource creation/modification pathway which triggers a process which compares the resource to active subscriptions. This technique might also be applied to SQL sources where database triggers can be inserted into the database and are used to initiate subscription processing for individual resources. This situation is so similar to ETL processing that it is not specifically considered further.

The following diagram depicts the extensions required to the FHIR Proxy processing model.



The diagram only shows those FHIR Proxy components already described by design paper 003 which are relevant to subscription processing. Both processing models depend on subscription resources

which held in the Local FHIR Repository and which are maintained using the standard FHIR API endpoint.

### **3.3.1 Periodic Subscription Processing**

The process is triggered by a timer working to a schedule. It queries active subscriptions lodged in the local FHIR store. For each subscription it extracts the criteria and constructs a search request which it passes to the standard FHIR Proxy “Execute Search” pathway. Result bundles are dispatched to “Package & Queue Subscription Results” process for distribution.

Note that for efficiency, this process may want to group its query of subscription resources to ensure that multiple subscriptions to the same data point lodged by different subscribers are processed only once.

### **3.3.2 Trigger On-Demand Subscription**

A RESTful interface that allows a third-party system or interface engine to process a subscription for a single resource. The POST request will have an HTTP body which contains the resource to process. This endpoint allows updates to an external document store to be reflected in subscriptions.

### **3.3.3 On-Demand Subscription Processing**

Individual resources are matched to each subscription lodged in the Local FHIR Store. The process is executed for every resource passing through the “Create, Update or Patch Resource” pathway and for externally held resources through invocation of “Trigger On-Demand Subscription”. Logically, the process queries subscriptions lodged in the Local FHIR Store for the particular resource type (possibly grouping by criteria) and compares the resource properties to the subscriptions’ criteria query string. If it matches then it distributes the resource to subscribers using the “Package & Queue Subscription Results” process.

As discussed below, there are optimisation opportunities for this unit.

### **3.3.4 Package & Queue Subscription Results**

Resources are dispatched individually to each subscriber. Support is offered for either messaging or rest-hook channels. Each channel operates a first-in-first-out queue. Resources, with envelope data identifying the subscriber endpoint, are placed on the appropriate queue.

### **3.3.5 Dispatch Messages**

Implements a Reliable Messaging Source as described by design paper 006 – “Reliable Messaging Infrastructure”. Note that messages use the event code YH003 – “Subscription result”.

### **3.3.6 Dispatch REST Hooks**

Invokes the RESTful endpoint specified in the subscription channel. Note, that the endpoint must comply with the security requirements for the YHCR and must require that a JSON Web Token (JWT) issued by the regional Identity and Access Management (IAM) Server (design paper 005 – “Identity and Access Management”) is included in the HTTP Authorization header. This unit is responsible for obtaining and maintaining the validity of a JWT from IAM.

Requests which are rejected by the endpoint (all 4xx type HTTP responses other than 408) must be deemed to be undeliverable. The event must be logged, and the queued subscription result discarded. Requests which are temporarily undeliverable due to network failure or a 408 HTTP response should be retried after a delay.

Support for HTTP redirection through response codes 3xx is not required.

### 3.4 Validation of Subscription Management Requests

Subscriptions are created by data consumers or the regional FHIR aggregator POST'ing a *Subscription* resource to the FHIR proxy RESTful endpoint. Subscriptions are managed through PUT or PATCH requests.

The FHIR Proxy must validate subscription management requests and reject those which are unacceptable.

For new subscriptions validation must include:

- the FHIR criteria complies with any subscription constraints;
- the subscription channel type is either a rest hook or messaging and is supported at the data provider's maturity level;
- the subscription channel endpoint has been whitelisted for the deployment.
- the subscription status is "Requested".

The status of a *Subscription* which is accepted must be set to "Active" by the FHIR Proxy.

The only modification which is permitted to a *Subscription* by an external source is to move its status from "Active" to "Off".

A FHIR Proxy may choose to archive or purge *Subscriptions* with status "Off".

### 3.5 Processing Model Computational Costs

The computational cost of processing subscriptions varies with a number of factors, but the main cost drivers are the number of updated/new records in a period (N) and the number of subscriptions (M). Costs vary for the different processing models.

Assuming that i) the number of created or updated records is small relative to the overall size of the data source, and ii) the data is well indexed by FHIR search path then the cost for periodic processing against a SQL data source is:

$$\sim M \cdot \log(N)$$

Whereas for unoptimized on-demand processing of new or modified resources the cost is:

$$\sim M \cdot N$$

A possible optimisation for on-demand processing to index comparator values in subscription search terms. If most subscriptions are targeted at individual patients, then indexing the patient identifier in search terms changes the cost to:

$$\sim N \cdot \log(M)$$

In other words, in the absence of other factors, situations where the number of new or updated records is much greater than the number of subscriptions favour the periodic processing model. Whereas a large number of well targeted subscriptions are best processed on demand.

