Digital Imaging and Communications in Medicine (DICOM)

Supplement 223: Inventory IOD and Related Services

Prepared by:

DICOM Standards Committee, Working Group 33 (Data Archive and Management)

1300 N. 17th Street, Suite 900

Rosslyn, Virginia 22209 USA

VERSION: For Public Comment (draft 25)

Developed in accordance with: DICOM Work Item 2020-08-B

Table of Contents

240

Document History

Document Version	Date	Content
00	01-Sep-2020	Initial Draft
01	14-Sep-2020	Initial discussion of file access
02	30-Sep-2020	
03	13-Oct-2020	Approach without Part 10/11/12
04	16-Oct-2020	Complete for WG33 review
05	26-Oct-2020	Hyperlinks added and verified
06	27-Oct-2020	Added Study Update Datetime
07	01-Nov-2020	Added MAC, Completion status, filename in container
08 – FR	09-Nov-2020	Added Inventory to IOD Summary Table, Directory record
09	18-Nov-2020	Sup223. Added Equipment IE, explanation on Imaging Service Request, initial Part 17 outline
10	23-Nov-2020	Scope includes Empty Study Date and Patient ID; full draft of Part 17
11	03-Dec-2020	Scope includes Empty Accession Number; additional sections for Part 17
12	08-Dec-2020	Compression of inventory SOP Instances
13	09-Dec-2020	Clarifications in Part 17
14	01-Jan-2021	Added N-ACTION Resume request, more Part 17
15	07-Jan-2021	Minor edits; change group number for new data elements
16	12-Jan-2021	Clarify timepoint matching
17	14-Jan-2021	Change "Not For Clinical Use" to "Removed from Operational Use"; changed VRs of Study UID List and Patient ID List
18	19-Jan-2021	Clarify TAR spec, metadata updates recorded in Inventory, MAC calculation, URI format
19	04-Feb-2021	Allow multiple URIs, base URI at Series level, base URI optional; add Original Attributes at Series and Instance levels; add Instance Availability
20	01-Mar-2021	Tree of inventory objects, file access macro; Part 17 on inventory tree
21	12-Mar-2021	Part 17 on inventory tree and reconciliation, miscellaneous editorial fixes
22	17-Mar-2021	Miscellaneous editorial fixes
23	25-Mar-2021	Change to Original Attributes Sequence moved to CP2081; Completion Status Type 1, add PARTIAL; add open issues; rename some attributes; separate Study Datetime Range and Study Update Datetime Range; allow Inventory instances with no records, remove NOMATCH status; disallow lower-level records; add match on Series Description, match on Patient ID together with Issuer
24	30-Mar-2021	Clarify Inventory Creation conformance; produced Inventory must have DICOM access method; changes to Part 17, open issues;
25 – PC	31-Mar-2021	Add N-ACTION Pause, request to change reporting interval

Scope and Field of Application

This Supplement introduces a new composite Inventory IOD, and SOP Classes to initiate creation of Inventory SOP Instances and to transfer/query/retrieve them.

A use case of steadily increasing significance is porting large DICOM repositories from one image management system (PACS or VNA) to another. Users typically replace their PACS after ~12-15 years, often with change of vendor. Replacement requires migrating historical data to the new system. Thus, every year, 5-10% of user organizations may be doing a PACS data migration.

Old images are now routinely retained "forever", and data set sizes are increasing with 3D/4D and multimodality studies. Archives in many institutions now store over a billion instances with data volumes over one petabyte. Migration approaches need to operate at large scales, and handle both on-premises and remote (e.g., cloud-based) storage.

Migration often occurs when either the source system or the destination, or both, are in clinical operation, but systems designed and configured to handle the throughput of regular operations might not have capacity for the additional massive input/output requirements of migration. With a data transfer rate of 1 terabyte / day (quite high for even the most advanced PACS), the time to transfer a petabyte archive is 3 years. Performance constraints exist on both the source and destination systems.

Similar needs arise when healthcare institutions merge previously disparate repositories into an enterprise repository - the old archives need to be migrated. This is an increasing need with the accelerated pace of healthcare organization consolidations. Conversely, large sets of archived data may sometimes need to be migrated out of a repository to support business divestment or realignment in a healthcare organization.

There are also research use cases, including artificial intelligence and machine learning, where bulk access to the archive is desirable, and such uses might leverage some of the same mechanisms developed for migration. PACS audit and quality control may also utilize some of the standardized functionality developed for migration, such as an archive inventory and metadata to identify the data produced by a particular unit or by a particular modality.

The current DICOM Standard does not address the use case and technical interoperability requirements for migration of a full enterprise repository data set, and it is currently ill-suited for the major performance issues of migration.

The current Standard is designed for routine daily department operational workflows – acquisition, storage, analysis, and reading of imaging studies associated with individual patients. The Standard is optimized for identifying and transferring the objects for, at most, just a few studies or patients at a time. Network Query/Retrieve operations are synchronous, and the network connection must remain open until the operation is complete. The number of items in a response therefore is typically restricted; the Standard explicitly defines A70x – Out of Resources error statuses for these restrictions. Even media-based data exchange is specified only for the use cases of limited file sets, basically what can fit on a DVD.

A key requirement for migration (and other use cases) is the ability to have an inventory of all studies, series, and instances in an archive. While the current Query Service (DIMSE or equivalent DICOMweb) could be used, limitations on number of responses and the synchronous protocol require the use of a possibly very large number of partial Query requests. This makes producing such an inventory difficult.

While the current Standard Retrieve and Storage Services make moving data possible, they require significant overhead for the transfer of each object. A standardized method of direct filesystem access to stored object files is needed.

This Supplement specifies an Information Object Definition capable of encoding an inventory of all studies, series, and instances in a repository. This is functionally equivalent to a Query response that returns an
290 inventory of the entire repository database (or a described subset thereof). Inventory of non-patient objects are out of scope for this Supplement.

The Supplement also defines a mechanism to initiate the production of the inventory and allow it to proceed asynchronously. The inventory SOP Instance(s) would be available when the operation is complete. Access to the inventory SOP Instances would be similar to mechanisms specified for other DICOM Non-Patient Object
295 classes (such as Color Palettes).

Since archive systems may optionally support direct filesystem access to DICOM Part 10 compliant files, for all or some of their stored instances, the Inventory IOD allows a link to such accessible files. To support a common PACS implementation design, wherein the archive may retain metadata updates (e.g., changed patient IDs) in its database and not propagate them to the stored instances, the inventory provides the current metadata for the
300 stored instances which may have not been updated.

