

DICOM Correction Proposal

STATUS	Letter Ballot
Date of Last Update	2025-01-19
Person Assigned	David Clunie <dclunie@dclunie.com>
Submitter Name	Jörg Riesmeier <dicom@jriesmeier.com>
Submission Date	2024-04-24

Correction Number	CP-2421
Log Summary: Fix inconsistent spelling and use of terms in Part 4	
Name of Standard PS3.4	
<p>Rationale for Correction:</p> <p>Various terms are used in PS3.4 in an inconsistent manner. This CP tries to fix some of them, e.g., there are no “Query Keys”, but “Key Attributes”, “Matching Keys” and “Return Keys”. Terms like “query level” and “query request” should also be written in a consistent manner throughout Part 4.</p> <p>Furthermore, the names of the UPS Service Class and the underlying UPS Information Model are rather inconsistent throughout Part 4. It is proposed to fix this in order to avoid any confusion.</p> <p><i>Editorial Changes:</i></p> <ul style="list-style-type: none"> • “Key Attribute” should always be written with capital initial letters. • “Matching Key Attribute” and the short variant “Matching Key” should always be written with capital initial letters. <i>Please be warned that PS3.4 Annex C also uses the variant “matching Key Attributes”, because “Matching Key Attributes” are not defined in this Chapter, but only on Annex K (and later).</i> • “Return Key Attribute” and the short variant “Return Key” should always be written with capital initial letters. • The terms “response” and “request” when used in combination with DIMSE messages, e.g. “C-FIND request”, should be written in lower case. This also seems to be the most commonly used spelling in Part 4, but there are some variations. 	
Correction Wording:	

For reference, PS3.4 Section C.4.1.3.2.2 (unchanged, see 2nd item under “Editorial Changes”)

C.4.1.3.2.2 Relational Search Method

A query using the relational method may contain any combination of keys at any level in the hierarchy. Starting at the top level in the Query/Retrieve Information Model, continuing until the Query/Retrieve level specified in the C-FIND request is reached, the following procedures are used to generate matches:

- a. The key match strings contained in the Identifier of the C-FIND request are matched against the values of the Key Attributes for each entity at the current level.
- b. If no Key Attribute is specified at the current level and the current level is not the level specified in the C-FIND request, the match shall be performed as if a wild card were specified for the Unique Key Attribute for the current level (i.e., all entities at the current level shall match).
- c. If the current level is the level specified in the C-FIND request, then for each matching entity (a matching entity is one for which the Attributes match all of the specified match strings in the Key Attributes), construct an Identifier. This Identifier shall contain all of the Attributes generated by this procedure at higher levels on this recursion path and all of the values of the Key Attributes for this entity that match those in the C-FIND request.

- d. Otherwise, if the current level is not the level specified in the C-FIND request, then for each matching entity construct a list of Attributes containing all of the **matching Key Attributes** and all Attributes that were prepared at the previous level for this entity. Then perform this procedure at the next level down in the hierarchy for each matching entity.
- e. Otherwise, if there are no matches, return a response with status equal to Success and no Identifier.

Note

1. The above description specifies a recursive procedure. It may recur upon itself multiple times as it goes down the hierarchical levels, and at each level, it may recur multiple times (one for each matching entity). This may result in a large number of Identifiers being generated.
2. It is not required that the above defined procedure be used to generate matches. It is expected that implementations will incorporate different algorithms for performing searches of the databases. For a given query, the set of matches shall be equivalent to that which would be generated by the above procedure.

<p><i>Change PS3.4 Section C.6.4.1.1 and C.6.4.1.2</i></p>
--

C.6.4.1.1 Record Key

Record Key (0008,041B) is defined at the Study, Series, and Instance query levels. It is an implementation-specific unique identifier within the level of the entity record in the SCP. The SCP of the Repository Query SOP Class shall return non-zero length values of Record Key (0008,041B). The content of Record Key (0008,041B) is opaque to applications other than the SCP.

The SCP shall construct the Record Key (0008,041B) value such that for each value the SCP can determine its order with respect to all other such values. C-FIND Response Identifiers shall be returned in the ordering of Record Key (0008,041B) values. The SCP shall be able to determine from a given value the next entity record to be returned that matches the given Query Request Identifier, without repeating any records.

