SNOMED CT®
The Release Format 2 Value Proposition

The case for early migration to RF2 for IHTSDO, SNOMED CT Extension Providers and Application Developers

Date   2013-04-07

Version   1.02

Amendment History

Version	Date	Editor	Comments
0.01	20120410	Harold Solbrig	First draft released to I&I Committee for comment
0.03	20120517	David Markwell	Reformatted to match IHTSDO template. Updates to increase visibility of benefits and added rationale for early migration to RF2. Outstanding actions Section 5.3 and 5.4 in need of further tidy up. For final review by SIA team and Executive Team before release to community.
1.00	20120623	David Markwell	Revisions of 5.3 and 5.4.
1.01	20121127	David Markwell	Further revisions to Section 5 and 6.
1.02	20130321	David Markwell	Revision based on comments from Alejandro Lopez Osornio, Mikael Nyström and John Gutai.

Review Timetable

Review date	Responsible owner	Comments
YYYYMMDD	Person/group responsible	Summary of action
		(remove or add rows if necessary)

Table of Contents

1 Purpose

This document summarizes the main reasons why all organizations that implement SNOMED CT or maintain extensions to SNOMED CT should support the new standard release format, known as Release Format 2 (RF2).

The document is intended to be read by decision makers who need to understand the reasons for switching over to the new format.

- IHTSDO Affiliates: Vendors and Developers responsible for making decisions about updates to application design to accommodate use of the use of terminology distribution files that conform to the new format.
- IHTSDO Members: National Release Centers making decisions about the format used for distributing National Extensions to SNOMED CT
- IHTSDO Members: Healthcare Provider Organizations setting requirements for applications that use SNOMED CT.

2 Introduction

This paper summarizes the main benefits that support for RF2 offers to those implementing SNOMED CT in software application and those developing SNOMED CT extensions at National or local levels. It summarizes the key changes and highlights the added value that each change offers to those involved in implementation and deployment. Gaining these benefits inevitably requires some changes.

Key elements of the design of SNOMED CT remain unchanged. Moving to RF2 does not affect the structure or content of records that include SNOMED CT codes or expressions. However, RF2 enables more consistent and efficient distribution and maintenance of the content of the terminology. It also adds a consistent but highly flexible mechanism for maintaining and distributing derivatives based on SNOMED CT. This mechanism supports subsets, term preferences, mapping to/from other code systems as well as many other requirements. Each of these changes brings significant benefits which are identified later in this document.

To support backward compatibility, RF2 release files can be converted to the original RF1 format by using a utility provided by the IHTSDO. However, this service is a short-term expedient to cover the transition period and should not be regarded as reason to delay moving to RF2. The result of RF2 to RF1 migration is an effective downgrade as features not supported in RF1 must be removed. Furthermore, the IHTSDO and organizations that provide SNOMED CT Extensions must maintain "compatibility" files to support the process of conversion.

3 Background

The original SNOMED CT release format[1] has met the requirements for distribution of the terminology since it was first released in 2002. Since then it has been adopted by twenty IHTSDO Member countries and used in more than fifty countries. Experience of use in large and small-scale installations has led to enhancement of its content and to an evolving understanding of the requirements for more effective implementation. Implementers have reported limitations in version management and in the approach to subsets, cross-maps and other derivatives. Analysis of these issues led to detailed review of the structure of the files used to distribute SNOMED CT. The outcome of this review was a consensus in favor or a redesigned set of file structures, known as Release Format 2 (RF2).

Some of the major changes over the past ten years include:
- Translation of SNOMED CT into multiple languages;
- Introduction of maps to external code systems such as ICD-O, ICD-9-CM, ICD-10;
- An increase in the use of country, application and domain specific maps and subsets;
- Growth of sophistication and innovation amongst application developers and within clinical and research communities;
- Increased use of national and local extensions, with the intention of sharing and transferring content between them;
- A broader understanding of the value of the formal logic that forms the foundation for SNOMED CT definitions.

The net result of these changes has been more demanding requirements for support of versioning and extensibility. RF1 was unable to meet these requirements as a result of its built-in limitations. These included:
- Overloading of fields serving multiple purposes;
- Lack of a simple consistent way to identify changes between releases;
- Inability to support anticipated future enhancements in the formal logic model;
- Absence of explicit representation of metadata specified by enumerated lists in documentation (e.g. values for concept status, description type, characteristic type, etc.).

