SNOMED CT Overview

Open Clinical History uses **SNOMED CT** as its primary clinical terminology foundation.

SNOMED CT provides a standard way to represent clinical concepts such as:

- disorders
- findings
- symptoms
- procedures
- anatomical structures
- clinical situations
- observable entities
- substances
- organisms

Rather than relying only on free-text labels such as:

```
heart attack
```

Open Clinical History can associate an event with the corresponding structured SNOMED CT concept.

This allows clinical information to be represented in a way that is more consistent, searchable and interoperable.

What Is SNOMED CT?

SNOMED CT is a comprehensive international clinical terminology maintained by **SNOMED International**.

Each clinical concept has a unique numeric identifier.

For example, conceptually:

```
Clinical concept
      |
      v
```

```
SNOMED CT identifier
      |
      v
Preferred clinical term
      |
      +--> synonyms
      |
      +--> relationships to other concepts
```

A SNOMED concept is therefore more than a text label.

It can include:

- a unique concept identifier
- a Fully Specified Name
- preferred terms
- synonyms
- hierarchical relationships
- clinical attributes
- historical and terminology metadata

Why Open Clinical History Uses SNOMED CT

Clinical records contain enormous variation in wording.

The same condition may appear as:

```
myocardial infarction
heart attack
acute MI
previous MI
non-ST elevation myocardial infarction
```

A text-only system may treat these as unrelated strings.

SNOMED CT provides a common clinical terminology underneath those expressions.

Conceptually:

```
Different source wording
        |
        v
Clinical interpretation
        |
        v
SNOMED CT concept
        |
        v
Consistent structured history
```

This provides Open Clinical History with a terminology layer that is independent of the wording used in the original document.

SNOMED CT Does Not Replace the Source Record

Open Clinical History retains the source evidence from which a clinical event was extracted.

The SNOMED concept is a structured representation of that event.

It is not intended to replace the original wording.

For example:

```
SOURCE

"Patient sustained severe bruising to the left shoulder
following a motor vehicle accident."
```

may become a structured event containing:

```
Title:
Severe contusions to left shoulder
```

```
SNOMED CT:
Contusion (disorder)

Anatomy:
Left shoulder
```

The source evidence remains available for review.

SNOMED CT Concepts

A SNOMED CT concept represents one clinical meaning.

Each concept has a stable numeric identifier.

For example:

```
125667009
```

may identify a particular clinical concept independently of the language used to display it.

This is important because text descriptions can change or multiple synonyms can refer to the same concept.

Conceptually:

```
125667009
    |
    +--> Preferred term
    |
    +--> Synonym
    |
    +--> Another synonym
```

The identifier provides the stable reference.

Fully Specified Name

The **Fully Specified Name**, commonly abbreviated:

```
FSN
```

describes the concept unambiguously.

It normally includes a semantic tag.

For example:

```
Contusion (disorder)
```

or:

```
Kidney structure (body structure)
```

The semantic tag helps distinguish concepts that may otherwise have similar wording.

Common semantic tags include:

```
(disorder)
(finding)
(procedure)
(body structure)
(observable entity)
(situation)
(substance)
```

Preferred Terms and Synonyms

SNOMED CT concepts can have multiple descriptions.

For example, a concept might contain:

```
Fully Specified Name
Preferred term
```

```
Synonym
Synonym
Synonym
```

This is particularly useful for clinical records because clinicians do not always use the same terminology.

Open Clinical History builds local terminology indexes from these descriptions so it can search efficiently during patient processing.

SNOMED CT Relationships

One of the most important features of SNOMED CT is that concepts are connected to other concepts.

These relationships allow Open Clinical History to understand more than simple terminology.

IS-A Relationships

SNOMED CT contains hierarchical:

```
IS-A
```

relationships.

For example, conceptually:

```
Left kidney structure
        |
        IS-A
        v
Kidney structure
        |
        IS-A
        v
Urinary system structure
        |
```

```
        IS-A
        v
Body structure
```

This means that a very specific anatomical concept can still be understood as belonging to broader anatomical groups.

Open Clinical History uses these relationships when constructing its anatomical hierarchy.