Options for scheduling of subscription processing can outweigh processing cost considerations. But both models have flexibility in this regard: periodic processing can be time-sliced to be near

continuous or can concentrate processing costs within a small time-window. Similarly, on demand processing can occur in real time or be messaged to a queue for perioding processing.

The model regional FHIR Proxy must allow deployment specific control over all options.

---

## 4 Regional Subscription Manager Technical Design

### 4.1 Creating Regional Subscriptions

Subscriptions are created using the RESTful interface of the regional FHIR Aggregator Service (design paper 010). Unless otherwise specified, it is implied that regional subscriptions apply to all data providers. A consumer can target a subscription at a specific provider by including a source tag in the search term i.e.:

```
Tag=https://yhcr.nhs.uk/Source|[provider]
```

Where [provider] is the data provider's registered name with the YHCR.

The regional subscription manager propagates subscriptions to relevant data providers including the YCHR FHIR Store (design paper 018).

Subscriptions targeted at specific providers are propagated to those providers regardless of the subscription criteria.

Subscriptions are only propagated to subscribers at an appropriate maturity level. Note that the onboarding process for providers (design paper 020 "Onboarding for Data Providers") can result in subscriptions which have been previously lodged regionally being propagated to the newly onboarded provider.

The Regional Subscription Manager maintains a map relating subscriptions made regionally to those propagated to data providers.

A subscription will not be issued more than once to a data provider. Different consumers may make the same subscription. In such cases, the subscription results will be relayed to both subscribers, but the data provider will not be aware of the second and subsequent occurrences.

Note that every consumer's subscriptions exist as unique resources in the regional FHIR store. *Linkages* resources relate these subscriptions to subscriptions propagated to local data providers.

#### 4.1.1 Criteria which Reference Patients

Subscriptions which reference patients are only propagated to those providers who have registered contact with the patients with the Regional Identity Exchange (design paper 004). Patient references to regional *Patient* resources are substituted in the propagate subscription with references to local *Patient* resources at the provider making use of the *Linkage* resources maintained centrally (see design Paper 001 – "A Unified Distributed Data Model for FHIR").

A FHIR subscription criterion might reference more than one patient. If the criterion implies that all patients must be known to a provider for the subscription to match data points, then the subscription will only be propagated to those providers for which there are *Linkages* to all patients.

If the criterion implies that any patient may be known by the provider, then the subscription will be propagated to providers that have had contact with any of the patients. The criterion will be rewritten to reference only those patients for which *Linkages* exist for these providers.

On an ongoing basis, providers have contact with new patients. Each contact is registered with the Regional Identity Exchange. A new registration may result in new subscriptions being propagated to

the provider. For each new contact the Regional Subscription Manager searches subscriptions which reference the patient and issues:

- i) A query composed from the subscription criteria;
- ii) A new subscription to the provider.

Query results from i) are treated as subscription results and are distributed by the Regional Subscription Manager to subscribers.

#### **4.2 Subscriber Identity Management, Security and Audit**

A data provider must obtain authorisation from the regional Identity and Access Management (IAM) server to be able to create a subscription. The process of making a claim to IAM is covered by design paper 005 – “Identity and Access Management”. The reason for access in the claim determines the scope of access rights. The same rules, as specified by design paper 005, which relate to the release of patient identifiable data apply to subscription rules as to direct searches. Subscriptions will be rejected which are inconsistent with the reason for access.

Release of data for indirect care require patient consent. The same consent processing applies to data released through subscription as released through direct query. Consent processing is detailed by design paper 008 – “Consent Management”.

Audit records are written for subscriptions created regionally and at local data providers, and for all data released as a result of subscriptions. Audit records are also written for rejected subscriptions and data withheld because of consent rules.

#### **4.3 Cohort Management**

A patient cohort is maintained in the regional FHIR Store as a *List* resource. Lists can be created and managed using the standard FHIR APIs on the Regional FHIR Aggregator Service.

Lists can be referenced in subscription criteria and FHIR searches made against the regional aggregator using the “in” modifier. Note that this is a custom extension to FHIR which is only supported by the YHCR. The following example searches for care plans for all patients in a particular list:

```
/CarePlan?patient:in=https://yhcr.nhs.uk/List/f08d70c0-a711-4e0d-993e-542fb3e3ca9e
```

Referencing a list in a subscription acts as an instruction to the Regional Subscription Manager to issue individual subscriptions to every patient in the list to local data providers. It also causes the Regional Subscription Manager to act on changes to the list membership:

- if a patient is removed from the list then subscriptions are revoked by data providers;
- if a patient is added to the list then new subscriptions are propagated to data providers.

Cohorts may also be referenced in direct searches made against the FHIR Aggregator. The operation of the aggregator is considered in design paper 010 – “FHIR Aggregator Service”.

#### **4.4 Rest Hooks and Messaging Channel Mediation**

The Regional Subscription Manager will use an appropriate channel type for its own purposes for subscriptions propagated to data providers. This may be different from the channel type for subscriptions made on it by data consumers.

The Regional Subscription Manager implements a messaging and queuing system which allows subscription results received over a rest hook to be reissued to a consumer requiring a messaging delivery and vice versa.

