INTERWEAVE

CONNECTING CARE

Cookbook for Regional Interoperability Detailed Design Paper #020

Onboarding for Data Providers

PRELIMINARY DRAFT

Version 1.1 – 23rd December 2019

Abstract Interoperability Cookbook Anchor Points

Section	Title
4	Requirements for Data Providers

Table of Contents

Version Control

Version	Release Date	Released By	Reason for Release
1.0	29/09/2019	R Hickingbotham	Preliminary draft
1.1	23/12/2019	R Hickingbotham	Improvements to content

Reviewers

Initials	Name	Role	Organisation

Release Notes for this Version

1) Improved definition of the DataProvider resource

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 020 - "Onboarding for Data Providers") defines the process for bringing a new data provider onto the System-Of-Systems. It covers the steps which are followed from when a potential data provider first approaches the YHCR up to the point where there is a live connection between SoS and data relating to active patients/clients is flowing to data consumers.

The paper is written from the perspective of the operators of the System-of-Systems. The technical, governance, or transformational activities undertaken by the data provider are crucial to prepare an organisation for participation but other design papers deal with these considerations such as 004 – "Conceptual Design for a FHIR Proxy Server" or 014 – "Governance for Data Providers". This paper focuses on the actions which will be taken by operators to:

- establish the identity of the data provider and its responsible officers;
- create and secure TCP/IP connections with the provider;
- provide access to the provider to a series of environments with increasing requirements for compliance with technical standards;
- validate the providers self-assessment of maturity;
- test the providers adherence with mandatory governance requirements;
- help the provider certify the quality of its data.

1.2 Onboarding and Risk Management

Onboarding is a key risk activity for the YHCR. There is the potential to compromise both security and data integrity. The SoS will ultimately be hosted in the open internet and its security is wholly dependent on the integrity of the public key infrastructure and consistency of the configuration of firewalls, DNS, and the participant registry. Risks are mitigated by automating the onboarding process and minimising the need for manual configuration of components. This paper includes specifies requirements for the YHCR Onboarding Suite: a software product which automates the workflow underpinning the onboarding process and manages configuration of the resources responsible for security.

1.3 The Participant Registry

The participant registry is an important source of configuration data for the SoS. It identifies the data providers and data consumers which have access to each of the SoS environments and contains all the details needed to manage interactions with the participants. This paper specifies the data elements which are managed for data providers which is hence termed the provider registry. A parallel design paper 021 – "Onboarding for Data Consumers" details participant registry requirements for data consumers.

1.4 Relationship of this Document with Other Standards

This paper does not require knowledge of any particular standard.

1.5 Intended Users of the This Document

Operators of the YHCR, developers of the regional onboarding suite.

2 Environments of the System of Systems

The SoS operates several semi-autonomous environments. These are designed for two distinct purposes:

i) to enable the System-of-System developers to develop, test, and deploy software code;
ii) to help participants to develop, assure, and take live compatible technology.

2.1 Types of Environment

Different sets of environments are provided for developers of the System-of-Systems, and for participants to develop and test participating technologies. Both sets operate on parallel principles:

- code/connecting technologies are developed and tested in Development/Sandpit environments;
- code releases/ connecting technologies which are candidates for live operation are migrated to a System Testing/Staging environment where they are tested at scale;
- code releases/ connecting technologies which pass system testing/accreditation are promoted to a Live environment.

A consequence of these principles is that the environments used by developers will at times be operating different versions of SoS code.

In contrast, all of the environments which are used by participants to develop and test compatible technology run identical System-of-System code so providing a consistent pathway to live. Each of the participant environments is updated whenever new code is promoted to Live.

Code releases, and onboarding tasks should be synchronised to minimise risk and disruption to participants. Specifically:

- releases must be tested in the System Test environment with connections to all data providers who have a presence in the Live environment;
- the migration of participants to Live must be suspended once testing of a release candidate has commenced in System Test;
- the impact of a code release applied to Live must be assessed against the participants working in the Staging environment. If there is risk then the participants should regress to retest in the Sandpit.