Record Key (0008,041B) values are used as the Prior Record Key (0008,041C) value in a subsequent ~~Query Request~~ **Query Request** (see Section C.6.4.5.3). The SCP may establish implementation specific conditions after which a Record Key (0008,041B) value is not valid, i.e., will no longer allow continuation of a sequence of Query operations. The SCP shall be able to determine from a given Prior Record Key (0008,041C) value whether that value is still valid for determining the next record to be returned.

Note

1. The structure, content, and ordering method of Record Key (0008,041B) values is SCP implementation-specific, and is opaque to the SCU, i.e., the SCU should not attempt to parse those values for components or semantics. Values may be permanent, or may be constructed dynamically during query processing. Only the SCP can determine from the value of one Record Key (0008,041B) what would be the next appropriate record to return. For example, an SCP may use encrypted representations of an internal database primary key as the Record Key (0008,041B), and such may appear to the SCU to be random unordered values.
2. The intention of the ordering and use requirements for the Record Key (0008,041B) is to allow an SCU to obtain the complete inventory matching the Key Attributes in a sequence of Queries. See Section YYYY.2.2 "Record Key and Continuation" in PS3.17.

C.6.4.1.2 Removed From Operational Use

The Removed from Operational Use (0008,0405) Attribute is defined at the Study, Series, and Instance query levels.

Enumerated Values:

- Y** The instances are not to be used for operational purposes.
N The instances may be used for operational purposes.

A value of Y indicates the Study, Series, or Instance has been removed from operational use related to patient care, although it may be retained in the repository system for other reasons (e.g., for audit of patient radiation exposure). At the Study and Series level, the Attribute indicates whether the entire Study or Series has been removed from operational use. A value of Y at the Study level supersedes any value specified for subsidiary Series and Instances, and a value of Y at the Series level supersedes any value specified for subsidiary Instances.

While defined at the Study, Series, and Instance levels, an SCP might not support this Attribute at some, or any, of those levels. E.g., an SCP may only manage this Attribute at the Instance level, and is not required to infer a value for the Series or Study level.

Note

1. The meaning of "operational use related to patient care" is implementation or site specific, but generally includes diagnostic, clinical, and therapeutic uses, as well as administrative uses necessary for providing care (e.g., insurance authorization).
2. Studies, Series, or Instances might be marked removed from operational use by actions associated with the processing of specific Key Object Selection Document SOP Instances, e.g., in accordance with [IHE RAD TF-1] Image Object Change Management Integration Profile (IOCM). Those Key Object Selection Document SOP Instances, and their Series, may themselves be marked as removed from operational use. The Context Group for Reason for Removal Code Sequence (0008,0406) includes the Key Object Selection Concept Codes specified in IOCM.
3. The semantics of the Removed from Operational Use (0008,0405) Attribute allows the SCP to include such entities in the Repository Query response without constraint. An SCP might exclude entities marked as removed from operational use from the C-FIND ~~R~~esponses of other Query/Retrieve SOP Classes (e.g., see [IHE RAD TF-2] Section 4.66.4.1.3.1 "Access to Rejected Instances").
4. Removed from Operational Use (0008,0405) is independent of Instance Availability (0008,0056). A composite instance may have been removed from operational use but is still accessible at the rapidity specified by Instance Availability (0008,0056). Conversely, an instance may not have been removed from operational use but is UNAVAILABLE for retrieval.
5. If the SCP retains records of deleted Studies, Series, or Instances, even though the actual ~~I~~nstances are physically deleted, it may include those entities in the C-FIND ~~R~~esponse with an appropriate Reason for Removal Code Sequence (0008,0406) value. Such instances may have an Instance Availability (0008,0056) value "UNAVAILABLE" (see Section C.4.1.1.3.2).

Change PS3.4 Section C.6.4.1.4

C.6.4.1.4 Metadata Sequence and Updated Metadata Sequence

An SCP may manage a set of metadata Attributes of the SOP Instances in the repository for response to ~~Q~~query requests. Metadata Sequence (0008,041D) in a Response Identifier shall contain all SOP Instance Attributes at the ~~Q~~query level that are managed by the SCP, excluding bulk data elements (such as pixel, waveform, and surface mesh data) and non-SOP Instance Attributes specified in Section C.3.4 or in Table C.6.4.1-1.

