Cookbook for Regional Interoperability Detailed Design Paper #005

Identity and Access Management

Version 1.7 – 3rd May 2023

Abstract Interoperability Cookbook Anchor Points

Section	Title
3.1.1	Identity and Access Management Server (IAM)
5.2	Interactions with IAM

Table of Contents

Version Control

Version	Release Date	Released By	Reason for Release
1.0	17/02/2019	R Hickingbotham	Preliminary draft
1.1	12/11/2020	R Hickingbotham	Internal discussion draft
1.2	05/07/2021	T Davey	Add ability to include ASID in token
1.3	26/10/2021	T Davey	Add new Reasons for PDS Trace and Clinical Safety Testing
1.4	04/04/2023	T Davey	Clarify use of “LCL:” as a prefix
1.7	25/04/2023	T Davey/S Lowsley	Align Roles with National RBAC

Reviewers

Initials	Name	Role	Organisation
TD	Tim Davey	Technical Architect	PA Consulting
GH	George Hope	Technical Architect	Leeds City Council

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 005 - “Identity and Assessment Management” (IAM) describes the functionality of a regional service which is at the heart of the security model for the YHCR. This service authorises a user or system to interact with regional services and the FHIR endpoints offered by the data providers which contribute to the YHCR.

IAM also helps to enforce the security model by acting as a proxy to all regional services. In this role IAM consistently validates that authorisation has indeed been granted to service users and data requests are appropriate.

A third function of IAM is to harmonise user identities so that a user has a single regional persona regardless of the mechanism through which the user accesses the YHCR. For example, a care professional employed by a community care trust may work some of his week in a neighbourhood team which is located in as primary care clinic. This care professional may access the YHCR through both the trust’s own Electronic Patient Record system and the General Practise system. One objective of IAM is to reconcile these apparently different identities and to associate each with a single regional user identity.

Harmonising user identities facilitates auditing of a user’s YHCR usage and simplifies automatic detection of potentially inadmissible usage. It also would, in the future, allow legitimate relationships, roles, and privileges to be enforced at a regional level.

1.2 The Regional Security Model

It is a key tenant of the YHCR that responsibility for authentication rests with local care organisations. In line with this, IAM acts as an authorisation service rather than an authentication service. A data consumer authenticates a user using an appropriate local mechanism or the National Identity service. The data consumer is responsible for assigning the user a role and enforcing access controls and legitimate relationships with patients. The data consumer also determines whether the user has access to YHCR, and if so, under which YHCR role the user is performing in his/her duties.

The local data consumer registers the user with IAM in an authorisation request. The relationship between IAM and the data consumer is a trusted one: it requires pre-registration of the consumer with the YHCR and for YHCR to issue an X.509 certificate to the system’s operators. This certificate can be used to sign the data consumer’s authorisation request and enables IAM to trust the registration.

On receipt of a valid authorisation request, IAM issues the data consumer an access-token. This is structured block of data which contains some information about the user, and which has been signed using a private key held by IAM. The access-token has an expiry data with a short (15 minute) validity window. The access-token is included in all future interactions with regional services or direct interactions with data providers. The IAM signed access-token all data providers trust to establish a train of trust to the data consumer and is sufficient, subject to additional local controls, to authorise access to data available from the data provider.

1.3 Regional Roles, Privileges, and Legitimate Relationships

The following user roles have been defined for the YHCR:

- Citizen;
- System or Robot;
- Administrator (of YHCR systems).
- Auditor
- Authorised Carer
- National Role 0
- National Role 1
- National Role 2
- National Role 3
- National Role 3plus
- National Role 4

Identification of these roles does not preclude new roles being defined in the future. In fact, as use is made of the regional persistence capability and a new breed of applications manage processes that span care settings then more precise definition of user roles and the rights to data that a role implies will need be modelled centrally.

This design paper anticipates the future requirements for regional control of data access by role and regional enforcement of legitimate relationships with patients and clients. It proposes a data model for access control and species management APIs.

1.4 Relationship of this Document with Other Standards

The authorisation component of IAM is derived from the Oauth2 specification and is compatible with SMART on FHIR. Relevant standards include: OAuth2 (RFC 6749), Assertion Framework for OAuth 2.0 (RFC 7521), JSON Web Token (RFC 7523), and JSON Web Signature (RFC 7515).

1.5 Intended Users of the This Document

This document is targeted at the developers of the regional IAM service and developers of local data consumers who looking to authorise access to regional data.