The Regional Subscription Manager behaves as a Reliable Messaging Source (design paper 006 – “Reliable Messaging Infrastructure” when messaging subscription results.

#### **4.5 Capability Discovery**

Understanding the capabilities of data providers is key to data consumers interacting with them in a predictable manner. Knowledge of the alignment of data providers with the maturity model is fundamental to data consumers and registration details are published from the regional IAM service (design paper 021 – “Onboarding Data Providers”). The Regional Subscription Manager also implements the subscription \$constraints FHIR operation described for data providers in section [3.2.2](#).

The output of the operation is the same as when invoked locally: a *Bundle* of *Lists* of *SearchParameters* which represent constraints imposed on subscriptions by a data provider. In common with all resources obtained through regional services, the provenance of each resource can be located to a data provider from a ‘Source’ tag attached to the resource. Note, that in this case, the source of each List will be the same as the source of the *SearchParameters* that it contains

#### **4.6 Revoking Subscriptions**

A regional subscription can be revoked by a consumer by modifying its “status” property to “Off”. The modification must be made by the consumer which created the subscription. This is the only management operation possible.

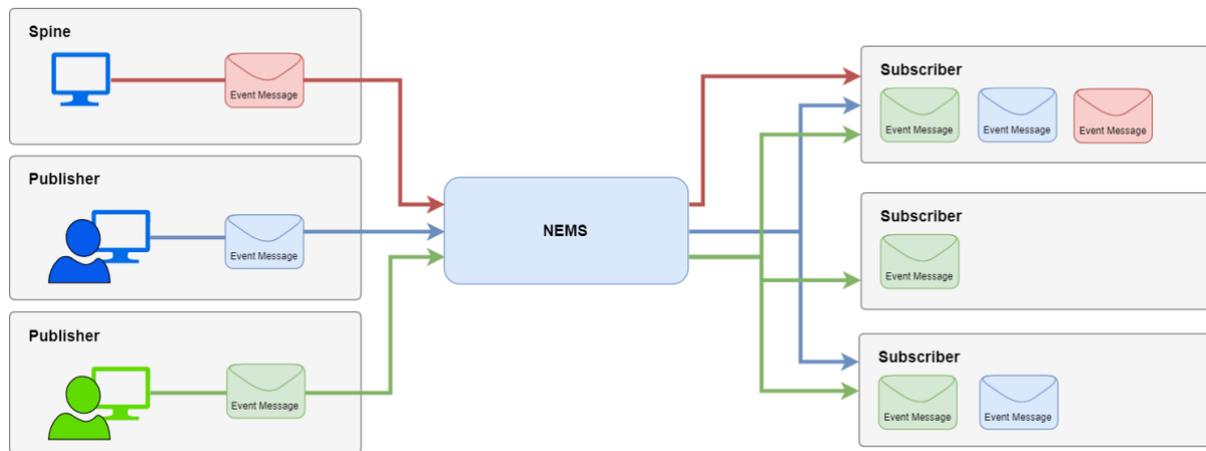
If this the only or last active subscription which is linked to subscriptions propagated locally, then the Subscription Manager will revoke local subscriptions.

Note the timing of revocations cannot be relied on. There may be in-flight subscription results which will be delivered to a data consumer after their associated subscription has been revoked. Data consumers must be designed to accommodate this situation.

## 5 Subscriptions and the National Event Management Service

NEMS is a service operated by NMS Digital which allows data about certain patient related events to be published to the service and for the subscribers to the service to be informed of the events as they occur.

A schematic which is used by NHS Digital to describe the service is as follows:



The service uses FHIR messaging to receive publication of events and to notify subscribers.

There is a clear overlap between NEMS and the YHCR Subscription infrastructure.

It may be possible for the Regional Subscription Service to register subscriptions in NEMS when a data consumer subscribes locally to a concept which is supported by NEMS.

NEMS is currently in its infancy and only supports events relating to failsafe alerting in pilot form and this paper proposes NEMS integration as a future possibility rather than seeking to direct immediate implementation.

NEMS uses FHIR Subscriptions but currently only supports subscriptions to *Bundles* as opposed to raw FHIR resources. This implies that subscriptions at a national level are to pre-defined packaged information about patients. The content of the package is determined by an event code that is specified in the *MessageHeader* of the *Bundle*. An example of a criteria in subscription made to NEMS is expressed is:

```
/Bundle?type=message&serviceType=UHV&Patient.identifier=http://fhir.nhs.net/Id/nhs-number|9434765919&MessageHeader.event=PDS002
```

The NEMS *Subscription* resource profile is based on the HL& STU3 resource and the few constraints which it introduces can be easily accommodated by the YHCR.

NEMS subscription results are delivered over a MESH channel. Support for MESH is considered in design paper 006 – “Messaging Infrastructure” and the Regional Subscriptions Manager could act as a mediator between MESH and rest-hooks or messaging delivery using the reliable messaging infrastructure.

The YHCR is well positioned to behave as a proxy to data consumers in the region making subscriptions in the above format to NEMS.