The environment topology is illustrated below:

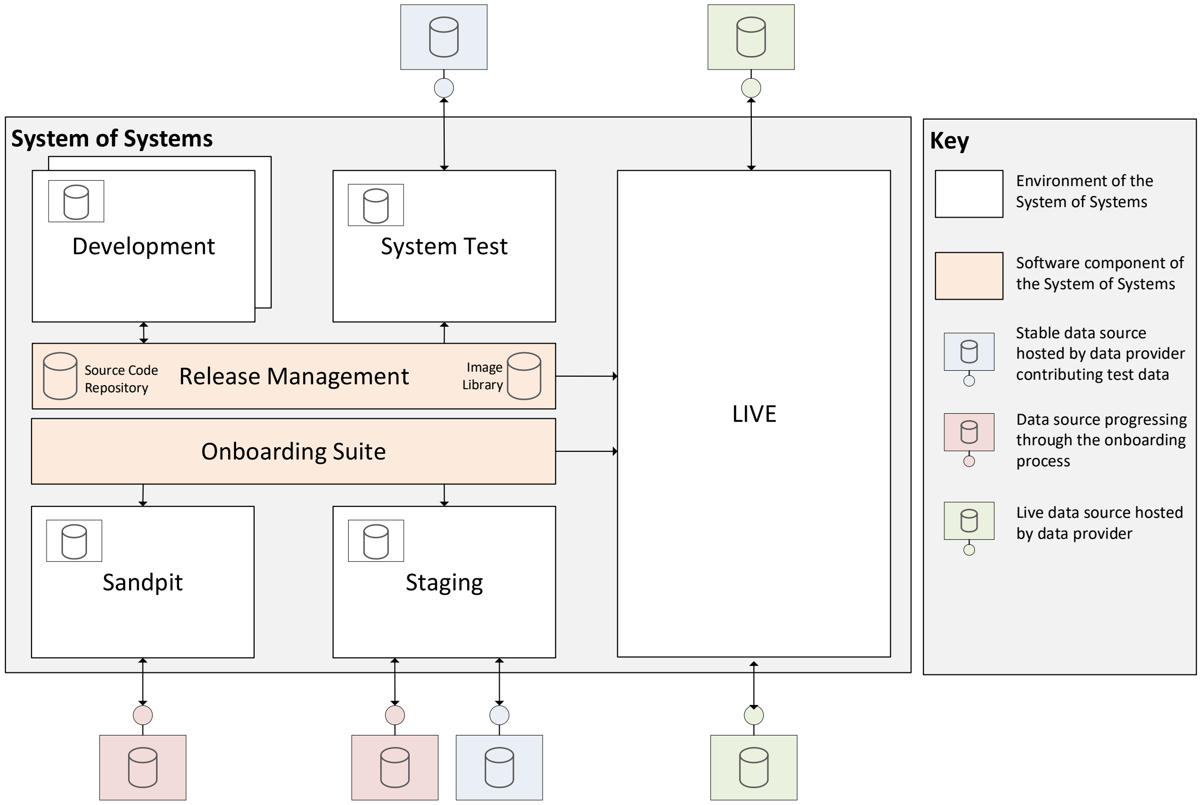
System of Systems
Development
System Test
Source Code Repository
Release Management
Image Library
Onboarding Suite
Sandpit
Staging
LIVE
Key
Environment of the System of Systems
Software component of the System of Systems
Stable data source hosted by data provider contributing test data
Data source progressing through the onboarding process
Live data source hosted by data provider

Release management software controls the movement of SoS code from one environment to another. It draws code from an image repository.

The movement of compatible technology from one environment to another is controlled by an onboarding suite.

2.2 Dependencies on Data Providers

The System Test and Staging environments are intended to be very close reproductions of the Live environment in that they will enable interactions with the same data providers and data consumers as in Live but using test data. Data providers will need to support this configuration by maintaining access to test endpoints following acceptance into live operation. Test endpoints should be capable of serving a high volume of test data relating to a population of test patients; demographics for which will be provided by the operators of the System-of-Systems. Note that test endpoints connected to the System Test and Staging environments may be the same instance of the provider's technology providing access to the same data.

2.3 A Data Provider's Pathway to Live

2.3.1 Sandpit

On registration with the System-of-Systems, a data provider gains access to a Sandpit environment. The environment is configured to enable the registrant to connect a FHIR RESTful server to act as a data source to a test suite which itself acts as a data consumer. On connection, automated tests begin executing against the data source. These tests are correlated with the participant's stated target maturity level. A dashboard helps the participant visualise progress to full test coverage.