While this Supplement defines an IOD and Services to support migration, it is not in itself a complete standard for migration. Migration involves many other processes, some of which may be supported by other standards such as HL7, and some of which are not (yet) supported by any standards.

305

Open Issues

WG-33 requests specific public comment on these open issues. Additional background information may be found in the proposed Part 17 informative annex (hyperlinks provided), and in the Sup223 educational presentation available on the https://www.dicomstandard.org/comment web site.

1. Scope of Inventory – selecting studies to be included in inventory (Part 17 link)	
1.A	**Is general key attribute matching needed for Scope of Inventory?**
	The Scope of Inventory is currently limited to only the specified key (scoping) attributes; additional attributes are not permitted (e.g., in a Standard Extended conformance). This is unlike C-FIND, which has a general mechanism for search keys. Are there use cases for which a general key attribute matching capability is needed?
	Would such a use case be better satisfied by client-side (user) filtering of a broader inventory scoped only by the currently specified key attributes?
	As DICOM formally allows the receiver to disregard such Type 3 data elements, would a receiver of an Inventory instance be able to properly interpret the Scope of Inventory and Completion Status for arbitrary sets of key attributes that it might not have in its data model?
1.B	**Is the use-specific matching specification for the Scope of Inventory warranted?**
	The matching semantics introduce another attribute matching mechanism into DICOM after the C-FIND mechanism and the Attribute Value Constraint (a.k.a. "Selector" mechanism, PS3.3 Section 10.25). The Scope of Inventory mechanism, while similar to C-FIND, is tailored for its use cases, e.g., with multi-value matching for Modalities in Study, Attributes Missing Value matching, and matching on combination of Patient ID List and Issuer of Patient ID. These types of matching are not possible with the existing C-FIND or Selector mechanisms. Is this new mechanism warranted?
1.C	**Should matching on SOP Classes in Study be specified?**
	What are the use cases for creating an inventory based on SOP Classes in the Study? Should matching on multiple values be specified, and what should be the logic relationship amongst those matching values (AND or OR)? Or is selection by Modality sufficient for the use cases, with a subsequent client-side filtering for SOP Classes?
2. Accessing data in repository/archive (Part 17 link)	
2.A	**Should Study or Series level folders or container files allow "extraneous" objects?**
	Each SOP Instance in the repository can have one or more direct file access links. Additionally, a link may be provided to a folder or container file that has all the stored object for a Study or Series, and which may include objects from other Studies/Series. Would this affect use of Study/Series containers in migrations or other use cases? Should such links be allowed only for containers without instances of other Studies/Series (but potentially with non-DICOM files, such as reports)?
2.B	**Should specific non-DICOM file access protocols be prescribed normatively?**
	Non-DICOM file access protocols are left open. Should specific protocols be prescribed normatively? If so, which ones? How should different HTTP-based cloud storage protocols be differentiated (e.g., AWS vs. Azure)?
3. Security (Part 17 link)	
3.A	**What are the use cases for a de-identification profile for Inventory instances?**
	The current (and only) de-identification profile in DICOM, the Basic Application Level Confidentiality Profile, is oriented toward classic image SOP Instances, with PHI attributes in the top-level data set. It is not clear that that profile can be effectively applied to the Inventory IOD, where all PHI is within Sequence attributes. Further, it is not clear that the use cases for the Inventory are even appropriate for applying that profile. Is there is a need for a de-identification profile for Inventory instances? Which new data elements should be added to the tables of PS3.15?

3.B	**What is effect of user authorization rules on production of an Inventory?**
	A repository might limit disclosure or retrieval of SOP instances, studies, or patients following a variety of authorization policies, often based on the user's identity and/or instance attributes. How should the Standard address the application of such rules during the production of an Inventory? E.g., depending on their permissions, different users of the Inventory Creation service might get different content for the same Scope of Inventory – is this acceptable? How should such different outcomes be identified in the Inventory SOP Instances?
3.C	**Should the Inventory include confidentiality/privacy attributes?**
	A repository might have confidentiality/privacy attributes on certain patients (e.g., VIP status). Should the Inventory IOD include standardized confidentiality/privacy attributes in the Patient IE, and if so, what attributes? Should such attributes be handled external to the DICOM IOD, e.g., in the EMR?
	4. Production of inventory (Part 17 link)
4.A	**Should Pause/Resume of inventory production have remote client control?**
	The Inventory Creation service defines a "paused" status, e.g., to support a temporary shortage of resources, or other production block during a potentially lengthy (multi-day) inventory creation process. In addition to server reasons for pausing production, N-ACTION operations are defined allowing the client initiating Inventory Creation to request pause/resume for operational reasons outside the knowledge of the server.
	Is this a necessary remote control for DICOM network operations, or should all control be exercised directly on the server (PACS) administrative interface? What are the implications for server implementation for remote Pause/Resume controls?
4.B	**Should client be able to request a "fresh" inventory?**
	A server may create an inventory by reusing (by reference) a baseline inventory and creating only the incremental inventory of studies that have been added or changed since that baseline (Part 17 link). Should a client be able to request a "fresh" inventory created from the database without reuse of existing inventories? Note that a server could refuse to accept a request for a fresh inventory based on authorization or resource constraints.
4.C	**Is it appropriate to require the producer to include the full tree of references to subsidiary Inventory instances?**
	An inventory may be divided into a set of multiple SOP Instances linked in a tree structure (Part 17 link). Each SOP Instance includes a full tree of references to the subsidiary instances in the set, including access attributes for those instances (another Part 17 link). Is requiring this complete tree an undue burden on the producer of an inventory? As access attributes may become "stale" and become broken links, is inclusion of those links appropriate? [see Closed Issue 17]
4.D	**Is there a need for an inventory production monitoring capability?**
	The client that initiates Inventory Creation receives N-EVENT-REPORT status notifications on a regular basis. Is there a need for other clients to also receive such status notifications, e.g., for an inventory initiated locally at the PACS administrative user interface? Should such a service include a client-initiated subscription capability? (Compare the Unified Procedure Step Subscribe/Unsubscribe)

310

Closed Issues

WG-33 has taken a position on each of the following issues, as noted in each item. However, public comment is welcome on them if there is a concern with the approach taken in the draft Supplement.

