Medical Data Gateway Specification

An Imaging Exchange Framework for Australia

Version 1.0

21 July 2025

Christopher Skene
christopher.skene@aurabox.cloud

Table of Contents

Introduction

This document is the functional specification for a Medical Data Gateway, a component of the Imaging Exchange Framework, a proposal for managing the seamless exchange of medical imaging in Australia.

This document should be read in conjunction with An Imaging Exchange Framework for Australia, available at https://aurabox.cloud.

About the Authors

The IEF was created by Aurabox. Aurabox is an Australian company working at the forefront of medical imaging interoperability. Founded in 2021 by radiologist Dr Chaturica Athukorala and technologist Christopher Skene, Aurabox's mission is to solve the problem of access to medical imaging. Aurabox operates a SaaS platform for medical imaging interoperability in Australia, the UK, and Singapore.

Licence & Reuse

The Profile

- Operational layer.
- Recommended (but not mandatory) for providers wishing to enable discovery and sharing without deeply modifying their PACS/RIS.
- Facilitates compliance and adoption of IEF by providing a drop-in component.

The **Medical Data Gateway (MDG)** is the reference implementation profile for the Imaging Exchange Framework (IEF). It provides a formal, standards-aligned technical architecture for integrating imaging systems into Australia's federated digital health environment, consistent with Health Connect Australia principles.

The MDG serves as an interoperability broker and control point, bridging local imaging environments (e.g., PACS, RIS, VNA) with requesting systems, FHIR-based services, and Health Connect Australia's national infrastructure. It implements the core capabilities of the IEF, providing secure, policy-compliant pathways for the discovery, access, and exchange of medical imaging data.

The MDG provides a scalable, consistent pathway for imaging providers to participate in Australia's connected health ecosystem without requiring wholesale system replacement.

> Although it is not defined in the IEF or MDG specifications, an MDG could function as an organisation's entire FHIR gateway, or be provided by an existing FHIR gateway, as long as it is capable of meeting the requirements of the specification. T*his opens the possibility that an expanded MDG can act as the sole FHIR service connector for an organisation* – especially where the organisation is primarily concerned with medical imaging – potentially simplifying any implementation.

The MDG enables:

- Secure **Registration** of newly acquired imaging studies with national or jurisdictional record locator services.
- Standards-based **Discovery** of imaging via a FHIR interface
- Bulk **Retrieval** of imaging via JMIX.
- **Authorising and brokering access** consistent with Health Connect Australia's Common Services Layer.
- **Integrating with local imaging systems** (e.g. PACS, RIS, VNA) to retrieve studies.
- Optionally managing **event logs, consent assertions, and usage reporting**.

The MDG can be deployed in several ways:

- **Directly by imaging providers**, as an on-premises or cloud-hosted gateway;
- **Via a reference implementation** such as *Harmony*, developed by Aurabox;

- **Bundled within vendor infrastructure**, such as in PACS or RIS software offerings.

*A functional specification for the MDG is provided in **Appendix 2**.*

Capabilities

The MDG defines a set of technical capabilities that enable secure, scalable, and standards-based discovery, access, and transfer of medical imaging. These capabilities are designed to align with the architectural principles, interoperability patterns, and policy requirements of the IEF. Together, these capabilities deliver a modern, interoperable, and policy-aligned technical foundation for imaging exchange.

1. Imaging discovery using FHIR
2. Bulk imaging transfer using JMIX
3. Registration of patient records via FHIR
4. FHIR subscription & notification services

1. Imaging Discovery using FHIR

Imaging systems expose a FHIR-compliant API endpoint that enables authorised parties to discover available imaging studies for a patient. This supports standardised, privacy-aware discovery workflows using national identifiers such as the Individual Healthcare Identifier (IHI) and provides sufficient metadata to inform clinical decision-making and relevance assessment before retrieval.

a. **FHIR Services** will provide a compatible FHIR API endpoint that enables authorised parties to discover available imaging studies for a patient.