The participant also has query access to the environment's FHIR aggregator. Queries can be executed using a REST interface simulation tool such as Postman. The connection is secured over HTTPS and a certificate signed by the YHCR is required to use it. The environment is configured to aggregate data from a number FHIR endpoint simulators, and the participants own data source. The FHIR endpoint simulators supply representative test data and are a useful resource for data providers in determining expected behaviour.

Data from different data providers working in the sandpit is segregated. Only data from simulators and the participant's own data source is accessible by the participant through the sandpit aggregator.

If the data provider is also onboarding as a data consumer, then access to the same test aggregator, and with the same privileges over data, applies to the participant's role as a consumer as to its role as a provider. This means that the participant is able to view its own data and test both its provider and consumer functions in parallel. The use of environments by data consumers is detailed in design paper 021 – "Onboarding for Data Consumers".

2.3.2 Staging

A data provider can enter a Staging environment once all automated tests have completed in the Sandpit for the maturity level targeted by the participant.

Tests in the Staging environment simulate transaction volumes at production levels and are designed to test performance and resilience. Tests also validate the security of the data provider's endpoint and compliance with governance requirements.

2.3.3 Live

A connection to the live environment might be the first opportunity that data provider has to clinically validate the quality of the live data that they will be providing to the SoS. Settings in the Participant Registry can restrict access to data to certain data consumers and this facility is used to provide the participant with an opportunity to review live data supplied from their endpoint prior to releasing data more generally. A SoS supplied data browser is given sole access to the data source and can be used by a clinician to validate live data.

Once validated, the data source is made available more generally to live consumers of the SoS. A soft go-live can be supported whereby data is selectively released to nominated data consumers and staged over time.

2.4 Autonomy of Environments

Each environment operates its own private IP network and hosts a complete set of SoS components. There is no possibility of direct connection from one environment to another.

Each environment operates a domain name server and is responsible for resolving sub-domains in the form:

env-name.yhcr.nhs.uk

Provider endpoints are allocated a DNS address from the environments from which there is connectivity. If a provider endpoint serves more than one environment (i.e. System-Test and Staging) then it is allocated a DNS address from each.

Certificates for HTTPS connections ae issued from relevant environments at the point the data provider is granted access to them.

Separate firewalls/routers for each environment bridge public and private IP address.

Each environment manages its own participant registry. A data provider must be registered in the environment for transactions to be routed to the participant.

Features which are common between environments include:

- code images are drawn from a common image repository and a shared release manager promotes code into each environment;
- a shared onboarding suite controls participants migration through environments;
- shared monitoring and alerting tools monitors performance and probity of environments and report adverse conditions;
- common security software monitors threats.

3 The Onboarding Workflow

Onboarding workflow is controlled by the Onboarding Suite. It's function is to ensure that a process is rigorously followed and to automate configuration of components which are responsible for the security of the YHCR.

3.1 Anatomy of a Data Provider

A data provider implements endpoints which the System of Systems connects for:

- synchronous query;
- asynchronous result polling and retrieval;
- resource management (including placement of subscriptions).

Only the synchronous query endpoint is mandatory.

A data provider can optionally connect to the SoS from pre-defined static IP addresses to deliver:

- subscription notifications;
- reliable messaging transactional messages.

3.2 Onboarding Steps

Identify Participant	1	Register prospective participants to the YHCR in the Onboarding Suite.
	2	Obtain and record details of senior officers, contact methods and domain name. Corroborate details from an independent source and record method of corroboration. All email addresses must be at the registered domain name for the organisation.
	3	Identify senior officer responsible for participation in the YHCR (SRO).
	4	Receive notification of interest to participate in the YHCR.
		Record details of the product or service to be onboarded.
	5	Obtain and record contact details of technical personnel responsible for onboarding. Corroborate participation and technical contacts through an exchange of email with the SRO.
	6	Generate online accounts on the onboarding suite for technical officials.
	7	Generate a synonym for the provider which will be used as a uniform identifier by the SoS.
	8	Provide access to Onboarding suite portal to technical officials. Capture password to be use in certificate signing requests.
Sandpit Developmen	9	Capture IP addresses and port numbers of the endpoint of the participant's FHIR development server. Separate IP addresses may be provided for: • synchronous query; • asynchronous result polling and retrieval; • resource management (including placement of subscriptions);