Clinical Attribute Relationships

SNOMED CT concepts can also have relationships describing clinical meaning.

For example, a disorder may contain an attribute such as:

```
Finding site
```

pointing to a body structure.

Conceptually:

```
Clinical disorder
      |
      | finding site
      v
Body structure
```

Procedures can similarly contain anatomical relationships.

These relationships are extremely important to Open Clinical History because they help connect a clinical concept to the relevant anatomical image.

SNOMED CT and the Anatomical Viewer

Open Clinical History combines SNOMED terminology with its installed anatomical image catalogue.

The relationship is broadly:

```
Clinical event
      |
      v
SNOMED clinical concept
      |
      v
SNOMED anatomical relationship
      |
      v
SNOMED body structure
      |
      v
Open Clinical History body layer
```

For example:

```
Clinical concept
      |
      | Finding site
      v
Left shoulder structure
      |
      v
Installed left shoulder image layer
```

This allows the anatomical visualisation to be derived from structured clinical terminology rather than relying only on keyword matching.

Local SNOMED CT Database

Open Clinical History uses a **local SNOMED CT database**.

It does not send clinical terms to an external SNOMED service every time a patient history is processed.

The raw RF2 release is imported into the local database.

Conceptually:

```
SNOMED CT RF2 release
        |
        v
Open Clinical History importer
        |
        v
Local SNOMED source tables
        |
        v
Compact runtime indexes
```

This provides:

- predictable performance
- local terminology control
- reduced external dependencies
- faster clinical lookup
- terminology version consistency

RF2

SNOMED CT releases are distributed using:

```
RF2
```

or **Release Format 2**.

A release commonly contains:

```
Full
Snapshot
Delta
```

Open Clinical History currently uses the:

```
Snapshot
```

release data for normal imports.

Snapshot Release

A Snapshot contains the current effective state of the terminology at the release date.

It does not contain every historical version of every component.

For Open Clinical History this is usually the most useful representation because it provides the active terminology required for current clinical processing.

Full Release

The Full release contains the historical version history of SNOMED components.

It is substantially larger.

Open Clinical History does not require the Full files for its normal terminology build.

Delta Release

A Delta contains only changes since the preceding release.

Open Clinical History currently performs its terminology build from Snapshot files rather than incrementally applying Delta files.

RF2 Files Used by Open Clinical History

The current importer uses four primary Snapshot datasets.

Concept Snapshot

Typical filename:

```
sct2_Concept_Snapshot_...
```

This contains the SNOMED concept identifiers and their current active state.

Description Snapshot

Typical filename:

```
sct2_Description_Snapshot...
```

This contains:

- Fully Specified Names
- synonyms
- terminology descriptions

Open Clinical History uses these descriptions to build its terminology indexes.

Relationship Snapshot

Typical filename:

```
sct2_Relationship_Snapshot_...
```

This contains SNOMED relationships including:

```
IS-A
```

and clinical attribute relationships.

This file is required by the current Open Clinical History implementation.

It is used for:

- hierarchy construction
- anatomical relationships
- clinical concept-to-image mapping

Language Reference Set

Typical filename:

```
der2_cRefset_LanguageSnapshot...
```

This identifies preferred and acceptable descriptions for the relevant language edition.

For an Australian edition this can include:

```
en-AU
```

language preferences.

SNOMED Editions

SNOMED CT has an International Edition and may also be distributed as national editions or extensions.

For example, an Australian installation may use an Australian SNOMED CT release.

The national edition can contain:

- International core terminology
- national extensions
- national preferred terms
- local reference sets

Open Clinical History records the release being imported so the terminology and derived image database can remain consistent.

Open Clinical History SNOMED Pipeline

The SNOMED build broadly follows:

```
RF2 Snapshot
      |
      v
Import concepts
      |
      v
Import descriptions
      |
      v
Import relationships
      |
      v
Import language refset
      |
      v
Verify release
      |
      v
Build anatomical hierarchy
      |
      v
Build clinical terminology indexes
      |
```

```
      v
Build image mappings
```

The result is both:

1. a local representation of the SNOMED source terminology; and
2. compact indexes optimised specifically for Open Clinical History.

Compact Runtime Indexes

The complete SNOMED source tables can contain millions of records.