1	Direct filesystem access supported for what formats of target data?
	Part 10 conformant files, Pt10 files in ZIP, Pt10 files in TAR, Pt10 files in TAR+GZIP. Will consider other container file formats, e.g., TAR+BZIP2, or TAR+Z, if public comment indicates they are in use

2	Do all metadata updates need to be propagated down into stored instances accessed through filesystem?
	No, current metadata managed in Inventory; must be applied by client to retrieved objects
2a	How do we indicate that data retrieved via filesystem needs metadata updates to be applied?
	Flag considered unreliable/untrustworthy. Client must always assume that metadata in inventory needs to be applied to stored instances. Discussion added to Part 17
3	Do we need to be able to de-identify Inventory SOP Instances?
	Inventory de-identification as a service option is out of scope. Further discussion under Open Item 3.A
4	Support for different users that may desire different "views" of the metadata (e.g., different Patient IDs)
	Out of scope – may be satisfied by Other Patient IDs Sequence, or by different AE Titles providing different "views" of data. Discussion added to Part 17
5	How to handle deprecated objects (IOCM), and IOCM KO objects?
	Inventory includes SOP Instances identified as Removed from Operational Use, with specific flag attribute, optional reason code; inventory includes IOCM KO objects if included in archive (may be marked as processed)
6	Address security of "data at rest" in archive?
	New security profile out of scope. DICOM Secure File Format defined in Part 10. WG-14 may address further specifications in separate Change Proposal as needed. Discussion of security added to Part 17
7	Referencing Files, File Sets, etc. – specify in Parts 10/11/12?
	Shoehorning into existing Media Exchange is problematic; approach using two-page Annex in Part 3.
8	Do we need variable richness of metadata in inventory (e.g., additional Instance level attributes)? Can one size fit all, or do we need a few sizes, or totally variable?
	Inventory can be requested at Study, Series, or Instance level. Inventory attributes and match keys are specified for migration use case. Additional attributes, e.g., to support research inventories, are generally not supported in PACS, may require reading SOP Instances, and are out of scope, but PACS may add any additional attributes as Standard Extended conformance. Inventory can be used as a first pass to identify relevant studies based on modality and study description. Discussion added to Part 17
9	As inventory objects may become extremely large, efficiencies in production become important. Using current attribute tags for the Referenced Series and Referenced Instance Sequences, Study level attributes will be split, with some before and some after the Series and Instance records for a Study. Should we define a different attribute for structuring the records so that complete Study and Series records can be produced and encoded before moving on to the lower level records?
	New attributes specified for Inventoried Studies, Series, and Instances Sequences
10	Should a method be defined to set the Scope of Inventory to Studies with empty Study Date and Time, or with empty Patient ID?
	Yes. Allow match to empty Study Date, Accession Number, or Patient Name
11	How should non-DICOM protocol URI be specified to avoid repeating the protocol and server name for each instance?
	Base URI specified at Study and Series level that can optionally be used for all instances in that Study or Series.
12	Is it necessary for the inventory to record multiple storage locations for accessing an inventoried SOP Instance?
	Yes. Multiple URIs allowed per referenced Instance.

13	C.YY.1.2.7 and C.YY.1.3.1 specify that a folder or container file (e.g., ZIP, TAR) at the Study or Series level must contain all the SOP Instances of that Study or Series, i.e., only a single container is allowed to be identified. Is this an acceptable limitation? *Since the file location of each SOP Instance is recorded, the Study and Series level links are an optimization. There is a method to link to each object, so this approach is acceptable*
14	The Scope of Inventory allows specification of a Study Instance UID List to be included in the inventory (e.g., to support research uses where applicable studies have been identified elsewhere). The Study Instance UID List has VR UI, which limits the list to 65534 bytes (~2000 UIDs). Is this an acceptable limit, considering that a client application that needs more can simply request multiple inventories? *Limit considered acceptable*
15	For implementation specific reasons, the content of an inventory may need to be split into more than one SOP Instance. How should these multiple objects be linked to provide the complete list of objects for the Inventory? 1- Associatively by a Transaction UID 2 - In a meta-inventory of inventory objects 3 - Explicit cross-reference within the structure of the Inventory IOD *Approach 3 taken – tree of linked inventory instances. Discussion added to Part 17*
16	Inventory Creation completion N-EVENT-REPORT and Inventory Query Model C-FIND do not provide full list of Inventory SOP Instances included by reference. SCU client application would need to open root inventory object to obtain list. Is this acceptable? *Acceptable*
17	Since Inventory Creation SOP Class might not be used, hence no N-EVENT-REPORT with link to inventory objects, and Inventory C-FIND may not be implemented, how should client locate and access instances in inventory tree? 1 – leave out of scope (creating application uses some unspecified means to provide access info) 2 – (given root object) access info provided to next objects in tree, app opens each object to traverse tree 3 – (given root object) access links to all objects in tree provided in root object 3a – flat list 3b – tree structure *Approach 3b taken. Discussion added to Part 17*
18	The Patient IE in the Inventory IOD contains a basic set of patient demographics (name, sex, DoB). Would additional attributes be useful for patient matching during migration/consolidation of multiple repositories? Are those additional attributes available in PACS databases? *Additional attributes not needed*
19	Section C.YY.1.1.5 allows a study to appear more than once among all SOP Instances of a group, e.g., to support implementations that replicate a study to multiple subsystems, or that may store different series of a study on different storage devices. What are the implications for the inventory consumer if a study is listed more than once in an inventory? What if the contents of those two listings are different, either due to the study being split to two storage subsystems (each inventorying its own data), or due to the study changing during the production of the inventory (with the two listings at different timepoints)? *Migration client in general needs to perform reconciliations of inventory with other data; rather than requiring producer to perform merge, this is just another part of reconciliation for the consumer. Discussion added to Part 17*
20	The Scope of Inventory Sequence is specified with almost all attributes Type 2. Should these be specified with conditional inclusion (Type 1C), such that the data elements would be absent (rather than zero-length) if not used? *Type 2 attributes OK.*

21	Inventory explicitly identifies data elements for the physician roles: Referring, Reading, Consulting, and Physician(s) of Record. Which of these are actually managed in PACS implementations? What is the use case for any of the Physician roles in migration? Which are required?
	No use for physician roles for migration identified, but may be useful for other use cases. As all are Type 3 optional, no harm in leaving them in the specification.
22	Should Inventory instances with no Study records be allowed?
	Yes – allows affirmative declaration that there are no matches for Scope of Inventory at Content Date/Time. Added to Part 17.
23	The Storage Management Service Class, which includes the Inventory Initiation SOP Class, is designed to support additional SOP Classes for other aspects of long-term repository management. It is modeled on the Storage Commitment Service Class and SOP Class. Should the Storage Commitment SOP Class be merged into the Storage Management Service Class? (This would be only a redocumentation.)
	Not at this time.