		• reliable messaging delivery point; • reliable messaging message source (IP address only).
	10	Obtain electronic signature verifying IP addresses from technical official.
	11	Corroborate veracity of IP addresses with SRO
	12	In the Sandpit environment: • generate DNS address records for the development FHIR server IP addresses; • configure outbound firewall rules to permit access to server addresses/ports and an inbound rule to permit access to the SoS message delivery service from the reliable messaging message source; • register the data provider in the Sandpit participant registry as an inactive participant.
	13	Prove IP connectivity between the Sandpit environment and data provider endpoints.
	14	Accept server and client certificate signing requests.
	15	Confirm receipt of certificates with SRO.
	16	Apply certificates to Sandpit Participant Registry.
	17	Determine and record target maturity level and apply to Participant Registry.
	18	Populate PIX/MPI with linkages representing patient contact at participant with the test patient cohort.
	19	Activate participant in Participant Registry.
	20	Begin automated testing.
	21	Participant confirms testing is complete.
	22	Participant completes governance checklist for data providers.
	23	Confirm progression to Staging with SRO.
Staging	24	Capture IP addresses and port numbers of the endpoint of the participant's FHIR staging server. Separate IP addresses may be provided for: • synchronous query; • asynchronous result polling and retrieval; • resource management (including placement of subscriptions); • subscription notification source (IP address only); • reliable messaging message source (IP address only). Note that the participant may elect to migrate the same servers from development to staging in which case these IP addresses may be the same as in step 8.
	25	Obtain electronic signature verifying IP addresses from technical officer.
	26	Corroborate veracity of IP addresses with SRO
	27	In the Staging environment: • generate DNS address records for the development FHIR server IP addresses; • configure outbound firewall rules to permit access to server addresses/ports and an inbound rule to permit access to the SoS message delivery service from the reliable messaging message source;

		register the data provider in the Staging Participant Registry as an inactive participant.
	28	If the same servers are being used in Staging as in the Sandpit, then deactivate the participant in the Sandpit environment.
	29	Prove IP connectivity between the Sandpit environment and data provider endpoints.
	30	Accept server and client certificate signing requests.
	31	Confirm receipt of certificates with the SRO.
	32	Apply certificates to the Sandpit Participant Registry.
	33	Populate PIX/MPI with linkages representing patient contact at participant with the staging patient cohort.
	34	Activate participant in the Sandpit Participant Registry.
	35	Execute pipe cleaning tests aimed at proving functional connectivity.
	36	Consume capability statement an validate against maturity level alignment.
	37	Schedule volume testing.
	38	Schedule assurance testing.
	39	Execute volume tests.
	40	Execute assurance tests.
	41	Confirm progression to Live with the SRO.
System Test Representation	42	Capture IP addresses and port numbers of the endpoint of the participant's FHIR system test server. Separate IP addresses may be provided for: • synchronous query; • asynchronous result polling and retrieval; • resource management (including placement of subscriptions); • subscription notification source (IP address only); • reliable messaging message source (IP address only). Note that the participant may elect to use the same servers in staging and system test in which case these IP addresses may be the same as in step 20.
	43	Obtain electronic signature verifying IP addresses from technical officer.
	44	Corroborate veracity of IP addresses with SRO
	45	In the System Test environment: • generate DNS address records for the development FHIR server IP addresses; • configure outbound firewall rules to permit access to server addresses/ports and an inbound rule to permit access to the SoS message delivery service from the reliable messaging message source; register the data provider in the Staging Participant Registry as an inactive participant.

