

How to Import SNOMED CT

Open Clinical History uses a local **SNOMED CT RF2 Snapshot release** to provide clinical terminology, terminology search, anatomical hierarchy information and clinical concept-to-body-layer mappings.

The SNOMED distribution is large and the initial import can take some time.

Once imported, Open Clinical History builds optimised local lookup tables so the application does not need to process the raw RF2 files during normal operation.

1. Obtain a SNOMED CT RF2 Release

Download an authorised SNOMED CT RF2 release from your SNOMED distribution provider.

Open Clinical History expects an **RF2 Snapshot** release.

A typical release contains directories similar to:

```
Content
Language
Map
Metadata
Refset
Terminology
```

For example:

```
root@server:/var/www/private/demo/snomed# ls

Content  Language  Map  Metadata  Refset  Terminology
```

The exact directory structure can vary slightly between SNOMED editions and releases.

Open Clinical History searches the release directories for the required RF2 files rather than depending on exact release filenames.

2. Store the Release Outside the Web Root

SNOMED CT source files should **not be publicly accessible through the web server**.

Store the extracted RF2 release in a private server directory outside the application's public `DocumentRoot`.

For example:

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

Do **not** place the RF2 distribution somewhere such as:

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

or another directory that Apache or Nginx can serve directly.

The SNOMED files are licensed content and should not be exposed for direct download.

File permissions

The files must not be web-accessible, but the Open Clinical History PHP process and background import worker still need permission to read them.

For example, your final structure may look similar to:

```
/var/www/  
├─ private/  
│   └─ demo/  
│       └─ snomed/  
│           ├── Content/  
│           └── Language/
```

```
|      |— Map/
|      |— Metadata/
|      |— Refset/
|      |— Terminology/
|
|— openclinicalhistory/
   |— public application files
```

The important distinction is:

“ The SNOMED directory must be readable by the application, but must not be reachable as a public web URL.

3. Configure the RF2 Release Path

Open:

Admin → Configuration

Locate:

SNOMED CT RF2 → RF2 release path

Enter the absolute server path containing the extracted SNOMED release.

For example:

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

Save the configuration.

Open Clinical History accepts either:

- the release root; or
- a compatible RF2 **Snapshot** directory.

The importer searches below the configured directory for the required RF2 files.

4. Open the SNOMED Import

From the administration interface, open:

SNOMED Import

The page inspects the configured release before starting an import.

At the top of the page you should see information similar to:

```
Release: AU1000036_20260731
/var/www/private/demo/snomed

Source: app_config
```

The screen also displays detected information about the currently staged SNOMED database, including:

- active concepts
- relationships
- anatomical/site relationships
- presence of the International core

5. Check the RF2 File Inventory

Before starting the pipeline, check the **RF2 file inventory** at the bottom of the page.

Open Clinical History currently requires four RF2 Snapshot datasets.

Dataset	Purpose	Required
---------	---------	----------

Concepts	SNOMED concept identifiers and status	Yes
Descriptions	Clinical names, terms and synonyms	Yes
Relationships	SNOMED hierarchy and clinical relationships	Yes
Language refset	Preferred terminology for the selected language edition	Yes

Each required file should show:

```
ready
```

For example:

```
Import concepts      ready
Import descriptions  ready
Import relationships ready
Import language refset ready
```

If a required file shows:

```
missing
```

do not force the import. Check the RF2 package and configured release path first.

Are SNOMED Relationships Required?

Yes, in the current Open Clinical History implementation.

The relationship file is not simply retained as unused SNOMED reference data.

Open Clinical History uses SNOMED relationships for several important functions.

Anatomical hierarchy

The application follows SNOMED **IS-A** relationships to build its anatomical hierarchy.

Conceptually:

```
Body structure
  |
  v
Parent anatomical structures
  |
  v
More specific structures
  |
  v
Specific anatomical site
```

This allows a specific SNOMED body structure to inherit the appropriate Open Clinical History anatomical image layer.

Finding-site relationships

Clinical concepts can contain relationships to anatomical structures.

For example, conceptually:

```
Clinical disorder
  |
  | finding site
  v
Body structure
  |
  v
Open Clinical History image layer
```

These relationships are used when building the concept-to-image mapping.

Without the relationship dataset, Open Clinical History cannot currently build these derived anatomical mappings correctly.

Release verification

The importer also verifies that the RF2 release contains:

- a substantial relationship dataset
- SNOMED `IS-A` relationships
- anatomical/site relationships

A release without them will fail the current verification stage.

6. Start the Pipeline

Once all required RF2 files show as ready, click:

Start or resume pipeline

The import runs as a background process.

You can leave the page open to watch its progress.

The pipeline consists of the following stages.