2 User Authorisation and Access Token Management

IAM adopts OAuth2 for its authorisation function. OAuth2 actors and their alignment with components of the system of systems is illustrated below.

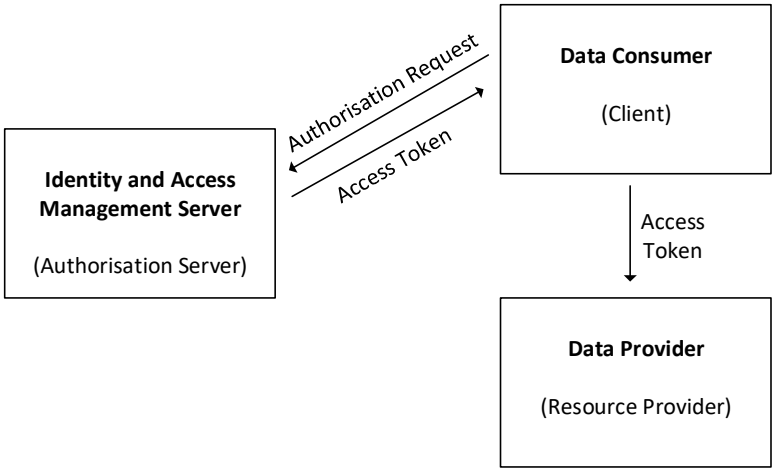

Not that the resource owner's authorisation is implied and that the authorisation server issues an access token without an authorisation grant from the resource owner.

2.1 Authorisation Service Request

The authorisation request follows the mechanism described in RFC 7521 for using assertions as an authorisation grant. In summary, the data consumer (client) generates an assertion based on a JSON Web Token (JWT) which is signed using a certificate issued by the YHCR certificating authority.

The client or data consumer identifies themselves using the Authorisation HTTP header using a basic scheme containing a base 64 encoded string of the format {ClientId:Secret}. The ClientId and Secret are allocated to the data consumer on registering with the YHCR.

The assertion is included as parameter in the grant_type header of an HTTPS POST request to the IAM authentication purpose.

```
POST /AuthService/oauth/token
Authorization: Yh4dZZxc987gGffd0078769Hgf3uHg2
Grant_type=urn:ietf:params:oauth:grant-type:jwt-
bearer&assertion={JWS}
```

The header is shown unencoded over multiple lines for clarity of reading.

A scope parameter is optional. If it is provided then it should follow the OAuth convention but this parameter will not be validated or enforced by the YHCR.

{JWS} is a compact serialisation of a JSON Web Signature which is defined by RFC 7515 as:

```
BASE64URL(UTF8(JWS Protected Header)) || '.' ||
BASE64URL(JWS Payload) || '.' ||
BASE64URL(JWS Signature)
```

The JWS Signature must be created using RSA and SHA256, ignoring unnecessary whitespace. The JWS Protected Header MUST be:

```
{"alg":"RS256"}
```

The JWS Payload MUST be a claim which conforms to the following structure:

```
{
    "jti":"36ee43c9f57e42bba265607508f0c8bc",
    "iss":"LCR",
    "aud":"IAM",
    "sub":523738395,
    "pat":{
        "nhs":1234567890,
        "fam":"Jones",
        "giv":"Jack",
        "dob":"19651206"
    }
    "ods":" 8JL372",
    "usr":{
         "fam":"Smith",
         "giv":"John",
         "rol":2,
         "ids": [{
             "sys":"ERS",
             "idc":"653990037"
         }],
         "org":"8JL372"
    },
    "rsn":1,
    "iat":50734946427,
    "exp":50734947327,
    "asid":"ABC123"
}
```

Where

Field	Explanation	Required
aud	The target audience of the JWT. This must be “IAM”	Yes
exp	The expiry date of the claim. This must use Co-ordinated Universal Time (UTC) and is expressed as the number of seconds since 1/1/1970.[3]	No
iat	The issue time of the claim. This must use Co-ordinated Universal Time (UTC) and is expressed as the number of seconds since 1/1/1970.[3]	No
iss	The application which issues the claim. A unique identifier which is assigned to a data consumer. This should be the same as the client id in the HTTP authorization header.	Yes
jti	The ID of the claim: a globally unique identifier which is used to prevent replay attacks.	Yes
ods	The ODS code of the organisation that issued the claim.	Yes
pat	The patient that is the subject of the access request.	(1)
pat.nhs	The patient’s NHS number	(1)
pat.fam	The patient’s family name (case insensitive).	(1)
pat.giv	The patient’s given name (case insensitive).	(1)
pat.dob	The patient’s date of birth (YYYYMMDD).	(1)