Relationship to other DICOM proposals

This Supplement is dependent on the Final Text completion of CP2081 and CP2092.

315   Material to be added as needed for alignment with Supplement 207 Revision of DICOM Conformance Statements.

DICOM PS 3.2 Conformance

Add SOP Classes to Table A.1-2 UID Values

Table A.1-2 UID Values

UID Value	UID NAME	Category
…		
1.2.840.10008.5.1.4.xxuid.1	Inventory Storage	Transfer
1.2.840.10008.5.1.4.xxuid.2	Inventory FIND	Query/Retrieve
1.2.840.10008.5.1.4.xxuid.3	Inventory MOVE	Query/Retrieve
1.2.840.10008.5.1.4.xxuid.4	Inventory GET	Query/Retrieve
1.2.840.10008.5.1.4.xxuid.5	Inventory Creation	Workflow Management

320

DICOM PS 3.3: Information Object Definitions

> *Add ZIP, TAR, GZIP and IHE to Section 2 Normative References*

2.1 INTERNATIONAL ORGANIZATION FOR STANDARDIZATION (ISO) AND INTERNATIONAL ELECTROTECHNICAL COMMISSION (IEC)

325 [ISO/IEC 21320-1] ISO/IEC. 2015. *Information technology - Document Container File - Part 1:Core*
https://standards.iso.org/ittf/PubliclyAvailableStandards/c060101_ISO_IEC_21320-1_2015.zip

…

2.3 INTERNET ENGINEERING TASK FORCE (IETF)

[RFC1951] IETF. *DEFLATE Compressed Data Format Specification version 1.3*
330 https://tools.ietf.org/html/rfc1951

[RFC1952] IETF. *GZIP file format specification version 4.3* http://tools.ietf.org/html/rfc1952

…

2.6 OTHER REFERENCES

[IHE RAD TF-1] IHE International. *IHE Radiology (RAD) Technical Framework, Volume 1 – Integration Profiles*
335 https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_TF_Vol1.pdf

[IHE RAD TF-2] IHE International. *IHE Radiology (RAD) Technical Framework, Volume 2 – Transactions*
https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_TF_Vol2.pdf

[POSIX] IEEE and The Open Group. *POSIX.1-2017 IEEE Std 1003.1™-2017*
https://pubs.opengroup.org/onlinepubs/9699919799/

340 [ZIP] PKWARE, Inc. *ZIP File Format Specification*
http://www.pkware.com/documents/casestudies/APPNOTE.TXT

Add Inventory to Section 7.13 DICOM Model of the Real World for Non-Patient-Related Information

345 **7.13.x Inventory**

Figure 7.13.x-1 shows the E-R diagram for the Inventory information model. The Inventory Information Entity provides an inventory of Studies, and their component Series and SOP Instances, managed by an information system (such as a Picture Archiving and Communication System – PACS). The Inventory Information Model includes contextual information about each Study through the Patient and Imaging Service Request IEs. It
350 includes information on the stored SOP Instances, including access mechanisms supported by the system.

Notes

1. This information model is similar to the Study Root Query/Retrieve Information Model (see PS3.4 Section C.6.2.1).

2. There is a potentially complex relationship between the Study and Imaging Service Requests in the real world (e.g., see [IHE RAD TF-2] *Section 4.6.4.1.2.3 Relationship between Scheduled and Performed Procedure Steps*).
355 However, the Inventory information model follows the basic Study information model and supports only a single Accession Number representing an Imaging Service Request (see Section C.7.2.1). Note that IHE specifies that if a Study has multiple associated Imaging Service Requests, the request attributes are encoded at the Series level.
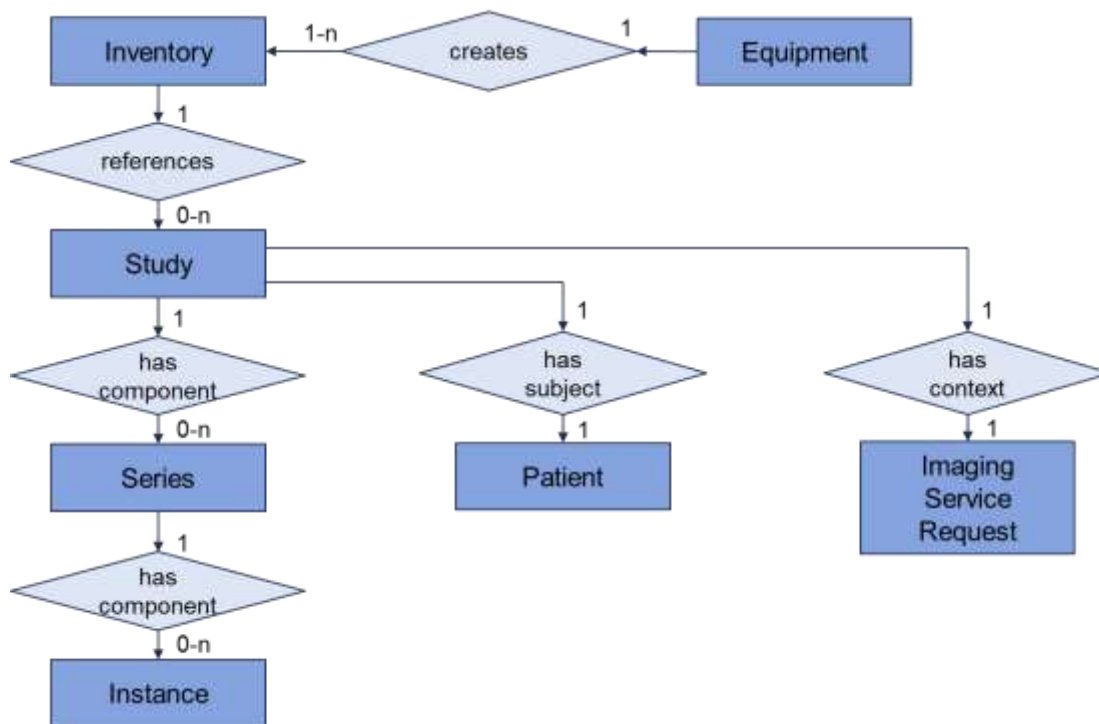
360 **Figure 7.13.x-1. Inventory Information Model E-R Diagram**

Add Inventory IOD to Section A.1.4 summary table

---

365 *– DICOM editor to select appropriate table (create new table for Non-Patient Objects with Color Palette and Hanging Protocol from Table A.1-3?)*

A.1.4 Overview of the Composite IOD Module Content

…

Table A.1-x. Composite Information Object Modules Overview -

IODs Module	Inventory
Inventory	M
Equipment	M
SOP Common	M

370

Add new section for Inventory IOD to Annex A Composite Information Object Definitions

---

A.XX INVENTORY IOD

A.XX.1 Inventory IOD Description

375 The Inventory IOD provides an inventory of all Studies, Series, and SOP Instances managed by an information system, or a specified subset thereof. It includes information about the available mechanisms to access such Studies, Series, and SOP Instances, including both DICOM and non-DICOM protocols. The Inventory provides selected Patient, Procedure, and Imaging Service Request attributes for the inventoried Studies.