Searching those tables repeatedly during every patient import would be inefficient.

Open Clinical History therefore builds compact clinical lookup tables.

Important examples include:

```
snomed_health_history_lookup
snomed_health_history_term_lookup
```

These contain the subset of terminology and descriptions useful to the patient-history pipeline.

Conceptually:

```
Millions of RF2 records
        |
        | build-time processing
        v
Compact terminology index
        |
        | runtime
        v
Fast patient-history matching
```

SNOMED Body Structure Candidates

Open Clinical History also builds a compact local catalogue of SNOMED:

```
(body structure)
```

concepts.

This is used when connecting the installed anatomical PNG catalogue to SNOMED.

Rather than allowing an LLM to invent anatomical concept IDs, Open Clinical History searches the local body-structure catalogue.

How a Clinical Event Is Matched

When Open Clinical History extracts a clinical event from a source document, it generates clinical search terms.

For example:

```
acute cerebral infarction
cerebral infarction
stroke
```

The terminology resolver then searches the local SNOMED indexes.

Conceptually:

```
Extracted clinical event
        |
        v
Generate search terms
```

```
        |
        v
Search local SNOMED indexes
        |
        v
Candidate concepts
        |
        v
Resolve best candidate
        |
        v
Verified SNOMED concept
```

Deterministic Matching

Where possible, Open Clinical History resolves terminology without an LLM.

Examples include:

- exact Fully Specified Name
- exact preferred term
- exact synonym
- curated local alias

A unique compatible exact match can therefore be selected deterministically.

Grounded AI Resolution

Some clinical language is ambiguous.

In these cases Open Clinical History can use an LLM to help choose between candidate SNOMED concepts.

However:

> The LLM is not allowed to invent a SNOMED identifier.

Instead:

```
Local SNOMED search
        |
        v
Candidate A
Candidate B
Candidate C
Candidate D
        |
        v
LLM evaluates supplied candidates
        |
        v
One supplied candidate selected
```

The available concepts come from the local SNOMED database.

SNOMED remains the terminology authority.

Why This Matters

Without candidate grounding, a generative model could potentially return:

- a nonexistent concept
- an obsolete concept
- a concept from the wrong semantic category
- a plausible-looking numeric identifier

Open Clinical History avoids treating model-generated identifiers as authoritative.

Concept IDs are sourced from the locally installed SNOMED release.

Confidence Thresholds

Grounded terminology resolution also uses configured confidence thresholds.

For example:

```
llm_grounded_resolution_min_confidence
```

controls whether a grounded result is considered strong enough to accept.

Low-confidence results can remain unresolved rather than being forced into the patient record.

This is deliberate.

A missing mapping is preferable to a confident-looking but incorrect clinical code.

Critic and Repair Processing

Open Clinical History can perform additional terminology checks for uncertain mappings.

Conceptually:

```
Initial resolution
      |
      v
Confidence / quality check
      |
      +---- suitable ------> accept
      |
      +---- uncertain
                |
                v
              critic
                |
                v
        revised search terms
                |
                v
        local SNOMED search
```

The repair process still operates against locally retrieved candidate concepts.

Learning Queue

If a clinical term cannot be resolved, it can be placed into the:

```
SNOMED Re-index / Learning Queue
```

rather than forcing a match.

Later, after:

- terminology indexes improve
- aliases are added
- SNOMED is updated
- matching logic improves

the term can be retried.

Conceptually:

```
Unresolved term
      |
      v
Learning Queue
      |
      v
Terminology improved
      |
      v
Re-index
      |
      v
SNOMED concept found
```

The original patient document does not need to be extracted again.

Curated Aliases

Open Clinical History can also maintain local aliases for terminology encountered in real clinical records.

For example, an organisation may consistently use a particular abbreviation or historical expression.

Conceptually:

```
Local expression
      |
      v
Curated alias
      |
      v
Known clinical concept
      |
      v
SNOMED CT
```

Aliases should be governed carefully.

They are intended to improve terminology recognition, not to override the clinical meaning of SNOMED.