rsn	The reason for the access request (see below).	Yes
sub	The subject of the claim. This must be unique identifier for the end user which has been authenticated by the data consumer and from which audit details of the user session can be obtained from the organisation issuing the claim.	Yes
usr	Details of the end user	Yes
usr.fam	The user's family name (case insensitive).	(2)
usr.giv	The user's given name (case insensitive).	(2)
usr.rol	The user's role (see table below).	Yes
usr.ids	One or more identifiers for the user. (see the table in 2.1.3 for valid identifier systems).	(2)
usr.ids.sys	The coding system used for the identifier.	(2)
usr.ids.idc	The identifier allocated to the user by the coding system.	(2)
usr.org	The ODS code of the organisation who employs the user in the capacity for which they are accessing the YHCR. (Note that this is the ODS code which will be passed to National Systems such as GP Connect)	Yes
asid	The asid (Accredited System ID") for accessing national systems such as GP Connect. This is optional and is only necessary to populate if (i) access to national systems and specifically GP Connect is required AND (ii) it is necessary to override the default ASID configured on the Data Consumer. (This might be the case, for example, in multi-tenant scenarios)	No

Notes:

(1) The requirement depends on the reason for access. A patient is not required if the reason is not patient-centric.
(2) A user name or identifier is not required for access by systems or robots. The iss field uniquely identifies the system or robot.
(3) System times should computed from clocks synchronised with a reliable time source using Network Time Protocol (NTP).

2.1.1 Reasons for Accessing the YHCR

The following reason codes have been defined:

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

5	Administration (Not patient-centric).
6	PDS Trace
7.1	Clinical Safety Testing (data) (This allows access to pre-release datasets which have been flagged as available for clinical safety testing only)
7.2	Clinical Safety Testing (UI)

The coding system is extensible and future use cases can be accommodated by appending to the above codes i.e.: 1.1.1 may be used to represent access for emergency, direct care from a walk-in centre.

Note that for non-patent-centric access regional infrastructure will restrict access to non-patient identifiable data.

Reason code 3 will be used by data consumers creating subscription to data points and reason code 2 for citizens when accessing their own care record.

2.1.2 Role Codes

The following user role codes are defined. These are aligned, where relevant, with the National RBAC Roles (see Appendix A for details)

Code	Role
1	National Role 4
~~2~~	~~Social Care Professional~~ **(DEPRECATED – DO NOT USE)**
3	Citizen.
4	System or Robot.
5	Administrator
6	Auditor.
7	Authorised Carer.
8	National Role 1
9	National Role 2
10	National Role 3
11	National Role 3plus
12	National Role 0

Notes:

- The coding system is extensible and future use cases can be accommodated (only if necessary) by appending to the above codes i.e.: 1.1 may be used to represent a clinical professional who is a nurse.
- National Roles 0 to 3plus have yet to be implemented for consumer use. These are represented in the table above as part of the preparatory work required to scope their access requirements. Prototyping will be done in the Interweave Portal in the first instance, so that this can be closely controlled and then extended to other consumers based on this experience and agreement with the interweave partners. This document will be further updated to reflect any decisions made on these roles.

- The National Roles are intentionally given numbers with no attempt to assign "user-friendly" names based on, for example, a job description. This reflects experience of the national team that it is better to simply define the appropriate set of permissions – rather than attempting to link with the many and varied job titles and their interpretations.
    - For example a "GP Receptionist" in one organisation might be granted National Role 1, and in a different organisation might be granted National Role 2. This difference may be entirely appropriate, based on the actual nature of the job, the individual's training, etc. Decoupling access permissions from job titles was a key insight of the national work
- "National Role 0" is different from the Interweave "Administrator" role
    - "National Role 0" defines a very limited set of access – for technical troubleshooting only, with no access to clinical data
    - The Interweave "Administrator" is a super-user role, with full access to all capabilities
- "National Role 4" is a cosmetic rename of the role previously called "Clinical Professional"
    - This has no technical impact as it retains the same number code "1" in messages – the change is purely to the on-screen display text
- The role previously called "Social Care Professional" (number code "2") has been deprecated and must no longer be used
    - It is now covered by National Role 4, which grants the equivalent full access to data
    - This is based on the philosophy described above where we are agnostic to job titles. It does not matter whether a person works in Health or in Social Care. It may be appropriate for some individuals to have National Role 4 and other individuals to have lesser access. This should be decided based on the specific circumstances, rather than a simplistic division between "health" and "social care"
    - In practice users found this distinction unhelpful anyway, and the "Social Care Professional" role has been almost entirely unused