In response this, a task force was formed with the goal of assembling a set of requirements for the new release format and proposing changes and enhancements that would transform the existing format into one that could serve the organization for its second decade and beyond. RF2 is the final product of this task force.

In January 2012, after several earlier preview releases, RF2 was adopted as the main release format for the SNOMED CT International Release. RF2 releases are produced directly from the master copy of SNOMED CT, maintained in the IHTSDO Workbench. To provide transitional support for applications that are not yet able to read RF2 files, the International Release of SNOMED CT is also made available in older release file format (RF1). RF1 releases are current[2] produced from the RF2 data using a conversion utility which, in order to fit the data into the RF1 format, inevitably results in loss of the additional features available to users of the RF2 release.

---

[1] The original SNOMED CT release format is now referred to as Release Format 1 (RF1).

[2] For practical reasons, the transitional period during which conversion is viable and useful will be limited.

4 RF2 Feature Summary

4.1 Overview

RF2 represents a *refactoring* of the original RF1 model with an emphasis on consistency, comprehensibility and extensibility. The conceptual and linguistic aspects of the RF1 model have been cleanly separated. A single, consistent versioning model has been implemented for all tables. A core set of extensible "reference sets" have been introduced to address today's functional requirements with an accompanying model that allows the definition of new structures in a consistent and approachable fashion. New capabilities have been introduced to support or the integration, assembly and dissemination of and application specific distributions.

4.2 Consistent Versioning

4.2.1 Description

RF2 introduces a consistent versioning mechanism. Every row in every RF2 distribution file includes the following information:

- A unique, permanent component identifier. This links modified versions of this component.
- A time stamp that identifies the point in time (version) when the state represented by row became effective. This distinguished different modified versions of the same component.
- A flag that indicates whether the identified versioned component is intended for "active use".
- A module identifier, which allows dependencies between components released in different modules to be recognized and resolved.

4.2.2 Benefits

A consistent and unambiguous versioning mechanism is critical for the effective maintenance and implementation of SNOMED CT. The RF1 approach to versioning varied between files and, version comparisons were possible they required the separate versions to be exhaustively compared.

The consistency of the RF2 approach to version simplifies the task of maintaining SNOMED CT content. It also provides a simple mechanism that allows implementers to access and compare different versions in real-time. As all the distribution represent versioning in exactly the same way, it is no longer necessary to write table specific code for processing updates. Instead common code can be used to provided version specific views of concept, description, relationship and reference set that represent subsets, maps and language preferences.

The RF2 versioning mechanism allows application developers and implementers to:

- Unambiguously identify which version of SNOMED CT is in current use;
- Determine what changed between two versions or between two points in time;
- Determine what *will* change when a new version is installed;
- Evaluate the impact of version updates on a running system;
- Allow dependencies between different releases and components to be recognized and resolved.

4.3 Flexible Updating

4.3.1 Description

The RF2 versioning mechanism includes three different release types. Each of these is designed to support a particular approach to installing and updating SNOMED CT content in a software application or terminology server.

The three release types are summarized here:

- **Full** – a complete view including every version of every SNOMED CT component ever released.
- **Delta** – this consists of only the entries that have changed from the previous release.
- **Snapshot** – a view of the latest version of every SNOMED CT component at the time of this release.

4.3.2 Benefits

Organizations responsible for SNOMED CT Extensions maintenance should track all changes made in each version. This requirement is well supported by a Full view of the content and derivatives they are maintaining. The Full RF2 view is the primary release type and can be generated directly by development tools in the IHTSDO Workbench. The Snapshot and Delta release files are generated by standard queries applied to the Full release and offer addition value to implementers.

Implementers benefit from these different release types as each supports a particular updating mode.

- A Full release allows an implementer to fully populate a database which can be used to view the state of SNOMED CT at any given point in time (e.g. before or after any particular release). Therefore, a Full release is ideal for initial installation of SNOMED CT content or a complete replacement of an earlier installation.
- A Delta release can simply be added to an existing database containing the Full view of the previous version. This provides a lightweight way to update SNOMED CT without changing existing content. It is particularly useful, as it provides a simple way to apply updates to individual Extensions without any impact on the International Release or any other installed Extension.
- A Snapshot release file is useful for those implementers who do not need to provide access to earlier versions. Only the latest state of every component at the time of the release is available.