Note

The set of Attributes managed by the SCP is implementation dependent. In some implementations the managed set of Attributes might include only those few required to be supported for ~~Query Key~~-matching, while in other implementations the set might include every non-bulk data Attribute. See Section C.6.2.1 in PS3.4.

An SCP may manage a set of metadata Attributes whose "updated" values differ from those in a stored SOP Instance accessible through a non-DICOM protocol specified in the File Access URI (0008,0409) or Folder Access URI (0008,0408). Although a stored SOP Instance shall be conformant to its IOD (per the requirements of the DICOM File Format), some Attributes in the file might not have current values (e.g., Patient's Name (0010,0010) may have been corrected or changed after the SOP Instance was stored). Updated Metadata Sequence (0008,041E) in a Response Identifier shall contain all Attributes at the ~~Q~~query level whose values are different from the values contained in the stored SOP Instance file.

An SCP that supports non-DICOM protocol URI references to stored SOP Instances shall support either the Metadata Sequence (0008,041D) or the Updated Metadata Sequence (0008,041E), or both, to provide current metadata values for SOP Instances accessed through the non-DICOM protocol.

Note

1. SOP Instances accessed through DICOM protocols are expected to have current values in all Attributes.
2. The SCP might not track whether Attribute values have changed, or which specific Attributes have changed values, and would therefore not support Updated Metadata Sequence (0008,041E). In this case, the SCU may request the Metadata Sequence (0008,041D) that contains all current Attribute values managed by the SCP, whether or not they have been updated. Determination of differences, if any,

between those returned Attribute values and values in the stored SOP Instance would be the responsibility of the SCU.

At any **Q**query level, Metadata Sequence (0008,041D) or Updated Metadata Sequence (0008,041E) may include the Original Attributes Sequence (0400,0561) describing the provenance of changes to Attributes at that level or at higher **Q**query levels.

If Metadata Sequence (0008,041D) and/or Updated Metadata Sequence (0008,041E) are present in a Request Identifier, their absence in the Response Identifier indicates they are not supported by the SCP (see Section C.2.2.1.3).

A zero-length value or a single empty Item in Updated Metadata Sequence (0008,041E) in a Response Identifier indicates support by the SCP, but that there are no differing metadata Attribute values.

Change PS3.4 Section V.2.2

V.2.2 Attributes Definition

Attributes are defined for each entity in the internal Entity-Relationship Model. An Identifier in a C-FIND request shall contain values to be matched against the Attributes of the Entities in a Substance Administration Query Information Model. For any **Q**query request, the set of entities for which Attributes are returned shall be determined by the set of Matching and Return Key Attributes specified in the Identifier.

Change PS3.4 Annex CC, in particular Section CC.1

CC Unified Procedure Step Service and SOP Classes (Normative)

CC.1 Overview

This Annex defines the Service and SOP Classes associated with a Unified Worklist and Procedure Step.

The Unified Procedure Step Service Class provides for management of simple worklists, including creating new worklist items, querying the worklist, and communicating progress and results.

A worklist is a list of Unified Procedure Step (UPS) instances. Each UPS instance unifies the worklist details for a single requested procedure step together with the result details of the corresponding performed procedure step. There is a one to one relationship between the procedure step request and the procedure step performed.

Unified Procedure Step instances may be used to represent a variety of scheduled tasks such as: Image Processing, Quality Control, Computer Aided Detection, Interpretation, Transcription, Report Verification, or Printing.

The UPS instance can contain details of the requested task such as when it is scheduled to be performed or Workitem Codes describing the requested actions. The UPS may also contain details of the input information the performer needs to do the task and the output the performer produced, such as: Current Images, Prior Images, Reports, Films, Presentation States, or Audio recordings.

The Unified ~~Worklist and~~ Procedure Step Service Class includes five SOP Classes associated with UPS instances. The SOP Class UID for any UPS Instance always specifies the UPS Push SOP Class. The separate SOP Classes facilitate better negotiation and logical implementation groups of functionalities.

The UPS Push SOP Class allows an SCU to instruct the SCP to create a new UPS instance, effectively letting a system push a new work item onto the SCP's worklist. It is important to note that the SCP could be a Worklist Manager that maintains the worklist for other systems that will perform the work, or the SCP could be a performing system itself that manages an internal worklist.