   In a practical sense, the FHIR Provider API will be expected to:
   a. Receive a compatible FHIR request from a third-party
   b. Confirm that it contains a valid HealthConnect authorisation token for the data it is requesting.
   c. Return a FHIR response listing the available ImagingStudy data containing the relevant Bulk Data URLs.

b. **FHIR Consumers** will consume the FHIR endpoints, providing imaging history back to relevant clinical systems, and providing Bulk Imaging Endpoints for the transfer service (below). FHIR Imaging Consumers are likely to be the same systems that consume other FHIR APIs in the Health Connect ecosystem.

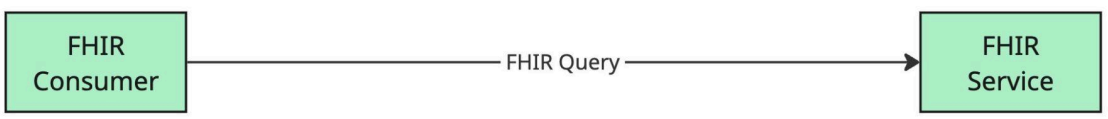

ImagingStudy Discovery via FHIR

2. Bulk Imaging Transfer using JMIX

To support high-throughput, system-to-system use cases, the MDG supports a *bulk transfer mechanism* enabling complete imaging studies to be securely packaged and delivered as compressed archives. This model simplifies both individual study transfers and the movement of large datasets for use cases such as longitudinal review, MDT preparation, research cohort assembly, or AI model training.

JMIX is the default and recommended transfer format within the Imaging Exchange Framework (IEF). It supports encrypted, policy-compliant distribution of imaging data across federated environments.

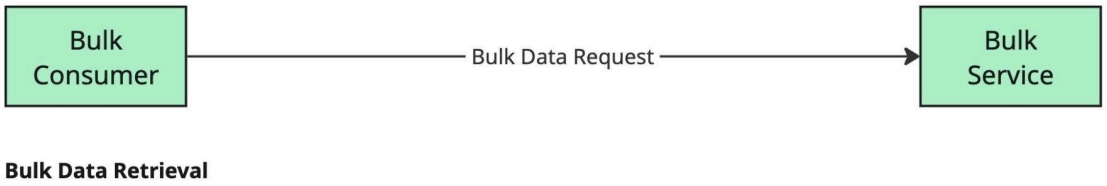

Bulk Data Retrieval

Dual-Mode Transfer Model

MDG Bulk Imaging Transfer implementations operate in a dual-mode structure:

- **Primary Endpoint (JMIX)**
  Bulk Data Services SHALL expose a JMIX-compatible endpoint that provides secure, verifiable access to full imaging studies via download URLs.

  The framework specifies either the new open specification, JMIX bulk imaging data model, which includes both secure transport format and simple retrieval API.

- **Compatibility Endpoint (DICOMweb, optional)**
  MDGs MAY expose a secondary DICOMweb-based endpoint as a *compatibility layer*. This allows for minimal integration with legacy viewers and clients that do not yet support JMIX. However, these endpoints MUST NOT be advertised as the default or preferred retrieval method and MUST be subject to the same authorisation and policy controls as JMIX endpoints.

Operational Flow – Bulk Data Service

A Bulk Data Service will:

- Receive a compatible FHIR or HTTP(S) request for an imaging dataset.
- Validate the presence and scope of a HealthConnect-compatible authorisation token.
- Return the dataset via the primary (JMIX) endpoint.
- Optionally expose a DICOMweb-compatible URL in the `ImagingStudy.endpoint` array **only when configured for compatibility**.

Operational Flow – Bulk Data Consumer

Bulk Data Consumers are systems that retrieve imaging from a Bulk Data Endpoint and forward it internally. These systems MUST:

- Extract the JMIX (or fallback) URL from the `ImagingStudy` resource.
- Revalidate or reuse a HealthConnect authorisation token.
- Retrieve the bulk data in JMIX format, unencrypt and validate it.
- Forward the dataset to relevant internal systems such as PACS, VNA, or analytical environments.

Where explicitly supported, Consumers MAY optionally use a DICOMweb endpoint, but SHOULD only do so where JMIX is not supported..

What is JMIX?