For implementation specific reasons, a complete inventory may need to be divided across multiple SOP
380 Instances. The Inventory IOD allows a SOP Instance to reference other SOP Instances whose content is logically included by such reference. A complete inventory thus consists of one root SOP Instance, with a tree of references to additional SOP Instances, the content of all of which comprises the inventory.

> Note:

> See additional explanatory information in Annex XXXX in PS3.17.

385 ### A.XX.2 Inventory IOD Entity-Relationship Model

The Inventory IOD uses the E-R Model specified in Section 7.13.X.

A.XX.3 Inventory IOD Module Table

Table A.XX.3-1 lists the Modules that make up the Inventory IOD.

Table A.XX.3-1. Inventory IOD Modules

IE	Module	Reference	Usage
Equipment	General Equipment	C.7.5.1	M
Inventory	SOP Common	C.12.1	M
	Inventory	C.YY.1	M

390 Note:

The attributes of the other Information Entities of the Inventory Information Model (such as Study, Series, Patient, etc.) are hierarchically encoded (within Sequence attributes) in the Inventory Module, and thus do not appear as separate Modules in the IOD.

395

Add new section for Inventory Creation IOD to Annex B Normalized Information Object Definitions

B.XX INVENTORY CREATION IOD

B.XX.1 Inventory Creation IOD Description

400 The Inventory Creation IOD describes the Attributes that may be present in an Inventory Creation Request or Response.

B.XX.2 Inventory Creation IOD Module Table

Table B.XX.2-1. Inventory Creation IOD Modules

Module	Reference	Description
SOP Common	C.12.1	Contains SOP common information
Inventory Creation	C.YY.3	Contains parameters for initiation of inventory production

405

Add new section for Inventory Modules to Annex C Information Module Definitions

C.YY INVENTORY MODULES

C.YY.1 Inventory Module

The attributes of the Inventory Module are shown in Table C.YY.1-1.

410 The attributes of the Information Entities of the Inventory Information Model (such as Study, Series, Patient, etc. – see Section 7.13.X) are hierarchically encoded within Sequence attributes in the Inventory Module.

Table C.YY.1-1 Inventory Module Attributes

Name	Tag	Type	Description
Content Date	(0008,0023)	1	With Content Time, time point at which the inventory creation began. See Section C.YY.1.1.1
Content Time	(0008,0033)	1	With Content Date, time point at which the inventory creation began. See Section C.YY.1.1.1
Include *Table C.YY.2-1 Scope of Inventory Macro Attributes*			Parameters for the studies that are included in the inventory. See Section C.YY.1.1.2
Transaction UID	(0008,1195)	1C	UID of the Inventory Creation SOP Class transaction that created this Inventory SOP Instance. Required if Inventory created by such a transaction, may be present otherwise.
Inventory Purpose	(00gg,0Fx1)	2	Purpose for which the inventory was created
Inventory Instance Description	(00gg,0Fx2)	3	Description of the content of this Inventory SOP Instance, such as limitation of the scope of this instance within the Scope of Inventory.
Inventory Level	(00gg,0Fx3)	1	Information Entity level for records included in inventory. Enumerated Values: STUDY – study records only SERIES – study and series records only INSTANCE – study, series, and instance records
Inventory Completion Status	(0400,06x5)	1	Degree of completion of inventory with respect to the defined Scope of Inventory. See Section C.YY.1.1.3
Number of Study Records	(0400,06x6)	1	Number of Inventoried Studies Sequence (0400,06x1) Items in this inventory, including those in SOP Instances referenced in the Incorporated Inventory Instance Sequence (0400,06x0). Note A Study that appears multiple times among the referenced Inventory SOP Instances will be counted for each appearance.
Incorporated Inventory Instance Sequence	(0400,06x0)	2	Inventory SOP Instances whose content matches the Scope of Inventory, and whose content is included by reference. See Section C.YY.1.1.4 Zero or more Items shall be included in this Sequence.
>Include *Table C.YY.2-3 "Inventory Reference Macro"*			

Name	Tag	Type	Description
Inventoried Studies Sequence	(0400,06x1)	2	Sequence of items describing Studies whose attributes match the Scope of Inventory. See Section C.YY.1.1.5 Zero or more Items shall be included in this Sequence.
Study			
>Study Instance UID	(0020,000D)	1	Unique identifier for the Study
>Modalities in Study	(0008,0061)	1	All of the distinct values used for Modality (0008,0060) in the Series of the Study.
>Item Inventory DateTime	(00gg,0Fx5)	1	Datetime of creation of the inventory information for this Study. All Study attributes in this Sequence Item are correct as of this datetime. The value shall be at or after the Content Date (0008,0023) and Content Time (0008,0033) of this SOP Instance. Note This attribute may be used for Study record reconciliation. See Section XXXX.6.8 in PS3.17
>Removed from Operational Use	(00gg,0Fx6)	1C	Flag that this Study has been removed from operational use. Enumerated Values: Y N Required if Study is not to be used for operational purposes. May be present otherwise. See Section C.YY.1.2.2
>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason for Study removal from operational use. Required if Removed from Operational Use (00gg,0Fx6) value is Y. See Section C.YY.1.2.2. Zero or one Item shall be included in this Sequence.
>>Include *Table 8.8-1 "Code Sequence Macro Attributes"*			Defined Context Group DCID 7xxx
>Number of Study Related Series	(0020,1206)	1	Count of stored Series in Study. See Section C.YY.1.2.3
>Number of Study Related Instances	(0020,1208)	1	Count of stored SOP Instances in Study. See Section C.YY.1.2.3
>Study Update DateTime	(0020,121x)	2	Datetime of last update to Study instances or metadata managed in the storage system. See Section C.YY.1.2.1
>Study ID	(0020,0010)	2	User or equipment generated Study identifier
>Study Date	(0008,0020)	2	Date the Study started
>Study Time	(0008,0030)	2	Time the Study started
>Study Description	(0008,1030)	2	Description or classification of the Study performed
>Anatomic Regions in Study Code Sequence	(0008,0063)	3	Anatomic regions and body parts examined in the study. See Section C.YY.1.2.4 One or more Items are permitted in this Sequence.
>>Include *Table 8.8-1 "Code Sequence Macro Attributes"*			
>Procedure Code Sequence	(0008,1032)	3	Type of procedure performed. One or more Items are permitted in this Sequence.