2.1.3 User Identification Systems

The following coding systems are supported:

Code	Description
ESR	Electronic Staff Record number.
ODS	Organisational Data Set code for practitioners.
SDS	Spine Directory Service identifier.
NHS	NHS number.
NI	National Insurance Number.
LCL:<ODS>	A local identifier allocated by the organisation identified by ODS code <ODS>

2.1.4 Validation of the JWT

The following validations are undertaken:

1) The HTTP Authorization header is present and the client Id and secret match that allocated to a data consumer.
2) The signature has been generated using the certificate allocated to the data consumer.

3) All required JWT properties are preent.
4) The JWT “iss” property is the same as the Client Id in the HTTP Authorization header.
5) The JWT “aud” property is “IAM”
6) The JWT “jti” property has not been previously used in a authorization request with IAM.
7) The JWT “ods” property is an ODS code that represents an organisation known to the YHCR.
8) If present the patient nhs number, family name, given name and date of birth represents a patient known to the YHCR.
9) The JWT “rsn” code is valid and is consistent with the “usr.rol” code.
10) For citizen access user.ids is an NHS number and this is the same at pat.nhs.

2.1.5 Authorisation Service Endpoint

The endpoint address of the Authorisation Service is published in the YHCR Operations Guide.

2.2 Authorisation Service Response

If the authorization request is accepted, then the body of the HTTP response contains the access token as follows:

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache

{
  "access_token":"2YotnFZFEjr1zCsicMWpAA.54Ry006FfggHkd54.7643jGh4kkJd",
  "token_type":"bearer",
  "expires_in":900
}
```

The access token is a compact serialisation of a JSON Web Signature signed by IAM. The JWS payload is a precise copy the JWT provided by the client in the authorisation request in all but the following respects:

- exp is an expiry date determined by IAM calculated according to the policy published in the YHCR Operations guide;
- iat is the date time that IAM processed the request;
- jti is a globally unique identifier created by IAM.

If the request is not accepted, then an error response is returned in the following format:

```
HTTP/1.1 400 Bad Request
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache

{
   "error":"invalid_request",
   "error_description":"Unsupported user identification coding system”
}
```

All invocation of the Authorisation Request service are logged. *AuditEvent* FHIR resources are created in the regional FHIR store. The content of the *AuditEvent* is detailed in design paper 009 – “Auditing”.

2.3 Validate Token Service

A data provider or resource owner can validate an access token issued using the following service:

```
POST /Validate/oauth/token
Authorization: Yh4dZZxc987gGffd0078769Hgf3uHg2
Content-Type: application/json;charset=UTF-8

{
  "access_token":"2YotnFZFEjr1zCsicMWpAA.54Ry006FfggHkd54.7643jGh4kkJd"
}
```

The authorisation HTTP header uses a basic scheme containing a base 64 encoded string of the format {ClientId:Secret}. The ClientId and Secret are allocated to the data provider on registering with the YHCR.

An example of successful response is:

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache

{
  "token_valid":1
}
```

Where 0 indicates an invalid token and 1 a valid token.

A data provider is not normally expected to validate tokens with IAM. In normal circumstances it is sufficient to establish token validity locally by verifying the signature and expiry date of the token. Central validation is only necessary when there is a measured risk of tokens being hijacked by third parties and where a mechanism has been established for revoking tokens where misuse is suspected.

2.4 Revoke Token Service

A data provider or data consumer can revoke an access token using the following service:

```
POST /Revoke/oauth/token
Authorization: Yh4dZZxc987gGffd0078769Hgf3uHg2
Content-Type: application/json;charset=UTF-8

{
  "access_token":"2YotnFZFEjr1zCsicMWpAA.54Ry006FfggHkd54.7643jGh4kkJd"
}
```

An HTTP 200 response indicated that the token has been invalidated.

2.5 Refresh Token Service

IAM does not provide a token refresh service. A data consumer holding a token that is nearing its expiry date should obtain a new token using the authorization service.

2.6 Change Context Service

IAM does not provide a change context service. A data consumer which is moving context to a new patient should obtain a new token using the authorization service.