The **JSON Medical Imaging Exchange (JMIX)** format defines a secure, open, and scalable envelope for packaging and transferring complete imaging studies. It was developed to overcome limitations in legacy DICOM and DICOMweb-based transfers by supporting:

- Secure, encrypted transport
- Passive integrity verification
- Support for both clinical and secondary (e.g., research or teaching) use
- Federation across trust boundaries

JMIX is the default and recommended format within the MDG and IEF framework. Reference specifications and implementations are available at:

👉 https://github.com/aurabx/jmix

3. Registration via FHIR

In order to locate patient imaging within a federated system, it is assumed that HealthConnect will support some kind of patient registry, allowing Imaging Consumers to determine which Services hold patient data[1]. Services must register new patient imaging with this registry(s).

The capability only needs to support a notification when data is first added, however there may be other requirements based on a final HealthConnect implementation.

4. Subscription & Notification

Finally, a Service should be able to handle FHIR Subscription and Notifications for a given patient, identified by IHI. This assists with a range of lifecycle issues defined in the Use Cases section for this capability.

[1] In the event that this is not the case, the Registration model enables falling back on federated indexes.

A Reference Architecture: The Harmony Project

Harmony is an open-source reference implementation of the **Medical Data Gateway (MDG)** concept defined by the Imaging Exchange Framework (IEF). Developed by Aurabox, Harmony provides a practical, standards-aligned, and vendor-neutral pathway for healthcare organisations to participate in federated imaging exchange, consistent with the requirements of HealthConnect Australia.

Harmony acts as an interoperability broker between local imaging environments (e.g. PACS, RIS, VNA) and national infrastructure such as HealthConnect's discovery, identity, and consent services. It enables organisations to meet their IEF obligations with minimal disruption to existing workflows or technology investments.

Harmony is currently under active development with an initial release in late 2025.

Key Features

Harmony will provide the following core capabilities out of the box:

- **FHIR-Based Imaging Discovery:** Implements the IEF's ImagingStudy discovery endpoint, allowing authorised parties to locate imaging studies via standards-compliant FHIR APIs.
- **Bulk Imaging Transfer via JMIX:** Supports secure, efficient retrieval of complete imaging studies using the open JMIX standard, simplifying large dataset movement for clinical care, research, and secondary use.
- **Registration** of new studies with a Registry service based on HL7v2 messaging or internal webhook
- **FHIR pass-through** for other FHIR use cases.
- **DICOM Integration:** Connects to backend DICOM systems behind the firewall, enabling them to work with HealthConnect
- **Secure Identity, Consent, and Access Control:** Integrates with HealthConnect Australia's identity, consent, and authorisation services, with robust enforcement of privacy, security, and audit requirements.
- **Vendor Neutrality and Extensibility:** Designed to interface with any standards-compliant PACS, VNA, or imaging source, Harmony can be deployed as a stand-alone MDG or integrated with third-party solutions.
- **Internal job queue** supporting structured retry, logging, and other internal event management.

Federated Capabilities through Networks

While Harmony can be deployed independently by any organisation, when paired with a network like **Aurabox**, it extends its functionality to support:

- **Network-wide Imaging Discovery**, including from remote PACS systems.
- **Peer-to-Peer Imaging Transfers** powered by encrypted WireGuard connections, enabling secure, direct exchange between trusted parties without centralised infrastructure.

- **Declarative Usage Control**, Consent Management, and Identity Assurance aligned with the Australian Privacy Principles and HealthConnect policies.

This approach provides organisations with flexibility to adopt Harmony on their own terms, while offering a seamless pathway to participate in broader imaging networks such as the Aurabox platform.

Open Source

Harmony will be released under an open-source licence, with source code, documentation, and deployment resources available via a github repository.

Organisations are encouraged to adopt or extend Harmony as a practical, standards-aligned solution for imaging exchange, whether as a fully independent deployment or in conjunction with trusted service providers.