The UPS Pull SOP Class allows an SCU to query a Worklist Manager (the SCP) for matching UPS instances, and instruct the SCP to update the status and contents of selected items (UPS instances). The SCU effectively pulls work instructions from the worklist. As work progresses, the SCU records details of the activities performed and the results created in the UPS instance.

The UPS Query SOP Class allows an SCU to query a Worklist Manager (the SCP), but does not otherwise interact with the UPS instances.

The UPS Watch SOP Class allows an SCU to subscribe for status update events and retrieve the details of work items (UPS instances) managed by the SCP.

The UPS Event SOP Class allows an SCP to provide the actual status update events for work items it manages to relevant (i.e., subscribed) SCUs.

Each of these services has an equivalent HTTP operation defined by the UPS-RS Worklist Service (see Section 11 “Worklist Service and Resources” in PS3.18).

While ~~an SCP of the Unified Worklist and Procedure Step Service Class~~ ~~SCP~~ is not required to support UPS-RS, an SCP may choose to support one or more of the UPS-RS services as an origin server. In this scenario, an SCP/origin server shall follow the same internal behavior for all Workitems irrespective of whether they originated with a DIMSE request or an HTTP request. A DIMSE request and its equivalent HTTP request with the same parameters shall yield the same response.

For example:

- A Workitem instance created via DIMSE N-CREATE can be retrieved via HTTP requests and vice-versa
- A Workitem instance created via DIMSE N-CREATE can be updated, have its state changed or be canceled via HTTP requests and vice-versa
- A C-FIND request and an HTTP SearchForUPS request with the same parameters shall return the same set of results
- An N-EVENT-REPORT SCU that also supports HTTP subscriptions will record whether a given subscriber uses DIMSE or WebSockets and send the appropriate form of notification to that subscriber
- A change made to a Workitem instance will result in the same event notifications regardless of whether the change was requested via DIMSE or HTTP
- A Global Subscription request or a Filtered Global Subscription request will subscribe an SCU (or user agent) to instances created both via DIMSE and via HTTP requests
- A DIMSE event subscriber will receive notifications for relevant changes made via HTTP requests
- An HTTP event subscriber will receive notifications for relevant changes made via DIMSE requests

The mapping between UPS DIMSE operations and UPS-RS services is defined in Section 11 “Worklist Service and Resources” in PS3.18.

Change PS3.4 Section CC.2.8.1.1

CC.2.8.1.1 E/R Model

In response to a given C-FIND request, the SCP might send several C-FIND responses, (i.e., one C-FIND response per matching worklist item). Each worklist item describes a single task and its related information.

The Unified Procedure Step ~~Query~~ Information Model is represented by the Entity Relationship diagram shown in Figure CC.2.8-1.



Figure CC.2.8-1. Unified Procedure Step E-R Diagram

There is only one Information Entity in the model, which is the Unified Procedure Step. The Attributes of a Unified Procedure Step can be found in Table CC.2.5-3.

Change PS3.4 Section CC.2.8.1.3.1

CC.2.8.1.3.1 Request Identifier Structure

An Identifier in a C-FIND request shall contain:

- Key Attributes values to be matched against the values of Attributes specified in the SOP Class identified by the Affected SOP Class UID.
- Conditionally, the Attribute Specific Character Set (0008,0005). This Attribute shall be included if expanded or replacement character sets may be used in any of the Attributes in the Request Identifier. It shall not be included otherwise.
- Conditionally, the Attribute Timezone Offset From UTC (0008,0201). This Attribute shall be included if Key Attributes of time are to be interpreted explicitly in the designated local time zone. It shall not be present otherwise, i.e., it shall not be sent with a zero-length value.

Note

This means that Specific Character Set (0008,0005) is included if the SCU supports expanded or replacement character sets in the context of this service. It will not be included if expanded or replacement character sets are not supported by the SCU.

The Key Attributes and values allowable for the query shall be defined in the SOP Class definition corresponding to the Affected SOP Class UID for the corresponding Unified ~~Worklist And~~ Procedure Step Information Model.

Change PS3.4 Section CC.2.8.3.1

CC.2.8.3.1 Worklist Search Method

The following steps are used to generate match responses.