3 Regional Proxy Services

IAM acts as a proxy to regional services and so centralises enforcement of access control. All IAM implemented services require a valid bearer-token to be presented in the HTTP Authentication header or in another protocol specific manner. IAM validates the bear-token signature, its expiry date, and that it has not been revoked.

The reason for accessing YHCR as defined by 2.1.1 implies the scope of access allowed. The scope is enforced in different ways by different regional services but presented here, logically, as an IAM function.

There is an overlap in the concept of implied scope and the scope definition which might optionally be included in the grant type for the Authorisation Request. The grant type scope is potentially a useful device for refining the scope defined by the “reason for access”. I.e.: the grant type scope should be a subset of the scope permitted by the “reason for access”, but this is not validated by IAM and the grant type scope is not used to further constrain rights to data.

3.1 Patient Centric Access

Many of the reasons for access are patient centric and patient related resources should only be accessed for patient which is currently in context. Specifically:

- resource read and version reads must be for a resource that references the patient as its subject;
- resource searches must reference the patient as one of its search terms;
- resources can only be created or modified that reference the patient as their subject.

Of the resources identified by the cookbook the following are patient related:

Appointment	A booking of a healthcare event involving a patient, practitioner, location or device.
AppointmentResponse	Confirmation or rejection of an attempt to book an appointment
AuditEvent	A record of a security audit event.
BodySite	Used by ProcedureRequests and Observations to define an anatomical location for a particular patient.
CarePlan	Describes an intention of how care will be delivered to address a particular condition for a patient or group of patients.
ClinicalImpression	An assessment aimed at determining the problems affecting a patient.
Condition	A problem, diagnosis or other issue pertaining to a patient or group of patents.
Consent	A statement of a patient’s acquiescence to a consent policy.
DiagnosticReport	The findings and interpretations of diagnostic tests applied to a subject (usually but not always a patient).
Encounter	An encounter with a patient or group of patients.
EpisodeOfCare	A period of care during which an organisation has a responsibility to a patient.
FamilyMemberHistory	Health events pertaining to a person related to a patient.
Group	A group of patients.
Immunization	The record of a vaccination being given to a patient.
MedicationRequest	The prescription of a medication.

MedicationStatement	A report by a patient or a care professional of a past medication administration.
Patient	A patient.
Person	A person.
Procedure	A medical procedure performed on a patient.
ProcedureRequest	A record of a request for diagnostic investigations, treatments, or operations to be performed.
Questionnaire	A set of questions.
QuestionnaireResponse	Responses to questions by an individual.
ReferralRequest	A request to refer a patient to a healthcare service.
RelatedPerson	A link to a person who is related to another.
RiskAssessment	An assessment of the likely outcome(s) for a patient or other subject as well as the likelihood of each outcome.

The following resources are not patient related and may be freely retrieved or managed (subject to data source policies) under the authorization grant.

CareTeam	An assembly or practitioners as a team.
Goal	An objective in a care plan.
HealthcareService	A service available at a location.
Location	A physical location.
Medication	The definition of a medication including details of packaging and batch identification.
Organization	An organisation.
Practitioner	A practitioner.
PractitionerRole	The role a practitioner undertakes in an organisation.
Schedule	Part of the mechanism for booking appointments for a clinic/practitioner.
Slot	A time period against which an appointment can be booked.
Substance	A homogeneous material with a definite composition.

The following resources may be patient related, and resource specific logic is needed to determine whether access is possible.

Communication	Some form of communication sent from one party to another.
CommunicationRequest	A request to receive a communication (less formal than a subscription).
Composition	A structural resource used to embed the content of an immutable document.
Flag	Things to be aware of for a patient, medication, location etc.
List	A structural resource representing a list of other resources.
Subscription	A structural resource representing an expression of interest in a data point.
Task	Tracks the request and execution of a task issued to an organisation or individual.

3.2 Access for the Purpose of Direct Care

By default, consent rules rule (as defined by design paper 008 “Consent Management”) do not apply. The exception is for data providers which manage patient supplied data. These providers are identified as such in the registration process (design paper 020 – “Onboarding Data Providers”) and any data released to a data consumer from these providers is tested against policy rules.

3.3 Access for the Purpose of Indirect Care

Rules for patient centric access to not apply. However, data will only be released if an explicit regional consent rule (as defined by design paper 008 “Consent Management”) permits it.