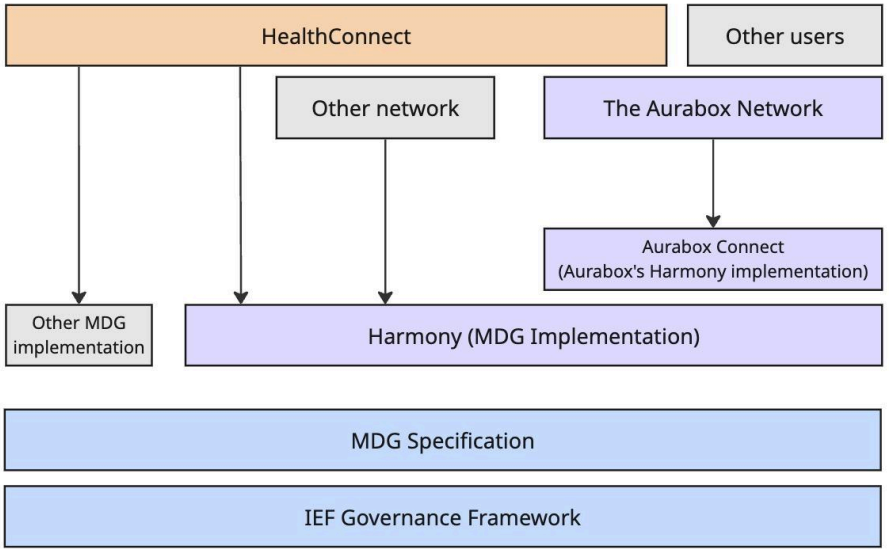

Relationship between the IEF, MDG, Harmony, Aurabox and Health Connect

Requirements

1) Automated, Standards-Based Resource Access

a) The system **MUST** support fully automated, pull-based access to health data resources, including imaging studies, reports, and associated metadata.

b) Access **MUST** comply with nationally recognised standards such as **FHIR**, and other protocols endorsed under the **AU Core Framework for Interoperability (AU CFI)** and HealthConnect

c) Exposure of resources **MUST** follow standardised **Health Connect** discovery and retrieval models including:

i) Discovered Information Exchange
ii) Directed Information Exchange

2) Standards-Compliant Authentication and Authorisation

a) **MUST** authenticate via a Health Connect Australia-approved identity and authorisation service (e.g. myID, HI Service), when such services are available.

i) Where feasible, this **MUST** include direct integration with the Digital Health Authentication and Participation Services provided by Health Connect Australia (e.g. myID, Healthcare Identifiers Service).

ii) Where direct integration is not technically or operationally feasible (e.g., due to local system limitations), the system **MUST** utilise an intermediate, Health Connect Australia-compliant service (e.g., a Medical Data Gateway) to broker authentication and authorisation.

iii) All authentication processes **MUST** comply with relevant national requirements for secure identity validation, including alignment to the Australian Privacy Principles, My Health Records Act, and Australian Signals Directorate Information Security Manual (ISM).

b) **MUST** expose a standards-compliant FHIR API for discovery of available imaging studies, supporting:

i) **MUST** support FHIR query parameters as defined by the appropriate AU ImagingStudy profile, including support for date range, modality, and body site when available.

ii) **MUST** provide a Capability endpoint that advertises the services FHIR capabilities

3) Automated, Policy-Based Access Control

a) The system **MUST** apply automated, policy-based access controls to all data access requests. Specifically:

   i) The identity and authorisation status of the requester.
   ii) The existence and validity of consent, preferences, or controls applicable to the resource in question.
   iii) The legal and policy basis for the access request, consistent with relevant legislative frameworks (e.g., My Health Records Act, jurisdictional privacy laws).

b) Where an access token specifies access to a specific IHI, the system **MUST NOT** allow access to other data.

c) Manual human intervention **MUST NOT** be required for routine access requests that meet these conditions.

4) Operational Automation

a) **MUST** enable full automation (i.e. no human intervention) for discovery, access, and delivery of imaging and health data within supported workflows.

b) Manual operator involvement **MUST NOT** be required for standards-compliant transactions

5) Edge Case Management and Manual Fallback

a) The system **SHOULD** provide defined, auditable pathways for manual intervention to support exceptional or edge-case scenarios where automated access cannot be granted.