Name	Tag	Type	Description
>>*Include Table 8.8-1 "Code Sequence Macro Attributes"*			
>Name of Physician(s) Reading Study	(0008,1060)	3	Names of the physician(s) reading the Study
>Physician(s) Reading Study Identification Sequence	(0008,1062)	3	Identification of the physician(s) reading the Study. One or more Items are permitted in this Sequence. If more than one Item, the number and order shall correspond to the value of Name of Physician(s) Reading Study (0008,1060), if present.
>>*Include Table 10-1 "Person Identification Macro Attributes"*			
>Consulting Physician's Name	(0008,009C)	3	Consulting physician(s) for this Patient Visit.
>Consulting Physician Identification Sequence	(0008,009D)	3	Identification of the consulting physician(s). One or more Items shall be included in this Sequence. If more than one Item, the number and order shall correspond to the value of Consulting Physician's Name (0008,009C), if present.
>>*Include Table 10-1 "Person Identification Macro Attributes"*			
>Physician(s) of Record	(0008,1048)	3	Names of the physician(s) who are responsible for overall Patient care at time of Study
>Physician(s) of Record Identification Sequence	(0008,1049)	3	Identification of the physician(s) who are responsible for overall Patient care at time of Study. One or more Items are permitted in this Sequence. If more than one Item, the number and order shall correspond to the value of Physician(s) of Record (0008,1048), if present.
>>*Include Table 10-1 "Person Identification Macro Attributes"*			
>*Include Table C.12-x "Original Attributes Macro Attributes"* [CP2081]			Recording of prior attribute values and provenance of metadata changes at the Study level. See Section C.YY.1.2.5
>Retrieve AE Title	(0008,0054)	1C	AE Title from which this Study may be retrieved. Required if Retrieve URL (0008,1190) is not present. May be present otherwise.
>Retrieve URL	(0008,1190)	1C	Base URI of the Origin Server to retrieve Study data through the Studies Service (see Section 10 in PS3.18) Required if Retrieve AE Title (0008,0054) is not present. May be present otherwise.
>Stored Instance Base URI	(00gg,0Fx9)	3	Base URI for accessing SOP Instances within the study through a non-DICOM protocol. See Section C.YY.1.2.6.
>Folder URI	(00gg,0FxA)	3	Access URI for a folder containing all SOP Instances for this Study. See Section C.YY.1.2.7.
>Stored Instance File URI	(00gg,0FxB)	3	Access URI for a container file containing all SOP Instances for this Study. See Section C.YY.1.2.7

Name	Tag	Type	Description
>Container File Type	(00gg,0FxC)	1C	Type of container file. Required if Stored Instance File URI (00gg,0FxB) is present. See Section C.YY.1.2.8
>Instance Availability	(0008,0056)	3	Specifies how rapidly the Study becomes available for transmission after a retrieval request. See Section C.YY.1.2.9
>Storage Media File-Set ID	(0088,0130)	3	A human readable Identifier that identifies the Storage Media on which the Study resides
Imaging Service Request			
>Accession Number	(0008,0050)	2	A number that identifies the order for the Study.
>Issuer of Accession Number Sequence	(0008,0051)	3	Identifier of the Assigning Authority that issued the Accession Number. Only a single Item is permitted in this Sequence.
>>Include *Table 10-17 "HL7v2 Hierarchic Designator Macro Attributes"*			
>Referring Physician's Name	(0008,0090)	3	Name of the Patient's referring physician
>Referring Physician Identification Sequence	(0008,0096)	3	Identification of the Patient's referring physician. Only a single Item is permitted in this Sequence.
>>Include *Table 10-1 "Person Identification Macro Attributes"*			
Patient			
>Patient's Name	(0010,0010)	2	Patient's full name
>Patient ID	(0010,0020)	2	Primary identifier for the Patient
>Include *Table 10-18 "Issuer of Patient ID Macro Attributes"*			Identifier of the Assigning Authority that issued the Patient ID
>Other Patient IDs Sequence	(0010,1002)	3	A Sequence of identification numbers or codes used to identify the Patient. One or more Items are permitted in this Sequence
>>Patient ID	(0010,0020)	2	An identifier for the Patient
>>Include *Table 10-18 "Issuer of Patient ID Macro Attributes"*			Identifier of the Assigning Authority that issued the Patient ID
>Patient's Birth Date	(0010,0030)	2	Birth date of the Patient
>Patient's Sex	(0010,0040)	2	Sex of the named Patient
>Other Patient Names	(0010,1001)	3	Other names used to identify the Patient
>Inventoried Series Sequence	(0400,06x2)	2C	Description of all Series within the Study. Each Series shall appear in only one Item. Required if Inventory Level (00gg,0Fx3) has value SERIES or INSTANCE. Zero or more Items shall be included in this Sequence.
Series			
>>Series Instance UID	(0020,000E)	1	Unique identifier of the Series
>>Modality	(0008,0060)	1	Type of equipment that originally acquired the data used to create the SOP Instances in this Series

Name	Tag	Type	Description

>>Removed from Operational Use	(00gg,0Fx6)	1C	Flag that this Series has been removed from operational use.
			Enumerated Values:
			   Y
			   N
			Required if Series is not to be used for operational purposes. May be present otherwise. See Section C.YY.1.2.2