3.4 Administrative Access

Access is not permitted to any data which outside of the regional infrastructure. Patient identifiable data may be accessed if it is part of the Master Patient Index (design paper 004 – “Patient Identity Exchange (PIX/MPI)”) or held in the regional FHIR store (design paper 018 – “Regional FHIR Store”).

3.5 Citizen Access

Citizens will not have access to the same data as care professionals. There may be data that the citizen has contributed to the YHCR which is withheld from general access by care professionals: a constraint which enforced by a consent rule.

Citizens or patients will also not have the same view of clinical data as their carers: data may be withheld by care professionals because it is sensitive or requires interpretation for the citizen. Data may also be presented in a different manner to facilitate understanding by the patient.

These rules cannot be enforced regionally. Rather, it is the data provider which is responsible for formatting data for citizen access and releasing only the data that is appropriate for citizen consumption. The citizen role is a data item passed to data providers in the signed JWT and this can be seen as an instruction to alter the presentation of data.

Different roles have been allocated to citizens as patients (role 3) and citizens acting as carers (role 7). This distinction allows for further precision in the definition of consent policies.

3.6 Auditor Access

Auditors have access solely to AuditEvent resources (design paper 009 – "Auditing"). No other role has access to resources of this type. The audit record is distributed across data consumers, the system of systems and data providers. Access to the consumer held audit record is under the control and at the discretion of the data consumer (although good information governance practices such as separation of roles and technical solutions to ensure the integrity of the audit record are required and are a condition of the onboarding process). Access to AuditEvent resources in the centre and at data providers requires an IAM claim to be made using the Audit role. A token issued in response to this claim can be used to access AuditEvents in the System of Systems or any data provider through the standard operation of the FHIR Aggregator (design paper 010).

4 Identity Management

IAM maintains an identity model to enable local identities to be consolidated into a regional identity and so enabling a full audit of all YHCR usage to be assembled for individuals regardless of how they access the YHCR. A consolidated regional identity also allows role-based access rights to be controlled for regionally held data. IAM manages the following identity model:

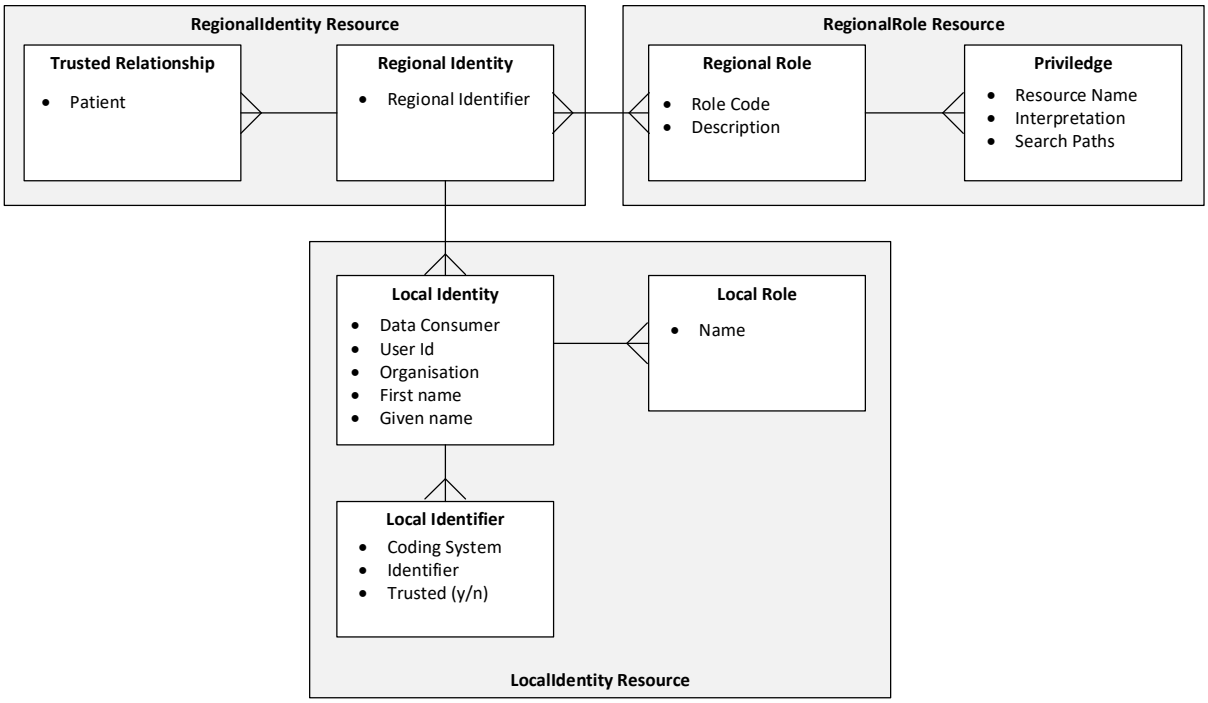

The model depicts the relationships between key data entities and the packaging of entities within a resource model.