	46	Accept server and client certificate signing requests for System Test.
	47	Confirm receipt of certificates with the SRO.
	48	Apply certificates to the Live Participant Registry.
	49	Activate participant in the Live Participant Registry with access restricted to a test data consumer
	50	Execute pipe cleaning tests aimed at proving functional connectivity.
	51	Accept server and client certificate signing requests.
	52	Confirm receipt of certificates with the SRO.
	53	Apply certificates to the System Test Participant Registry.
	54	Activate participant in the System Test Participant Registry.
	55	Execute pipe cleaning tests aimed at proving functional connectivity.
	56	Populate PIX/MPI with linkages representing patient contact at participant with the system test patient cohort.
	57	Activate participant in the System Test Participant Registry
Migrate to Live	58	Capture IP addresses and port numbers of the endpoint of the participant's FHIR live server. Separate IP addresses may be provided for: • synchronous query; • asynchronous result polling and retrieval; • resource management (including placement of subscriptions); • subscription notification source (IP address only); • reliable messaging message source (IP address only). Note that the Sandpit servers cannot be repurposed for live operation and must remain accessible from the Sandpit environment.
	59	Obtain electronic signature verifying IP addresses from technical officer.
	60	Corroborate veracity of IP addresses with SRO
	61	In the Live environment: • generate DNS address records for the development FHIR server IP addresses; • configure outbound firewall rules to permit access to server addresses/ports and an inbound rule to permit access to the SoS message delivery service from the reliable messaging message source; • register the data provider in the Staging Participant Registry as an inactive participant.
	62	Confirm with the participant the plan for batch on-take of historic patient contact.
	63	Accept server and client certificate signing requests for Live.
	64	Confirm receipt of certificates with the SRO.
	65	Apply certificates to the Live Participant Registry.

	66	Activate participant in the Live Participant Registry with access restricted to a test data consumer user interface through which the participant can browse live data.
	67	Execute pipe cleaning tests aimed at proving functional connectivity.
	68	Suspend all outgoing connection points to the data provider (subscriptions will be generated as historic patient contact is loaded and these should queue in the SoS).
	69	Batch load historic patient contact into PIX/MPI.
	70	Register an instance of the data browser as a data consumer with access restricted to data provided by the new participant.
	71	Set up accounts with the data browser as instructed by the participant.
	72	The participant reviews live data and completes a data compliance test.
	73	Schedule go-live date and soft go-live option.
	74	Confirm data compliance review results and go-live date with SRO.
	75	Decommission data browser instance and de-register the data consumer.
	76	Lift access restrictions in the Live Participant Registry in accordance with the go-live plan. Note that f the participant is re-entering the onboarding process in order to raise maturity then there may be an existing live service in operation for the data provider. Access restrictions lifted on this more mature service must be made in conjunction with parallel restrictions being placed on the less mature service effectively transiting the data provider from one service to another.
	77	Open outgoing connection points and monitor queue release.

3.3 Requirements of the Onboarding Suite

The onboarding suite is a software product which controls the process detailed above and interfaces with other components of the SoS to manage their configuration and ensure consistency of access controls in each environment.

3.3.1 Enforce Workflow

The Onboarding Suite ensures sequential flow through the steps detailed in 3.1.

It enables controlled change to contact and connection details. It ensures that changes do not impact the integrity of the workflow or security of the relationship with the participant. Changes to technical contact points in an organisations and details of the SRO require independent corroboration. Changes to IP addresses and receipt of resigned certificates must be authorised by the SRO before they are accepted by the Onboarding Suite and applied to components controlling security.

The target maturity level can be adjusted whilst the data provider is in the Sandpit environment but is locked thereafter.

The Onboarding Suite allows the workflow to be regressed to key control point effectively enabling progression through the environments to be backed out. The Onboarding Suite negates all configuration changes made to impacted environments and revokes certificates issued after the regression point. Revoked certificates cannot be used to gain access to an environment or a component within an environment. Key regression points include:

- steps confirming progression to the next environments so moving a participant out of an environment;
- receipt of notification of interest to participate in the YHCR so cancelling the participant as a data provider.

3.3.2 Segregate Duties

Role based access control must separate responsibilities for:

- administering users from all other functions;
- registering a participant from recording independent corroboration of its domain name;
- registering details of the SRO from independent corroboration of those details;
- registering details of a technical contact from recording of acknowledgement of the contact from the SRO;
- recording IP addresses from recording corroboration from the SRO of IP addresses and signed certificates.

Lifting access restrictions in the Live Participant Registry should be a dedicated senior role.

Individuals with application level access must not have system database administrative rights to the infrastructure running the Onboarding Suite.

3.3.3 Access for Technical Officers

Technical Officers at participant have access to the onboarding suite to:

- set and reset a password used in a certificate signing request (this action requires corroboration from the SRO);
- enter IP addresses for server's connection with environments;
- view the organisational details;
- view the configuration of all environments pertinent to their connection;
- view the status of the onboarding workflow.