An additional benefit of the RF2 approach to versioning is that it allows content to be maintained, distributed and updated in separate modules. Version interdependencies between modules are represented in a standard way allowing implementers to confirm they have installed the correct versions of all the modules required.

This module-based approach also permits components to be moved from an Extension to the International Release without changing their component identifiers. A new *moduleId* can be assigned to an existing component to indicate a change in "ownership" without creating a new identifier in the International Release namespace. This added flexibility is not supported in RF1 and thus changes of this type by those using RF2 will be invisible to those who have not updated to RF2.

4.4 Consistent mechanism for subsets, maps and other derivatives

4.4.1 Description

Experience of various types of SNOMED CT implementations has led to recognition of a wide range of requirements for references, and associations between, SNOMED CT components. The RF1 Subset mechanism addressed some of these requirements including: subsets, term preferences in different languages and alternative hierarchies. The RF1 Cross Map mechanism supported requirements for mapping from SNOMED CT to classifications. However, these mechanisms were inconsistent with one another, suffered from a lack of adequate support for versioning and could not be extended to meet newly identified requirements.

RF2 addresses these limitations by replacing RF1 Subsets and Cross Maps with a single flexible, extensible and fully versioned approach, known as the "Reference Set mechanism".

The Reference Set mechanism includes:

- Full versioning support using the same approach used throughout RF2.
- Use of SNOMED CT concepts and descriptions to name each Reference Set.
- A formal way to specify and disseminate new Reference Set structures to meet additional requirements.

Several References Sets patterns have been defined including the following:

- A *simple reference set* specifies a set of components (concepts, descriptions, relationships, other reference sets, etc.).   Examples of uses of simple reference sets include representing: a subset of descriptions for a user-interface pick-list, a subset of concepts applicable to a particular field or attribute in a communication specification (a value-set), or a set of reference sets required for a particular purpose.
- An *ordered reference set s*pecifies a set of components that can be ordered, nested or subdivided into groups. Ordered reference sets allow SNOMED CT to be represented to users in a consistent and structured ways that suit particular use cases. Examples of uses of ordered reference sets include representing: prioritization of descriptions in a search list and hierarchies used for user interface navigation.
- A *language reference set* specifies the applicability of descriptions in a given language or dialect. For each concept, a language references set also indicates one description of each description type that is preferred for use in that language or dialect. Language references sets allow different language communities to adapt SNOMED CT so the terms displayed are from the relevant language and spelt in accordance with national or local conventions.
- A *simple map reference set* supports simple (usually one-to-one) maps between SNOMED CT concepts and values in alternate coding schemes. Examples of these maps include links between SNOMED CT concepts and codes in antecedent terminologies include NHS Clinical Terms Version 3 (the Read Codes) and earlier versions of SNOMED.
- A *complex or extended map reference set* supports maps from SNOMED CT concepts to one or more codes in a target scheme. These maps contain additional information including the priority and correlation to the map between the source and target. In addition, a complex map may include a machine readable rule and/or human readable advice which can be used to choose between alternative maps at run-time.   An extended map type reference set is used to provide clear, unambiguous rule-based maps from SNOMED CT to another code system or classification (e.g. ICD-10).

- An *association reference set* support links between pairs of SNOMED CT components to represent a non-defining relationship between them. For example, "Historical association reference sets" are used to represent relationships between inactive components and the active components that replace them.
- The *module dependency reference set* represents the dependency between particular versions of different modules (a module is a collection of components that are part of the SNOMED CT International Release or part of an Extension).

New reference set structures are defined using *reference set descriptor reference set* – a reference set that uses the SNOMED CT core model (concepts, descriptions and relationship) to define the structure of present and future reference sets.

4.4.2 Benefits

Consistent versioning support in the Reference Set mechanism means that derivatives that relate to SNOMED CT can now be maintained more effectively. In particular, the RF1 requirement for a new Subset identifier for each version of a Subset no longer applies. This means that once released each Reference Set has a permanent identifier, which is makes it easier to implement updates and to track changes in these sets.