Regional identities are derived as a function of the Authorization Service. When a data consumer invokes the service IAM attempts to link the local identity of the user with a regional identity. The principle being that a regional identity allows the same individual accessing the YHCR from different data consumers to have the same regional persona with the same privileges over regional data.

IAM uses local identifiers (ERS number, NHS number NI number etc.) to match local identities. Local identities with same local identifiers are joined to a regional identity.

Identifiers can conflict and where a conflict is found then the status of a local identifier is downgraded, and it becomes non-trusted. Trust is managed so that a trusted local identifier is associated with exactly one regional identity. There may be other instances of the same local identifier associated with different regional identities but these have a non-trusted status.

Local Identities are uniquely indexed by the combination of a data consumer and the user id which was allocated by the consumer (the iss and sub fields in authorisation request claim). In other words, IAM maintains a local identity for every user of every data consumer.

IAM maintains identity demographics (organisation, first name and given name) and these are populated from the last demographic presented by a consumer. Local roles are an amalgamation of roles presented by the data consumer over different authorisation requests.

4.1 Legitimate Relationships and Privileges

At this stage the primary use of legitimate relationships and privileges is seen to be to control access to regionally held data. Whilst use cases are still to be defined, it is likely that the YHCR will be a catalyst for a new breed of patient centric applications which co-ordinate interactions with patients across care settings and independently from any traditional source of data. IAM allows a common set of privileges and legitimate relationships to managed for all applications using the regional FHIR Store (design paper 018 – "Regional FHIR Store") and these are enforced regionally.

Legitimate relationships are defined as a set of references to FHIR Patient resources and the implication is that the user should only have access to data for these patients. If the set is not defined, then it is implied that the user has access to all patients' data (subject to privileges).

Privileges are defined by role. A privilege is defined for a resource type and is tied to a management action (reading, creating, or updating) for the resource type. A set of FHIR search paths provide the details of which resources are covered by the privilege.

4.2 Regional Identity Management

Regional Identities are created during an authorisation request made on IAM.

4.2.1 A New Identity Presents

A new regional identity may be created, or the new local identity may be linked to existing regional identity. Linking a new local identity to an existing regional identity implies that it inherits existing legitimate relationships and privileges. Local identifiers for the presenting identity are used to determine the linkage action.

Local Identifier Test	Action
None of the local identifiers match any existing trusted local identifiers	Create a new regional identity and trust all local identifiers
Some or all of the local identifiers match trusted local identifiers which are associated with a single regional resource.	Link the local identity with the existing regional identity and trust all local identifiers.
Two or more local identifiers match trusted local identifiers associated with two or more regional identities.	Create a new regional identity and trust only those local identifiers which are not trusted for other regional identities.

4.2.2 A Known Identity Presents with New Local Identifiers

The local identity already has a regional identity. It is possible that new identifiers are present which conflict with those trusted for other regional identities. If this is the case, and the regional identity is shared by other local identities, then local identity must be detached (see section 4.2) from its regional identity and a new regional identity created. Local identifiers (old and new) are only trusted for where there is no conflict with other trusted local identifiers.

4.3 Detaching a Local Identity from a Regional Identity

Where a regional identity is shared by one or more local identities then one of the local identities may be detached. Detaching a local identity results in a new regional identity being created with the

same legitimate relationships, roles, and privileges as the previous regional identity. The detached local identity is reattached to the new regional identity and identifiers which conflict with those trusted at the prior regional identity are now untrusted.

Note that identity chains might exist at the original regional identity. Each linked local identity must be tested and possibly re-linked to a newly created regional identity or another regional identity created specifically for it.