3.3.4 Manage Component Configuration

A key role of the Onboarding suite is to manage the configuration of component which contribute to the security of the YHCR. The configuration of the following components will, as exclusively as is feasible, will be managed by the Onboarding Suite:

- entries in the participant registry;
- domain name server records;
- firewall configuration;
- certificate key chain.

Any manual configuration changes will be logged and documented by a change control board.

Note that the certificate key chain is used by all System of System components to enforce the public key infrastructure detailed in design paper 016 – "Securing the YHCR"

3.3.5 Log Access and Data Management

The Onboarding Suite logs all user access with session details and actions undertaken,

3.3.6 Alerting and Reporting

A management console, which is available to both participants and operators graphically summarises the number participants that have a presence in each environment, the time spent in the environment and time spent in the current environment.

Alerts escalate tasks that have be in the hands of operator for greater than a configurable period.

3.4 Raising Maturity Levels and Re-entering the Onboarding Process

A data provider onboards at a maturity level which defines the data coverage which can be expected from the data provider and the technical capability of their endpoints. Maturity levels are defined by design paper 023 – “The YHCR Maturity Model”.

It is anticipated that data providers will increase their maturity over time and may, at times, look to upgrade an operational service which has been previously onboarded to the YHCR.

The onboarding process allows an organisation to onboard and operate multiple ‘products’. An enhanced capability should be onboarded as a separate product which is associated with the product that it is replacing. The original product will be superseded as the newly onboarded product enters live and the Onboarding Suite will remove data consumers’ access to the superseded product in the live environment in line with access being released to the new product.

This process requires the data provider to operate different endpoints for different product versions during the course of onboarding a new version.

4 The Participant Registry for Data Providers

The provider registry covers both data consumers and data providers. Requirements here relate only to data providers; a parallel definition is provided by design paper 021 for data consumers.

4.1 The *DataProvider* Resource

The data model described below is used by components of the SoS to control the relationships between data providers and data consumers. It can also be accessed by consumers of the YHCR that interested in determining the registration details of other participants.

The data is not a FHIR resource, but it uses the notation and inheritance of model of FHIR and therefore can be served from FHIR server which is configured with an appropriate schema.

The data structure represents a data provider which may be associated with an Organisation which is also known to the YHCR as a health & care provider. Normally only one data provider will be associated with an organisation. The data provider can operate one or more products that each have a presence in one or more environments. In each environment the product must have a FHIR endpoint (the endpoint that serves standard synchronous FHIR queries) and may have endpoints for asynchronous queries and data management. The provider may also support messaging and subscription interaction patterns, in which case it also registers client applications that connect to the SoS. The details of endpoints and client applications may change over time and can be superseded.

The participant registry for data providers is defined as follow:

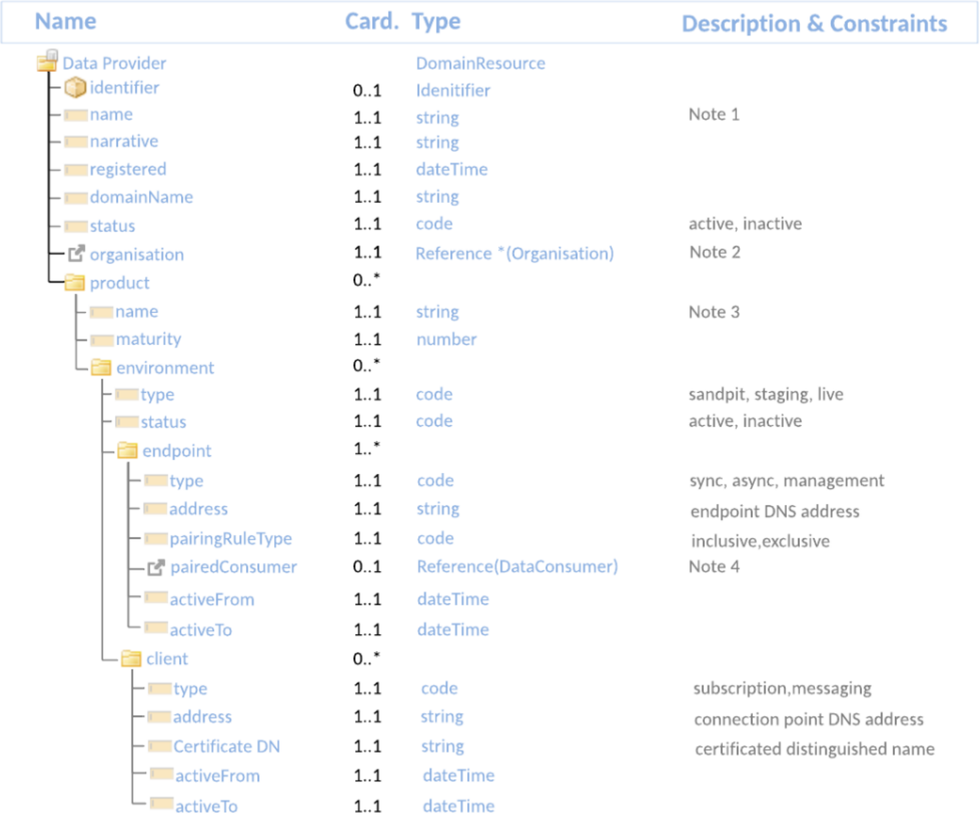

Name	Card.	Type	Description & Constraints
Data Provider		DomainResource	
identifier	0..1	Idenitifier	
name	1..1	string	Note 1
narrative	1..1	string	
registered	1..1	dateTime	
domainName	1..1	string	
status	1..1	code	active, inactive
organisation	1..1	Reference *(Organisation)	Note 2
product	0..*		
name	1..1	string	Note 3
maturity	1..1	number	
environment	0..*		
type	1..1	code	sandpit, staging, live
status	1..1	code	active, inactive
endpoint	1..*		
type	1..1	code	sync, async, management
address	1..1	string	endpoint DNS address
pairingRuleType	1..1	code	inclusive,exclusive
pairedConsumer	0..1	Reference(DataConsumer)	Note 4
activeFrom	1..1	dateTime	
activeTo	1..1	dateTime	
client	0..*		
type	1..1	code	subscription,messaging
address	1..1	string	connection point DNS address
Certificate DN	1..1	string	certificated distinguished name
activeFrom	1..1	dateTime	
activeTo	1..1	dateTime	

Notes:

(1) The name of the data provider is unique among *DataProvider*. An ODS code is preferred but not mandated.

(2) The association to an *Organisation* is optional. Whilst most of the organisations operating a data provider might also be expected to be providers of health and care services there may be exceptions. If it is supplied, then it must to be a reference to an Organisation resource held by the YHCR (reference design paper 013 – "Interfaces with the Organisational Data Service (SDS)". The onboarding suite may choose to use details of a FHIR resource to store organisational details but it is not mandated and the potential operational benefits of decoupling the definition of an organisation and its officers from the perspective of onboarding to that from the perspective of a health & care provider.

(3) The name of the product is unique within the *DataProvider*. In combination with the provider name, the product name uniquely identifies the data source and is used in the meta-data source tag to identify the provenance of FHIR resources. A data source is specified as:

[provider name]-[product name] ie: RV9-FHIR Bus

(4) Paring rules allow the FHIR aggregator to restrict access to a particular data provider to certain data consumers. Rules are either inclusive (only identified data consumers have

access to the provider) or exclusive (all consumers have access to the provider except those listed).

Search parameters comprise:

Name	Type	Description	Expression
name	string	The name of the data provider	DataProvider.name
identifier	token	The unique id for a particular data provider	DataProvider.identifier
status	code	The status of a data provider	DataProvider.status
pairedConsumer	reference	Data consumers that are paired with this data provider	DataProvider.product.environment. endpoint.pairedConsumer (DataConsumer)
organisation	reference	The organisation that operates the data provider	DataProvider.organisation (Organisation)
productName	string	The name of a product operated by a provider	DataProvider.product.name
environment	code	The environment in which the product is operating	DataProvider.product.environment.type
connectionType	code	The status of a product in an environment	DataProvider.product.environment. endpoint.type | DataProvider.product.environment. client.type

4.2 APIs for Interacting with *DataProviders*

DataProvider resources will be persisted in a regional FHIR Store and will be accessible from a URL published in the Operations Guide by regional components and any consumer of the YHCR using standard FHIR STU3 search syntax.

Resources will be managed by the Onboarding Suite. Management APIs will not be available to any other component or participant.

Note that the content of a DataProvider resource should be based on the environment from which it was obtained:

- references to *Organisations* must resolve to the organisation within the same environment as the *DataProvider;*
- only one environment branch should be populated – that which corresponds to the environment from which the *DataProvider* was obtained.