The Reference Sets specified so far address a range of existing or very-near-future needs identified by the IHTSDO community. However, we expect the requirements of the IHTSDO user community to continue to evolve. The flexible approach to Reference Set specification is thus a significant advantage as it means that new requirements can be addressed as they arise without fundamental redesign.

Advantages for implementers include standard internally consistent support for various types of customization using subsets, prioritization, display order, language preferences and alternative hierarchies using a common mechanism. The consistent versioning model ensures that implementer have opportunities to develop and utilize effective version sensitive applications. The extension of the mechanism to cover mapping to/from SNOMED CT, and built-in future flexibility, add value to implementing and optimizing Reference Set enabled software.

4.5 Concept enumerations

4.5.1 Description

The RF1 release format used documented enumerated value sets in its definitions. As an example, the concept status field identified 9 different concept status values (0,1,2,3,4,5,6,10 and 11) each of which had a name and functional description that could only be discovered by reading the accompanying documentation.

In contrast to SNOMED CT itself, this approach did not follow best practice for coding a terminology.

- It lacked unique identifiers (1 in the *conceptStatus* field means "Retired" while 1 in the *isPrimitive* field means "Primitive").
- It was language specific – translations of the descriptions and definitions can only be accomplished by translating the reference document.
- It contained some implicit relationships between values that were noted in the text but not represented formally.

RF2 addresses these issues by encoding all of these attributes as SNOMED CT concepts that are defined in the newly created *SNOMED CT model component* module.

4.5.2 Benefits

The transformation of these fields into SNOMED CT concepts derives several benefits. The same tools that are used to search for clinical concepts can be used to search for metadata concepts. Metadata concept descriptions and definitions can be translated and distributed in multiple languages leveraging already existing mechanisms. Implicit relationships among various values can be made explicit. Where appropriate, metadata value sets can be extended using the same mechanisms that are used to extend SNOMED CT itself.

4.6 Consistent representation of terms and other text strings

4.6.1 Description

In RF1 terms and other text string associated with concepts appeared in three different files. The concept file contained the fully specified name; the description file contained synonyms, preferred term and another copy of the fully specified name (including where available translations of these terms); finally a "supplementary" textual definition file contained longer, more detailed definitions of some concepts.

In RF2 all text strings associated with concept are represented in using the description file structure. For backward compatibility, RF2 retains the current 255 character length restriction on terms for fully specified names, and synonyms (including preferred terms). However, textual definitions have a different description type and are allowed to exceed this length. The specification of the description file allows for very long terms for future description types (up to 32kb) and supports the inclusion of markup in terms (e.g. HTML, DITA, etc.).

4.6.2 Benefits

Greater consistency in representation of text associated with concepts enables more integrated development, maintenance and distribution of SNOMED CT content. In addition, these changes make it possible to use a single common mechanism for supporting translation and setting language and dialect preferences. The added flexibility offered by support for longer term strings and use of markup will ensure that this consistent approach can also be applied to meet a range of anticipated future requirements.

Implementers will benefit from a single representation of text allowing presentation of different terms to be managed by reusable software functions. Similarly, the extension of this common approach to language preferences will also facilitate a single solution for language specific configuration.

5 When should my organization migrate to RF2?

5.1 Introduction

The previous section sets out the main benefits of migration to RF2. While accepting the validity of these benefits, some organizations may feel that RF1 meets their current requirements. As a result, they may choose to delay updating their software and processes to RF2. The IHTSDO provides a conversion utility and a set of compatibility files, which can be used to generate RF1 files from RF2 files. This makes it possible to continue to support RF1 during a transitional period but this cannot be regarded as a long term solution. This section sets out several practical reasons why all organizations implementing SNOMED CT and/or maintaining Extensions to SNOMED CT should make the change to RF2 as soon as possible.

5.2 Case for early RF2 migration – IHTSDO

The RF2 to RF1 conversion process is a short-term expedient that allows asynchronous migration of different Extensions and different implementations. However, there are significant costs for the IHTSDO in maintaining two different formats even with the conversion utility. These costs arise from two sources:

- The need to maintain the compatibility package used by the conversion tool. This is required to support identifier consistency between versions and to retain metadata appropriate to RF1 Subsets and Cross Maps.
- The need to maintain documentation, training materials and implementation support to deal details that differ between RF1 and RF2. The differences between the versions lead to confusions which will inevitably persist while the two release versions are both supported.