- Match the Matching k~~Key~~ ~~match~~ Attributes contained in the Identifier of the C-FIND request against the values of the ~~Key~~ Attributes for each worklist entity.
- If there are no ~~matching keys, then there are~~ no matches, return a response with a status equal to Success and with no Identifier.
- Otherwise,
 - For each entity for which the Attributes match all of the specified ~~m~~Matching k~~Key~~ Attributes, construct an Identifier. This Identifier shall contain all of the values of the Attributes for this entity that correspond to the Matching Key ~~and r~~Return k~~Key Attributes~~ specified in the C-FIND request.
 - Return a response for each remaining Identifier.

Table CC.2.5-3 defines the Attributes of the Unified Procedure Step Information Model, the requirements for ~~key~~ ~~m~~Matching Keys, and the requirements for ~~r~~Return k~~Keys~~.

Change PS3.4 Section CC.3.1

CC.3.1 Service Class and SOP Class UIDs

All UPS ~~i~~nstances shall be created with the value of SOP Class UID set to "1.2.840.10008.5.1.4.34.6.1" (i.e., that of the UPS Push SOP Class).

Note

UPS Instances are all based on the Unified Procedure Step IOD and are all created either internally by the SCP, or in response to an N-CREATE request issued as part of the UPS Push SOP Class.

Once created, UPS instances may be operated on by DIMSE services from any of the UPS SOP Classes defined in the Unified ~~Worklist and~~ Procedure Step Service Class.

During association negotiation, the Abstract Syntax UID shall be the implemented SOP Class as shown in the following list:

- 1.2.840.10008.5.1.4.34.6.1 (UPS Push SOP Class)

- 1.2.840.10008.5.1.4.34.6.2 (UPS Watch SOP Class)
- 1.2.840.10008.5.1.4.34.6.3 (UPS Pull SOP Class)
- 1.2.840.10008.5.1.4.34.6.4 (UPS Event SOP Class)
- 1.2.840.10008.5.1.4.34.6.5 (UPS Query SOP Class)

CC.3.1.1 DIMSE Implications for UPS (Informative)

A SOP Instance may be created with one SOP Class UID (UPS Push) and later DIMSE Services may refer to it over an association negotiated for a different SOP Class UID. Further details on this can be found in Section 10 "DIMSE-N" in PS3.7.

For DIMSE-N Services, the Affected SOP Class UID (0000,0002) or Requested SOP Class UID (0000,0003), when present, will be the UID of the UPS Push SOP Class regardless of the negotiated Abstract Syntax UID. The SCU and SCP will not reject DIMSE-N messages on the basis of the Affected/Requested SOP Class UID being that of the UPS Push SOP Class, rather than one of the other four SOP Class UIDs as listed in the Abstract Syntax UID during association negotiation. The SCU and SCP may reject the DIMSE-N messages if the instance is not a UPS Push SOP Class Instance.

For DIMSE-C Services (C-FIND), the Affected SOP Class UID will always match the negotiated Abstract Syntax UID for the Presentation Context under which the request is made. This will be UPS Watch, UPS Pull or UPS Query. All of these SOP Classes represent the Unified Procedure Step Information Model described in Section CC.2.8.1.

For example, in a typical "Pull Workflow" message exchange, the C-FIND query from a "performing SCU" would use the UPS Pull SOP Class UID for both the negotiated Abstract Syntax UID and the Affected SOP Class UID (0000,0002), however the SOP Class UID (0008,0016) of the C-FIND responses themselves will be set to the UPS Push SOP Class UID by the SCP. All the subsequent N-ACTION, N-SET, and N-GET messages, would then use the UPS Pull SOP Class UID for the negotiated Abstract Syntax UID, and the UPS Push SOP Class UID for the Affected SOP Class UID (0000,0002).

<i>Change PS3.4 Section DD.3.2.4.3</i>
--

DD.3.2.4.3 Behavior

The SCU uses ~~the~~ N-DELETE to request the SCP to delete an applicable Machine Verification SOP Instance. The SCU shall specify in the N-DELETE request primitive the SOP Instance UID of the applicable Machine Verification instance.

The SCP shall delete the specified SOP Instance, such that subsequent operations of the same SOP Instance will fail.

The SCP shall return the Status Code of the requested SOP Instance deletion. The meanings of success, warning, and failure status classes are defined in Annex C "Status Type Encoding (Normative)" in PS3.7.

If an N-DELETE is not issued, the SOP ~~Class~~ Instance may be deleted on the SCP by a manual or automatic operation. This behavior is beyond the scope of the Standard.