4.3 Use of the *DataProvider* Resource by YHCR Components

All YHCR components manage data structures which involve data providers by reference to a DataProvider resource or to the combination of a provider name and product name.

4.3.1 FHIR Aggregator

The FHIR Aggregator restricts access to those data providers who have an inclusive pairing with the *DataConsumer* (design paper 021) who is requesting data and those data providers which have an exclusive pairing but don't exclude relationships with the consumer.

FHIR interactions are targeted at the appropriate active endpoint registered for the data provider.

4.3.2 Subscription Manager

Subscriptions are created on the active data management endpoint address for the data provider.

Subscriptions notification connections are accepted from the registered subscription client. The connecting IP address is validated as resolving for the registered domain name.

The Distinguished Name in the certificate used in establishing the TLS connection is validated against the active subscription client registration.

4.3.3 Reliable Messaging

Messages are accepted from the registered messaging client. The connecting IP address is validated as resolving for the registered domain name.

The Distinguished Name in the certificate used in establishing the TLS connection is validated against the active messaging client registration

5 Automated Compliance Testing

Automated compliance tests validate that a data provider meets the security and governance requirements of the YHCR, and its capability is aligned with its publicised maturity level.

Compliance testing begins in the Sandpit environment where tests focus on alignment with the maturity model. Maturity testing proves that:

- expected resource types can be retrieved for a set of test patients and if a resource type is not available an appropriate *OperationOutcome* resource is inserted into the search results (refer to design paper 017 – “Data Quality Reporting”;
- resources comply with the resource profile published at the given maturity level including compliance with coding systems;
- coded values exist in the coding system being used;
- the FHIR endpoint *CapabilityStatement* is consistent with the maturity level claimed;
- all features of the FHIR search API required for the maturity level;
- appropriate resource management actions (POST, PUT and PATCH) are supported as required by the maturity level;
- data can be queried asynchronously, and results retrieved (at appropriate maturity levels);
- subscriptions can be lodged, and subscription results are delivered.

Maturity testing will normally be concluded in the Sandpit environment, but it should be possible to repeat tests in the Staging environment.

Additional compliance tests run in Staging and Live which focus on security and governance. Tests include:

- a JWT is required to access data provider endpoints;
- the JWT must be of the correct format and must be signed;
- the scope rules which are derived from the JWT reason-for-access and user-type and define rights to access data types and to conduct non-patient centric searches are complied with;
- an expired JWT cannot be used to access data;
- an audit record is written for data accessed and this complies with the regional auditing standard (design paper 009 – “Auditing”).

The above tests are run in the live environment at a periodicity which is documented in the YHCR Operations Guide.

Appendix 1 – Maturity Matrix

Section	Narrative	Consultative	Draft	Normative
1 Introduction 1.1 Purpose of this Document	X			
1.2 Onboarding and Risk Management	X			
1.3 The Participant Registry	X			
1.4 Relationship of this Document with Other Standards	X			
1.5 Intended Users of the This Document	X			
2 Environments of the System of Systems 2.1 Types of Environment			X	
2.2 Dependencies on Data Providers			X	
2.3 A Data Provider's Pathway to Live 2.3.1 Sandpit			X	
2.3.2 Staging			X	
2.3.3 Live			X	
2.4 Autonomy of Environments			X	
3 The Onboarding Workflow 3.1 Anatomy of a Data Provider			X	
3.2 Onboarding Steps			X	
3.3 Requirements for the Onboarding Suite 3.3.1 Enforce Workflow		X		
3.3.2 Segregate Duties		X		
3.3.3 Access for Technical Officers		X		
3.3.4 Manage Component Configuration		X		
3.3.5 Log Access and Data Management		X		
3.3.6 Alerting and Data Reporting		X		
3.3 Raising Maturity Levels and Re-entering the Onboarding Process		X		
4. The Participant Registry for Data Providers 4.1 The *DataProvider* Resource			X	
4.2 APIs for Interacting with DataProviders			X	
4.3 Use of the DataProvider Resource by YHCR Components			X	

4.3.1. FHIR Aggregator				
4.3.2 Subscription Manager			X	
4.3.3 Reliable Massaging			X	
5. Automated Compliance Testing		X		