   Such scenarios **MAY** include:

   i) Incomplete or ambiguous identity validation.
   ii) Absence of verifiable consent or consumer preferences.
   iii) Technical failure of automated pathways (e.g., Health Connect service unavailability).

b) Manual fallback processes **MUST** adhere to security, privacy, and audit requirements and **MUST NOT** circumvent standard access controls.

6) FHIR-Based Imaging Study Discovery

For Services:

a) **MUST** support GET requests for the FHIR ImagingStudy resource, *using the Individual Healthcare Identifier (IHI) as the primary search parameter,* consistent with Australian clinical identifiers policy.
b) **MUST** conform to the latest applicable Australian FHIR Implementation Guide, including but not limited to specifications developed through the Sparked FHIR Accelerator or published as part of the AU Core Framework for Interoperability (AU CFI).
c) **MUST** return only metadata appropriate for discovery, excluding diagnostic images or bulk data transfer, to facilitate determination of clinical relevance without unnecessary data exchange[2].
d) **SHOULD** include key clinical metadata fields within the ImagingStudy resource, subject to availability in the source system.
e) **MAY** support consumption of upstream FHIR services to enrich or manage imaging information (e.g., demographic lookups).
f) **MUST** ensure that all responses comply with applicable privacy, consent, and authorisation requirements, including integration with Health Connect Australia's Common Services Layer for identity validation and access control, where available.
g) **SHOULD** support query parameters consistent with international FHIR standards and Australian guidance to enable filtered discovery[3], such as:
    i) Date range (date parameter).
    ii) Modality (modality parameter).
    iii) Body site or anatomic location (body-site parameter), where supported.
h) **MAY** support additional query parameters to limit available study information within the existing IHI query.

For Imaging Consumers:

i) **MUST** be able to construct and make FHIR queries compatible with HealthConnect. These queries:
    i) **MUST** support HealthConnect authentication and authorisation
    ii) **MUST** support GET requests for the FHIR ImagingStudy resource, using the Individual Healthcare Identifier (IHI) as the primary search parameter, consistent with Australian clinical identifiers policy.
    iii) **MAY** support additional parameters to reduce scope, but **MUST** fall back on supported queries on systems where these parameters are not supported
    iv) **MUST** conform to the latest applicable Australian FHIR Implementation Guide, including but not limited to specifications developed through the Sparked FHIR Accelerator or published as part of the AU Core Framework for Interoperability (AU CFI).
    v) **MUST** be capable of providing Bulk Data URIs from ImagingStudy resources, to a compatible Bulk Data Consumer, where the system

[2] Subject to clinical input.
[3] Subject to clinical input.

performs both functions.

7) Bulk Data API for Imaging Retrieval

For Services:

a) To support efficient, high-throughput access to imaging data, participating systems MUST implement a mechanism for the bulk retrieval of imaging studies.

   i) **MUST** provide bulk access to complete imaging studies via a download-capable endpoint in either JMIX (preferred) or DICOMweb
   ii) **SHOULD** support the **JMIX format** for bulk imaging exchange. JMIX is the preferred mechanism and is intended to become the default over time.
   iii) **MAY** expose **DICOMweb-compatible endpoints** (e.g., WADO-RS) as a compatibility layer for clients that explicitly require it. These endpoints:
      (1) **MUST** comply with the same authorisation and access controls used for JMIX.
   iv) **MUST** use secure, encrypted **HTTP(S) with TLS 1.2 or higher**.
   v) **MUST** support **authenticated, time-limited download URLs**.
   vi) **SHOULD** support **resumable transfers** using HTTP Range requests or equivalent mechanisms to support robust transfer in variable network conditions.
   vii) **MAY** act as a proxy between the bulk data endpoint and internal systems (e.g., PACS, VNA, repository), provided that the bulk download response is fulfilled accordingly.