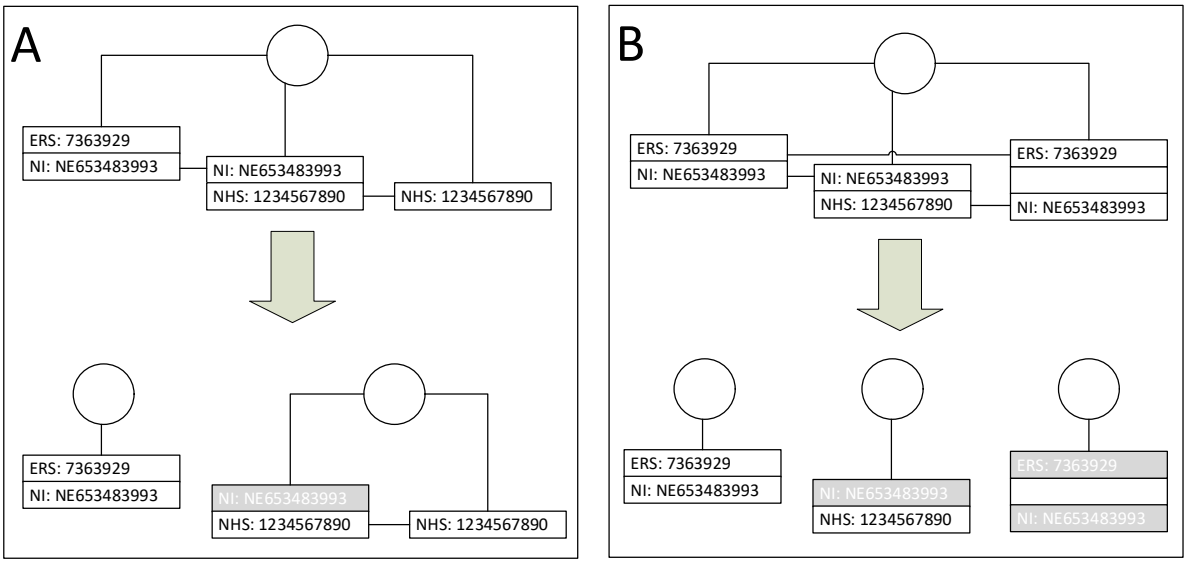

4.4 Merging Regional Identities

The effect of merging two regional identities is to create a new regional identity which is associated with a superset of legitimate relationships and regional roles drawn from the merged identities.

Note that the new regional identity is also associated with the superset of trusted local identifiers which were associated with the merging regional identifiers.

4.5 Audit Records and Identity Management

Regional services log all access to the YHCR and regional FHIR resources obtained through the regional FHIR Bus as *AuditEvents* in the regional FHIR Store (Design paper 009 – "Auditing").

AuditEvents reference IAM *LocalIdentity* resources. Whilst identity management activities performed by IAM will affect the linkage between regional identities and local identities, local identities are never changed and relinkage do not impact the integrity of audit records.

IAM maintains a history of all relinkages and the historic privileges and legitimate relationships that existed for a local identity through its association with a regional identity can be re-established at any time.

4.6 Security Implications of Automated Identity Management

The processes describe here result in hitherto unknown user identities being afforded privileges over patient data automatically through their association with other identities assumed by them at other organisations. It should be noted that privilege in isolation does not entitle the user to access the data.

4.7 Identity Management APIs and Management Console

REST APIs allow management of all IAM managed resources. In particular resources can be retrieved, queried, created, deleted (where data model integrity is preserved), and patched using the standard HTTP verbs and search syntax.

APIs preserve the reliability of the association between regional identities and trust local identifiers.

APIS can only be accessed by a client in possession of a IAM issued bearer token with a reason for access “Administration” and role “Administrator”.

A management console is offered which performs the following functions:

- role and privilege management;
- legitimate relationship management;
- deletion of local identities;
- detaching of local identities from regional identities;
- merging of regional identities;
- switching of the trust status on local identifiers (subject to integrity constraints).

5 Appendix A – National RBAC Roles

This appendix summarises the National RBAC Roles, as defined by NHS England in the document "***Shared Care Records, January Version 2.1, Access Controls Model (for direct care)***"

Role 0	grants access to systems but no access to clinical information
Role 1	grants access to basic demographic data
Role 2	grants access to demographic and administrative data including the ability to view and book appointments
Role 3	grants access to summary information only in the shared care record**.** User Interface Design will determine how this information is presented.
Role 3+	grants access to summary information with the option of accessing extended information, subject to additional controls. User Interface Design will determine how this information is presented.
Role 4	grants access to extended information. User Interface Design will determine how the information is presented. If individuals/people who draw on social care object to information being shared for direct care purposes, this should be discussed, and a decision made in line with clinical safety. Please see Appendix D for further information.