The best that the conversion utility can achieve is to provide RF1 style functionality. It can only provide this across release versions by using additional compatibility files which need to be maintained and quality assured with extensive testing to ensure the consistency of the RF1 files.

Experience in many industries illustrates that setting clear timeline for deprecation and withdrawal of support is essential to avoid a growing burden of multi-version maintenance coupled with holding back realization of the benefits of enhancements.

In the case of code systems, one argument for retaining backward compatibility relates to supporting existing stored data. However, the migration to RF2 does not change the codes used in records and thus this factor does not apply to this change. On the contrary, by supporting a full historical representation from the first release of SNOMED CT to the most recent release, RF2 provides a more effective way to understand historical record entries in the context of the time when they were created.

If one of the Extensions used by an implementer is only available in RF1 format this represents a major disincentive for migration, as this may prevent them from taking advantage of the documented benefits of RF2. For this reason IHTSDO Members that produce and maintain National Extensions should migrate at the earliest possible opportunity.

Specification changes, such as the recent provisions for components to move to the International Release from an Extension without changing identifiers, are only supported by RF2. Some Members

are keen to use these facilities and will do so as soon as they have adopted RF2. However, this will raise provenance traceability issues for RF1 users which cannot be resolved without reference to the relevant RF2 files.

These factors create a strong case for encouraging all Extension providers to migrate to RF2 at the earliest possible date. This will allow implementers to migrate and the community as a whole to gain the benefits of using RF2. As a result, in consultation with its Members the IHTSDO should set the earliest date for deprecation of RF1 as a base release format for all SNOMED CT Extensions.

5.3 Case for early RF2 migration – for providers of Extensions

Extension providers (including National Release Centers) can only take advantage of the new features described in the previous section when they have migrated to RF2. In particular, National Release Centers cannot take full advantage or the IHTSDO workbench and associated tools unless they are producing their Extension content in the RF2 format.

Until all Extension providers deliver RF2 release files, implementers that depend on their Extensions will be unable to take advantage of the enhancement in RF2 features, as they will need to continue to support RF1 and may end up offering user the lowest common denominator.

Extension developers may consider that, provided the International Release RF2 files are converted to the RF1 format, they can expect their implementers to be happy to consume their Extension in RF1. However, if an implementer is required to import an RF1 release this represents a major disincentive for migration since they will be unable to take advantage of the documented benefits of RF2. In addition, the RF2 to RF1 conversion process requires maintenance of compatibility files which add a maintenance burden to the IHTSDO and other Extension providers.

Maintaining multiple approaches during a transitional period is entirely appropriate. However, if transitional periods are stretched beyond the reasonable time required for software development a state of inertia can develop in which:

- Work-rounds are developed and maintained in RF1 for issues that are more completely resolved by RF2;
- Interdependent stakeholders in implementation are held back or demotivated by one another. For example, software vendors see no advantages in supporting RF2 while Extension developers only provide RF1 files and Extensions developer argue there is no demand because some implementers don't support RF2.

Within the IHTSDO community there is currently a now extensive knowledge of how to accomplish the RF1 to RF2 conversion. This knowledge includes an understanding of how to address issues arising during the conversion process RF2. The majority of these issues arise from errors in previous releases which are "hidden" by RF1 due to the lack of a fully history mechanisms but become visible when converted to RF2. The best time to migrate is while knowledge of these issues and the ways to manage them in a consistent way is available. The best time to migrate is before more errors hidden and perpetuated in RF1 add to future conversion.

RF2 was designed to make it easier content and derivatives developed by the IHTSDO community to be absorbed and integrated with the common resource the SNOMED CT International Release. Its intention is to make both the creation and the adoption of resources developed by the community of Members and Affiliates to be less expensive by supporting effective version tracking across a range of extensible purpose built file structures. In addition, RF2 allows separate modules and dependencies between modules to be recognized allowing effective integrations of components from different

Extensions. For organizations that are already authoring their own content today and are using their own tools, the up-front cost of revising these tools may appear high. However, the cost of bespoke solutions for issues not addressed by RF1 will be substantial and delays in migration to RF2 will increase the cost of transitional arrangement for supporting these solutions.