For Imaging Consumers:

b) Bulk Data Consumers are responsible for retrieving complete imaging datasets from third-party endpoints and delivering them to downstream systems. These systems may include PACS, VNA, clinical viewers, or analytical platforms.
   i) **SHOULD** support the retrieval and processing of **JMIX-formatted archives**, including decryption and integrity validation in accordance with the JMIX specification.
   ii) **MAY** support **DICOMweb retrieval** (e.g., WADO-RS) for compatibility with endpoints that do not yet provide JMIX. When doing so, Consumers:
      (1) **MUST** comply with the access controls and authorisation flows defined by the Service.
      (2) **SHOULD** optimise for high request volumes using batching, prefetching, or parallelisation strategies.
   iii) **MUST** use secure, encrypted **HTTP(S) with TLS 1.2 or higher** for all data transfers.

iv) **SHOULD** support **resumable downloads** for large study transfers to enhance reliability.

8) New Study Registration

a) **MUST** register newly acquired imaging studies with an authorised HealthConnect FHIR-based record locator or discovery service *or intermediate broker.*

b) **MUST** generate or update corresponding ImagingStudy resources with relevant metadata and access endpoints.

c) **MUST** support registration workflows triggered automatically from PACS, RIS, or VNA systems

9) Vendor Neutrality and Extensibility

a) **MUST** interoperate with standards-compliant PACS, VNA, or imaging systems
b) **SHOULD** support:
   i) Modular deployments (e.g. cloud, on-prem, hybrid)
   ii) Integration with third-party services like Aurabox
   iii) Extension to non-imaging health data exchange when appropriate

Imaging Consumer Interoperability:

c) **MUST** support Health Connect-compatible query and authorisation
d) **SHOULD** process returned resources using AU Core and JMIX standard

10) Identifier handling

Patient Identifiers (DICOM):

a) Implementers **MUST** store the IHI in the `Other Patient IDs Sequence (0010,1002)` as a distinct identifier entry (See Appendix 6 for more information)
b) Implementers **MUST NOT** use the following fields to store the IHI:
   i) Patient ID (0010,0020)
   ii) Patient Comments (0010,4000)
c) Implementers **MUST NOT** treat `PatientID (0010,0020)` as a globally unique identifier.

Patient Identifiers (Other data):

d) JMIX envelopes **SHOULD** include a `patient.identifiers` entry containing the IHI[4]

[4] see https://github.com/aurabx/jmix/blob/main/spec/envelope/metadata.md.

e) FHIR **MUST** include `ImagingStudy.subject.identifier` of type IHI[5].

Patient Reconciliation:

f) When processing or reconciling patient identities across systems, implementers **SHOULD** use the precedence table from Appendix 6.

[5] see https://build.fhir.org/ig/hl7au/au-fhir-core/StructureDefinition-au-core-patient-definitions.html#key_Patient.identifier:ihi.

Appendix: Dealing with locally scoped DICOM identifiers

In a federated imaging ecosystem, each participating institution may generate local `PatientID (0010,0020)` values independently. These identifiers:

- Are typically unique only within the local PACS or RIS domain
- Frequently lack global namespacing or issuer qualification
- Are often overwritten or regenerated on import/export

This results in a high probability of **collisions**, where:

- Two different patients may share the same `PatientID` value across systems
- A single patient may have multiple `PatientID` values across providers

To ensure safe and deterministic interoperability, the Imaging Exchange Framework (IEF) defines the following **collision resolution strategy**:

1. Principle: Never Rely on PatientID Alone for Identity

Implementers **MUST NOT** treat `PatientID (0010,0020)` as a globally unique identifier. It **MUST** only be used in conjunction with additional metadata, such as:

- `IssuerOfPatientID (0010,0021)`
- `OtherPatientIDsSequence (0010,1002)` entries
- Context from the **Data Source** (e.g. site ID, facility OID, or trusted certificate)
- Federated patient match logic (e.g. IHI, name+DOB, or FHIR-based reconciliation)

2. Recommended Identifier Precedence Strategy

When processing or reconciling patient identities across systems, implementers **SHOULD** apply the following precedence:

Priority	Source	Field	Notes
1	JMIX / FHIR envelope	`patient.identifier`	Globally scoped, immutable
2	OtherPatientIDsSequence	`(0010,1002)` with `TypeOfPatientID = NATIONAL`	Contains IHI
3	OtherPatientIDsSequence	`(0010,1002)` with `TypeOfPatientID = LOCAL`	For institutional matching
4	PatientID + IssuerOfPatientID	`(0010,0020)` + `(0010,0021)`	Namespaced local ID
5	PatientID only	`(0010,0020)`	Risky, fallback only

If multiple records resolve to the same `PatientID` but differ on higher-precedence identifiers (e.g. IHI), systems MUST treat them as distinct patients.