SNOMED CT and Laterality

Laterality may be represented by:

- the SNOMED concept itself
- a related body-structure concept
- explicit clinical source information
- Open Clinical History event metadata

Open Clinical History attempts to preserve laterality such as:

```
left
right
bilateral
midline
```

when resolving both terminology and anatomical imagery.

Laterality conflicts can trigger mapping warnings rather than silently selecting the wrong anatomical side.

SNOMED CT and Date Information

SNOMED represents **what** the clinical concept means.

It does not determine:

```
when the event happened
```

Open Clinical History stores temporal information separately.

For example:

```
SNOMED concept:
Contusion (disorder)

Event date:
21 April 1988

Outcome:
Resolved
```

These are different dimensions of the clinical event.

SNOMED CT and Clinical Status

Similarly, the terminology concept does not by itself establish whether an event is:

```
active
ongoing
resolved
historical
```

Open Clinical History derives and stores clinical status separately from the SNOMED coding.

This allows the same SNOMED concept to appear at different times with different clinical states.

SNOMED CT and Severity

Clinical severity is also stored separately.

For example:

```
Concept:
Pneumonia

Severity:
Major

Status:
Resolved
```

Severity belongs to the clinical event rather than redefining the SNOMED concept.

SNOMED CT and Source Evidence

Every imported event can retain:

```
source evidence
```

alongside the SNOMED coding.

This allows a reviewer to see both:

```
What did the source say?
```

and:

```
How did Open Clinical History classify it?
```

The terminology mapping therefore remains auditable.

SNOMED Version Matters

SNOMED CT changes over time.

New releases can:

- add concepts
- inactivate concepts
- add synonyms
- change preferred descriptions
- modify relationships
- introduce extension content

Open Clinical History therefore treats the installed SNOMED release as part of the system's terminology state.

Updating SNOMED CT

When a new authorised release is installed:

```
Download new RF2 release
        |
        v
Configure new release path
        |
        v
Run SNOMED Import
```

```
        |
        v
Rebuild Image/SNOMED database
        |
        v
Process Learning Queue
```

Rebuilding the image database is important because anatomical mappings depend on the locally installed SNOMED relationships.

Previously unresolved terminology may also become resolvable after an update.

Keeping Releases Separate

It is recommended that different SNOMED releases be kept in separate protected directories.

For example:

```
/var/www/private/snomed/
├── AU1000036_20260731/
├── AU1000036_20270131/
└── AU1000036_20270731/
```

This makes:

- rollback
- auditing
- troubleshooting
- release comparison

easier.

SNOMED Files Should Not Be Web Accessible

The raw RF2 distribution should be stored outside the public web directory.

For example:

```
/var/www/private/demo/snomed/
```

rather than:

```
/var/www/html/snomed/
```

The application and import workers need filesystem access to the release.

Website users should not have direct HTTP access to the raw terminology files.

Licensing

SNOMED CT is not simply a public-domain terminology dataset.

Use and redistribution are subject to SNOMED International licensing arrangements and applicable national licensing conditions.

Administrators are responsible for obtaining SNOMED CT through an authorised distribution mechanism appropriate to their location and intended use.

Open Clinical History does not redistribute the SNOMED CT RF2 release itself.

What Open Clinical History Stores

There are effectively three layers of terminology data.

1. Raw imported SNOMED data

Examples:

```
sct_concept
sct_description
sct_relationship
```

These represent the imported RF2 terminology.

---

2. Open Clinical History derived indexes

Examples:

```
snomed_health_history_lookup
snomed_health_history_term_lookup
snomed_body_structure_candidate
snomed_body_hierarchy
```

These are optimised for OCH processing.

---

3. Patient event coding

Committed clinical events store the SNOMED concept selected for that event.

Conceptually:

```
RF2 terminology
      |
      v
OCH terminology indexes
      |
      v
Patient event
      |
      v
SNOMED concept ID
```

What SNOMED CT Provides to Open Clinical History

SNOMED CT serves several separate purposes.

Function	SNOMED role
Clinical classification	Gives extracted events a structured clinical concept
Synonym handling	Connects different clinical expressions
Semantic typing	Distinguishes disorders, findings, procedures, anatomy, etc.
Anatomical hierarchy	Provides body-structure relationships
Clinical anatomy mapping	Connects disorders/procedures to body sites
Learning	Provides the terminology target for unresolved expressions
Interoperability	Provides standard identifiers usable outside OCH

SNOMED CT and Interoperability

Using SNOMED CT makes the clinical event more portable than an application-specific text label.