>>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason for Series removal from operational use. Required if Removed from Operational Use (00gg,0Fx6) value is Y. See Section C.YY.1.2.2. Zero or one Item shall be included in this Sequence.
>>>*Include Table 8.8-1 "Code Sequence Macro Attributes"*			Defined Context Group DCID 7xxx
>>Series Number	(0020,0011)	2	A number that identifies this Series
>>Series Description	(0008,103E)	3	Description of the Series
>>Series Date	(0008,0021)	3	Date the Series started
>>Series Time	(0008,0031)	3	Time the Series started
>>Body Part Examined	(0018,0015)	3	Identification of the part of the body examined. See Section C.YY.1.3.2
>>Laterality	(0020,0060)	3	Laterality of (paired) body part examined.
>>Request Attributes Sequence	(0040,0275)	3	Attributes from the Imaging Service Request. See Section C.YY.1.3.3 One or more Items are permitted in this Sequence.
>>*Include Table C.12-x "Original Attributes Macro Attributes"* [CP2081]			Recording of prior attribute values and provenance of metadata changes at the Series level. See Section C.YY.1.2.5
>>Instance Availability	(0008,0056)	3	Specifies how rapidly the Series becomes available for transmission after a retrieval request. See Section C.YY.1.2.9
>>Storage Media File-Set ID	(0088,0130)	3	A human readable Identifier that identifies the Storage Media on which the Series resides
>>Stored Instance Base URI	(00gg,0Fx9)	3	First part of URI for accessing SOP Instances within this Series through a non-DICOM protocol. See Section C.YY.1.2.6.
>>Folder URI	(00gg,0FxA)	3	Access URI for a folder containing all SOP Instances for the Series. See Section C.YY.1.3.1.
>>Stored Instance File URI	(00gg,0FxB)	3	Access URI for a container file containing all SOP Instances for the Series. See Section C.YY.1.3.1.
>>Container File Type	(00gg,0FxC)	1C	Type of container file. Required if Stored Instance File URI (00gg,0FxB) is present. See Section C.YY.1.2.8 for Defined Terms.
>>>*Include Table 10-9 "Request Attributes Macro Attributes"*			

Name	Tag	Type	Description
>>Inventoried Instances Sequence	(0400,06x2)	2C	Description of all SOP Instances within the Series. Each SOP Instance shall appear in only one Item. Required if Inventory Level (00gg,0Fx3) has value INSTANCE. Zero or more Items shall be included in this Sequence.
SOP Instance			
>>>SOP Class UID	(0008,0016)	1	SOP Class UID of the SOP Instance
>>>SOP Instance UID	(0008,0018)	1	SOP Instance UID of the SOP Instance
>>>Removed from Operational Use	(00gg,0Fx6)	1C	Flag that this SOP instance has been removed from operational use. Enumerated Values: Y N Required if SOP instance is not to be used for operational purposes. May be present otherwise. See Section C.YY.1.2.2
>>>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason for SOP instance removal from operational use. Required if Removed from Operational Use (00gg,0Fx6) value is Y. See Section C.YY.1.2.2. Zero or one Item shall be included in this Sequence.
>>>>*Include Table 8.8-1 "Code Sequence Macro Attributes"*			Defined Context Group DCID 7xxx
>>>*Include Table C.12-x "Original Attributes Macro Attributes"* [CP2081]			Recording of prior attribute values and provenance of metadata changes at the Instance level. See Section C.YY.1.2.5
>>>Instance Availability	(0008,0056)	3	Specifies how rapidly the Instance becomes available for transmission after a retrieval request. See Section C.YY.1.2.9
>>>Storage Media File-Set ID	(0088,0130)	3	A human readable Identifier that identifies the Storage Media on which the Instance resides
>>>File Access Sequence	(00gg,0FyF)	3	Non-DICOM protocol methods to access the SOP Instance in the DICOM File Format. See Section C.YY.1.4.1 and Section C.YY.1.4.2 One or more Items are permitted in this Sequence.
>>>>*Include Table C.YY.2-2 "Stored File Access Macro"*			
>>>>Expiration Datetime	(00gg,0Fy9)	3	Expected time after which SOP Instance file might no longer be accessed via this non-DICOM protocol URI. See Section C.YY.1.4.3

C.YY.1.1 Inventory Module Attributes

415 ### C.YY.1.1.1 Content Date and Content Time

Content Date (0008,0023) and Content Time (0008,0033) establish the time point at which the creation of the inventory began. Conceptually, all Studies received by the storage system before this DateTime, and that match the specified Scope of Inventory key attributes, are included in the inventory, and all patient and study updates through this DateTime have been applied to the attributes as recorded in the Inventory SOP Instance(s).

420 While Content Date and Time specify the nominal time for creation of the inventory, each Study record has its own Item Inventory DateTime (00gg,0Fx5) that specifies the time of creation of that record, and its attribute values are accurate as of that time.

Whether Studies and updates received after Content Date and Time might be included in the inventory is implementation specific.

425 **C.YY.1.1.2 Scope of Inventory**

The Scope of Inventory Sequence (00gg,0Fx0) defined by the Scope of Inventory Macro specifies the values of Study attributes that select Studies to be included in the inventory. and the Scope of Inventory specifies the selection parameters for the current SOP Instance and all its subsidiary incorporated Instances. The Scope of Inventory within those subsidiary incorporated Instances is relevant only to their respective sub-trees.

430   Note

For example, an Inventory SOP Instance may specify in its Scope of Inventory that it includes studies from the years 2015-2019, but it does that by linking five subsidiary Inventory SOP Instances, one for each year, and which each specifies in its Scope of Inventory the year to which it applies.

The content of an Inventory SOP Instance, together with its subsidiary incorporated Instances, is complete with
435 respect to the Scope of Inventory only if Inventory Completion Status (0400,06x5) has the value COMPLETE.

C.YY.1.1.3 Inventory Completion Status

Inventory Completion Status (0400,06x5) is the status of the inventory with respect to the defined Scope of Inventory as of the Content Date and Content Time for the system identified in the General Equipment Module (see Section C.7.5.1), including all SOP Instances referenced in the Incorporated Inventory Instance Sequence
440 (0400,06x0).

Enumerated Values:

COMPLETE   All Studies matching the Scope of Inventory are included in the Inventory SOP Instance and its subsidiary Incorporated Inventory SOP Instances

FAILURE   Production of the inventory failed; some Studies matching the Scope of Inventory might not be
445   included in the inventory

CANCELED   Production of the inventory canceled; some Studies matching the Scope of Inventory might not be included in the inventory

PARTIAL   Production of the inventory incomplete; some Studies matching the Scope of Inventory might not be included in the inventory

450   Note

1.   An inventory may be COMPLETE and have no Study records, i.e., for the equipment at the Content Date and Time, no Studies match the Scope of Inventory.

2.   The attribute tag (0400,06x5) places the encoding of this attribute after the end of the encoding of the Inventoried Studies Sequence (0400,06x1), i.e., at the end of the data set after the Study, Series, and Instance inventory
455   records.