Step	Stage	Purpose
1	Import concepts	Load the RF2 Concept Snapshot
2	Import descriptions	Load SNOMED names and synonyms
3	Import relationships	Load hierarchy and clinical relationships
4	Import language refset	Load preferred-language terminology
5	Verify release	Confirm that the imported release is usable
6	Validate body image mappings	Validate configured anatomical SNOMED seeds
7	Build anatomical hierarchy	Build body-structure ancestry
8	Build concept-to-image mappings	Connect clinical concepts to anatomical layers
9	Build health-history concept lookup	Build optimised concept lookup data
10	Build health-history term lookup	Build optimised terminology search data
11	Validate health-history catalogue	Check the resulting clinical catalogue
12	Complete	Pipeline finished

7. Be Patient

SNOMED CT is **large**.

A release can contain millions of rows.

A typical Australian release may contain approximately:

Concepts	hundreds of thousands
Descriptions	millions
Relationships	millions
Language refset	millions

The raw RF2 files themselves may also be hundreds of megabytes in size.

The first import therefore takes time.

Some later stages must also analyse and build indexes over millions of imported records, so apparent pauses during the process are not necessarily failures.

Avoid restarting the server or database while the import is running.

8. Monitoring Progress

While an import is running, the SNOMED Import page displays:

- current pipeline stage
- worker process
- elapsed time
- rows processed
- worker heartbeat
- worker output
- error messages

The page updates as the background worker progresses.

For large RF2 files the row count can increase substantially before a stage completes.

9. Safe Restart and Resume

The SNOMED pipeline is designed to be resumable.

If an import is interrupted, return to:

SNOMED Import

and select:

Start or resume pipeline

The importer checks:

- the persistent pipeline state
- the selected RF2 release
- the source-file information
- existing database tables
- current row counts

Completed and still-valid stages are skipped.

A failed or stale stage is rebuilt, followed by any stages that depend on it.

This avoids unnecessarily re-importing millions of SNOMED rows after an interruption.

Rebuild Every Stage

The SNOMED Import screen also contains:

Rebuild every stage, including completed tables

Selecting this forces the entire pipeline to run again.

This includes re-importing the large RF2 source tables.

Do not select this for a normal restart.

Use a full rebuild when, for example:

- installing a new SNOMED CT release
- deliberately replacing all staged terminology
- troubleshooting potentially inconsistent source tables
- testing changes to the SNOMED import process

For an interrupted import, normally use **Start or resume pipeline** without selecting the rebuild option.

10. Successful Completion

The import is complete when all pipeline stages show:

complete

and the final stage reports:

Pipeline complete

At this point Open Clinical History has both:

1. the imported SNOMED CT RF2 source tables; and
2. the derived lookup and anatomical mapping tables used by the application.

Normal patient-history processing can then use the local SNOMED terminology database.

Updating to a New SNOMED Release

When a newer SNOMED release becomes available:

1. Download the new authorised RF2 Snapshot release.
2. Extract it into a new private directory.
3. Update **Admin** → **Configuration** → **SNOMED CT RF2** → **RF2 release path**.

4. Open **SNOMED Import**.
5. Confirm the RF2 inventory identifies the new files.
6. Run the pipeline.

Keeping separate release directories is preferable to overwriting the previous release immediately.

For example:

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

This makes rollback and troubleshooting considerably easier.

Troubleshooting

No RF2 Snapshot release found

Check:

- the configured RF2 path
- that the release has been extracted
- that the PHP/worker user can read the directory
- that the release contains RF2 Snapshot files

The path must be a filesystem path, not a URL.

RF2 file shows **missing**

Check the file inventory.

The current importer expects files matching the standard Snapshot families:

```
sct2_Concept_Snapshot...  
sct2_Description_Snapshot...
```

```
sct2_Relationship_Snapshot...
der2_cRefset_LanguageSnapshot...
```

The release-specific identifiers and dates in the filenames may vary.

Open Clinical History discovers them by their RF2 filename prefixes.

Existing table rows are shown but the RF2 file is missing

The **Table rows** column reports what is currently in the Open Clinical History database.

It does not prove that the corresponding source file exists in the currently selected release.

For example:

```
Import relationships
```

```
File:      missing
Status:    missing
Table rows: 5,061,383
```

means that relationship records exist in the database, but the selected RF2 directory does not currently contain a relationship Snapshot file recognised by the importer.

The existing rows may have come from an earlier release.

Do not treat this as a valid new release until the source inventory also reports the relationship file as:

```
ready
```

Important Security Note

SNOMED CT distributions may be subject to licensing and redistribution conditions.

Do not expose the raw RF2 files through the Open Clinical History website or another publicly accessible directory.

Store them in a protected server-side location and provide access only to the application and authorised administrators.

Revision #1

Created 2026-08-25 09:31:00 UTC by Admin

Updated 2026-08-25 09:44:50 UTC by Admin