For example:

```
Open Clinical History event
        |
        v
SNOMED CT concept
        |
        +--> FHIR resource
        |
        +--> clinical integration
```

```
        |
        +--> analytics
        |
        +--> another SNOMED-aware system
```

SNOMED coding does not automatically make two systems interoperable, but it provides a shared clinical vocabulary that significantly improves the ability to exchange structured meaning.

Relationship to FHIR

SNOMED CT and FHIR solve different problems.

Broadly:

```
FHIR
=
how healthcare information is structured and exchanged

SNOMED CT
=
how clinical meaning is represented
```

For example, a FHIR Condition might contain a SNOMED CT code.

Conceptually:

```
FHIR Condition
      |
      v
code
      |
      v
SNOMED CT concept
```

Open Clinical History can therefore use SNOMED terminology as part of future or current standards-based interoperability without treating SNOMED itself as the transport format.

SNOMED CT Is Not an AI Classification System

SNOMED CT is a clinical terminology.

Gemini or another LLM may assist Open Clinical History in interpreting free text, but SNOMED itself is independent of the LLM.

Conceptually:

```
LLM
=
helps interpret clinical language

SNOMED CT
=
provides the controlled clinical terminology
```

The two functions should not be confused.

Open Clinical History Design Principle

A central design principle is:

> **AI may help interpret clinical language, but it does not define the terminology catalogue.**

Open Clinical History therefore prefers:

```
Source evidence
      |
      v
Clinical extraction
      |
      v
Local SNOMED candidates
      |
      v
Grounded resolution
      |
      v
Verified local SNOMED concept
```

over:

```
Source text
      |
      v
Ask AI for a SNOMED ID
      |
      v
Trust whatever it returns
```

The latter is deliberately avoided.

Limitations

SNOMED CT is extremely comprehensive, but terminology resolution is not always straightforward.

Challenges include:

- ambiguous abbreviations
- historical terminology
- very specific narrative descriptions
- local clinical language
- incomplete source information
- differences between SNOMED editions
- concepts that do not map neatly to a visual body layer

A clinical event can therefore legitimately remain unresolved.

Open Clinical History treats unresolved terminology as a condition to investigate, rather than automatically forcing a questionable match.

Related Open Clinical History Pages

SNOMED Import

Used to import and validate an RF2 Snapshot release.

See:

How to Import SNOMED CT

Image/SNOMED Summary Build

Builds the connection between SNOMED body structures and the Open Clinical History anatomical PNG catalogue.

See:

Building the Image / SNOMED Database

SNOMED Re-index / Learning Queue

Retries terminology that could not be resolved during patient processing.

See:

SNOMED Re-index / Learning Queue

Patient History

Displays the final SNOMED-coded clinical events alongside their chronology, source evidence and anatomical mappings.

See:

Reviewing and Understanding a Patient History

Summary

Within Open Clinical History, SNOMED CT provides the controlled clinical terminology that sits between:

```
Raw clinical language
```

and:

```
Structured longitudinal clinical history
```

The complete relationship can be viewed as:

```
Raw source document
        |
        v
Clinical event extraction
        |
        v
Local SNOMED candidate search
        |
        v
Grounded terminology resolution
        |
```

```
        v
SNOMED CT concept
        |
        +----------------------+
        |                      |
        v                      v
Clinical classification   Anatomical relationships
        |                      |
        v                      v
Patient event             Body-image mapping
        |                      |
        +-----------+----------+
                    |
                    v
       Longitudinal clinical history
```

SNOMED CT is therefore not just a label attached to an event.

In Open Clinical History it provides:

- clinical meaning
- terminology consistency
- semantic relationships
- anatomical hierarchy
- structured classification
- a foundation for interoperability

while the original clinical source remains available to show **where that structured interpretation came from**.

Revision #1
Created 2026-08-25 10:43:51 UTC by Admin
Updated 2026-08-25 10:44:09 UTC by Admin