5.4 Case for early RF2 migration – for SNOMED CT implementers

RF2 provides valuable additions to SNOMED CT that cannot be adequately represented in RF1:

   a) A full lifetime versioning audit trail of all SNOMED CT components and standards-based derivatives;
   b) A consistent representation of metadata;
   c) An extensible structure that supports yesterday's, today's and tomorrow's requirements for adding information to SNOMED CT components, linkages between SNOMED components and linkages between SNOMED CT components and codes in other code systems, classifications and terminologies.

Organizations that are currently using RF1 subset mechanisms may already have a significant investment in tooling that meets many of their current requirements. The utility for converting RF2 release files to RF1 may appear to offer a way to continue to use RF1 format files in their applications. However, in reality this is only a temporary stopgap for a transitional period. It will not convert the beneficial features of RF2 into an RF1 representation. Users of RF1 files derived from an RF2 release only have access to content that can be represented in the RF1 format. They do not have access to:

   * The full historical versioning data provided by RF2;
   * Flexible approaches to updating using Delta release files;
   * Cross maps that use the "Complex Maps Reference Set" pattern, including the ICD-10 maps (developed by IHTSDO) and ICD-10-CM maps (developed by NLM in the US based in the ICD-10 maps.
   * Other extended derivatives that go beyond the capabilities of the RF1 "Subset Mechanism" (including the planned artifacts linking SNOMED CT and LOINC).

IHTSDO developed the RF2 distribution format to correct shortcomings in the original RF1 format, many of which were identified by implementers. Put simply, RF1 cannot support the evolving needs of the SNOMED CT implementer community. The future development of SNOMED CT will increasingly use features only available in RF2. The IHTSDO continues to support a transformation from RF2 to RF1, but transformation inevitably loses functionality supported by RF2 and the functionality gap will continue to grow in subsequent releases.

As future releases use more of the features RF2, the list of gaps in an RF1 implementation will continue to grow. It is likely that some implementers who do not implement RF2 will invest more effort in proprietary workarounds for these limitations than is required to migrate to the RF2 standard. The sooner an implementer has completed the transition the sooner implementers can benefit from the valuable enhancements that RF2 brings to SNOMED CT.

6 Support for migration to RF2

6.1 Support for migration to RF2 – SNOMED CT Implementers

The most important point for an implementer to note is that migration from RF1 to RF2 does not affect instances of records or documents in which SNOMED CT codes or expressions are stored. The logical content of SNOMED CT and the identifiers used to refer to that content are unchanged.

For an applications that accesses SNOMED CT via a terminology server, provided and maintained by a third party, the only prerequisite for migration is for the terminology service provider to support RF2. Thereafter, additional services supported by RF2 (as delivered by that terminology server) may be where they offer benefits to users of the implemented application.

Providers of terminology servers and other application that directly use or import the SNOMED CT release files need to adapt their release file import routines. They are also likely to need to revise or extend the internal data structures and indexing strategies they apply to the imported data. The IHTSDO provides advice on these issues in the Technical Implementation Guide. In addition, an illustrative example using a MySQL database into which an RF2 Full release can be imported is available to Members and Affiliate Licensees to demonstrate some implementation techniques.

6.2 Support for migration to RF2 – Extension Providers

The initial RF2 file content for each Extension must be generated from the full set of all the RF1 files ever released for that Extension. This process can be carried out and supported by tools originally used to generate the RF2 International Release.

Experience of conversions done in the past indicates that previously unrecognized errors in the historical RF1 data may be revealed during this process. These errors may disrupt conversion requiring some manual conversion. The team of people involved in earlier conversions can provide advice on this process based on experience of the best way to resolve these inconsistencies.

Once the conversion has been done and the initial RF2 baseline has been released for each Extension it is important to ensure there is no need for further manual intervention. A process that requires conversion of a complete set of RF1 files into RF2 for each release carries a high risk of introducing new inconsistencies as a result of manual exception handling. Therefore, after the initial baseline RF2 release, RF2 aware release processes and tools should be used for maintenance and subsequent releases. The IHTSDO Workbench provides an appropriate set of RF2 maintenance and release tools to support this process.