3. Storing Source Metadata in JMIX or FHIR

To support traceability and safe resolution:

- JMIX envelopes **SHOULD** include a `patient.identifiers` entry containing the IHI[6]
- FHIR **MUST** include `ImagingStudy.subject.identifier` of type IHI[7].

4. Collision Warnings and Manual Review

If a gateway or receiving system detects multiple `PatientID` records that:

- Share the same value
- Differ in IHI or other identifiers
- Originate from different facilities

It **MUST** raise a **collision alert**, log the discrepancy, and where possible:

- Prefer records with matching IHIs
- Defer automated linkage
- Flag the encounter for **manual resolution** with full auditability

6. Practical Implications

Scenario	Resolution
Two studies from different sites with `PatientID = 123456`	Treated as distinct unless higher-level match (IHI or registry match) occurs
A single patient with different MRNs across sites	Reconciled via IHI or matched using patient demographics
Study lacks IHI or registry identifier	Resolved using `IssuerOfPatientID` + site ID or flagged for manual review
Teaching set with scrubbed identifiers	Treated as non-identifiable, no matching attempted

[6] see https://github.com/aurabx/jmix/blob/main/spec/envelope/metadata.md.
[7] see https://build.fhir.org/ig/hl7au/au-fhir-core/StructureDefinition-au-core-patient-definitions.html#key_Patient.identifier:ihi.

Appendix: Handling IHIs in DICOM

1. Preferred DICOM Tag: Other Patient IDs Sequence (0010,1002)

Implementers **MUST** store the IHI in the `Other Patient IDs Sequence (0010,1002)` as a distinct identifier entry. This method avoids interference with the PACS-reserved `Patient ID (0010,0020)` field and enables systems to safely manage multiple identifiers for the same patient.

Each sequence item SHOULD include:

- **Patient ID** (0010,0020): set to the 16-digit IHI (e.g. `8003601234567890`)
- **Issuer of Patient ID** (0010,0021): set to the IHI OID namespace: `1.2.36.1.2001.1003.0`
- **Type of Patient ID** (0010,1002): set to `"NATIONAL"`

Example:

```
Shell
(0010,1002) SQ  # Other Patient IDs Sequence
  > (0010,0020) LO [8003601234567890]
  > (0010,0021) LO [1.2.36.1.2001.1003.0]
  > (0010,1002) CS [NATIONAL]
```

This approach ensures that:

- Local PACS/RIS systems can retain their internal Patient ID management.
- IHIs are preserved through federation and indexing workflows.
- Systems implementing FHIR or JMIX can extract and reference the IHI for cross-system discovery.

2. Fields to Avoid

Implementers MUST NOT use the following fields to store the IHI:

1) **Patient ID (0010,0020)**
   This tag is commonly overwritten or re-mapped by PACS and RIS systems. Using it to store the IHI risks data loss, identifier collisions, and failure in local workflows.

2) **Patient Comments (0010,4000)**
   This field is free-text, non-machine-readable, and unsuitable for structured patient identification. It MUST NOT be used to store IHIs.

3. Use in Modern Transport Contexts

Where imaging is transmitted using **JMIX** or **FHIR**, the IHI SHOULD also be included in:

- The `ImagingStudy.subject.identifier` field (FHIR)
- The `patient.identifier` property in the JMIX envelope
- Optional provenance or assertion metadata, cryptographically bound to the imaging payload

This ensures the IHI remains available for systems operating outside of DICOM-only environments, and can be redacted as required for privacy.

4. Redaction Requirements

The IHI is considered personally identifiable information (PII) under the **Australian Privacy Principles (APPs)**. Implementers MUST ensure:

- IHIs are excluded from **de-identified** datasets (e.g. for research, teaching, public sharing)
- Automated workflows applying redaction or anonymisation remove the IHI from `(0010,1002)` and any envelope metadata
- De-identification is logged and auditable where required