C.YY.1.1.4 Incorporated Inventory Instance Sequence

For implementation specific reasons, the content of an inventory may be divided into more than one Inventory SOP Instance. The Inventory SOP Instances are organized as a tree of referenced SOP Instances, with one SOP Instance as the root. The Incorporated Inventory Instance Sequence (0400,06x0) specifies the links from
460 an Inventory SOP Instance to its subsidiary Instances in the tree. The Incorporated Inventory Instance Sequence definition recursively includes itself, thus allowing representation of the complete tree of referenced Inventory SOP Instances.

Referenced Inventory SOP Instances shall have the same Inventory Level (00gg,0Fx3) as the current SOP Instance.

465 The inclusion of an Inventory SOP Instance reference in the Incorporated Inventory Instance Sequence means that the Items of the Inventoried Studies Sequence (0400,06x1) in that referenced SOP Instance are logically included in the current SOP Instance. Therefore, all inventoried Studies in the referenced SOP Instances shall be within the Scope of Inventory of the current SOP Instance. However, the Item Inventory DateTime (00gg,0Fx5) of those Study records may be prior to the Content Date (0008,0023) and Content Time
470 (0008,0033) of the current SOP Instance.

Notes

1. Reasons for splitting an inventory might include practical limits on the maximum size of an individual SOP Instance, parallel processing across multiple separate subsystems, or reuse of historical inventories. See Section XXXX.2.3 in PS3.17.

475 2. These requirements apply to every Inventory SOP Instance. Thus, each Inventory SOP Instance will include in its Incorporated Inventory Instance Sequence the complete tree of references to the subtree of which it is the root.

C.YY.1.1.5 Inventoried Studies Sequence

The Inventoried Studies Sequence (0400,06x1) references Studies whose attributes match the Scope of
480 Inventory. Within the tree of incorporated Inventory SOP Instances, a given Study may be referenced multiple times among the Inventoried Studies Sequence Items. The Items may have different content, but each Item shall have a complete record of the contents of the Study as known by the creator of that Item.

Notes

1. Differences in content may occur due to changes to the metadata or content (SOP Instances) of the Study during
485 the production of the inventory, or due to different Series of a Study being stored on different media, or for other reasons. The application using an Inventory may need to reconcile such multiple occurrences (see Section XXXX.6.8 "Study Record Reconciliation" in PS3.17).

2. If there are multiple records for a study among the tree of incorporated SOP Instances, the Number of Study Records (0400,06x6) value will reflect the number of records, not the number of unique Study UIDs.

490

C.YY.1.2 Study Attributes

C.YY.1.2.1 Study Update DateTime

Study Update DateTime (0020,121x) is the datetime of the last update to Study contents (SOP Instances) or to Study metadata within the storage system. Metadata may change either due to coercion of study related data
495 (Patient ID, Study Description, Accession Number, etc.) to correspond with values set by the enterprise or department master information system (HIS / EMR / RIS), or due to other updates applied by the storage application entity.

As this is an attribute of the Study as managed within the storage system, the initial Study Update DateTime is the time that the Study was first received in the storage system, and is updated upon each new SOP Instance
500 received for the Study, as well as for updates of patient, procedure, and other metadata.

Notes

1. This attribute, like Number of Study Related Series (0020,1206) and Number of Study Related Instances (0020,1208), is defined at the Study level of the Inventory Information Model, but is not part of the Study attributes in SOP Instances. It is an attribute that may be managed by the storage application entity, and might not have a
505 value for some or all of the studies in the inventory.

2. Study Update DateTime is not the time of last user access to the Study contents.

If DateTime Range Match Attribute (00gg,0Fy3) within the Scope of Inventory Sequence (00gg,0Fx0) has a value of 0020121x, indicating matching to Study Update DateTime, but Study Update DateTime is zero-length, then by implication this Study was matched to Study Date and Study Time (see Section C.YY.2.1.1.2).

510 **C.YY.1.2.2 Removed from Operational Use and Reason for Removal Code Sequence**

The Removed from Operational Use (00gg,0Fx6) attribute is defined at the Study, Series, and Instance levels. 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 storage system for other reasons (e.g., for audit of patient radiation exposure). At the Study and Series level, the attribute indicates that the entire Study or Series has been
515 removed from operational use.

Reason for Removal Code Sequence (00gg,0Fx7) provides a coded reason for Study, Series, or Instance removal from operational use. The defined Context Group is CID 7xxx.

Notes

1. The meaning of "operational use related to patient care" is implementation or site specific, but generally includes
520 diagnostic, clinical, and therapeutic uses, as well as administrative uses necessary for providing care (e.g., insurance authorization).

2. The content of the Inventory may not match the information obtained through the Query/Retrieve Service since Studies, Series, or Instances in Inventory SOP Instances marked Removed from Operational Use typically do not appear in Query/Retrieve responses.

525 3. Studies, Series, or Instances might be marked Removed from Operational Use by actions associated with the processing of specific Key Object Selection SOP Instances, e.g., in accordance with [IHE RAD TF-1] *Image Object Change Management Integration Profile* (IOCM). Those Key Object Selection SOP Instances, and their Series, may themselves be marked as Removed from Operational Use. The Context Group for Reason for Removal Code Sequence includes the Key Object Selection Concept Codes specified in IOCM.

530 4. While defined in this specification at the Study, Series, and Instance levels, an Application Entity might not manage this attribute at some, or any, of those levels. E.g., an AE may only manage this attribute at the Instance level, and is not required to infer a value for the Series or Study level.

5. If the Application Entity retains a record of deleted Studies, Series, or Instances, even though the actual Instances are physically deleted, it may include them in the Inventory with a Reason for Removal code of (11yyy1, DCM,
535 "Data not accessible from storage"). Such instances may have an Instance Availability (0008,0056) value "UNAVAILABLE".

C.YY.1.2.3 Number of Study Related Series, Number of Study Related Instances, and Inventoried Series Sequence

540 Depending on the specified Inventory Level (00gg,0Fx3), an Inventory SOP Instance might only contain references at the Study level. In such case, the Number of Study Related Series (0020,1206) and Number of Study Related Instances (0020,1208) shall give accurate counts of stored Series and SOP Instances, respectively, but the Inventoried Series Sequence (0400,06x2) will be absent.

The counts of Series and Instances shall include those that are marked as Removed from Operational Use.

545 **C.YY.1.2.4 Anatomic Regions in Study Code Sequence**

The Anatomic Regions in Study Code Sequence (0008,0063) consolidates the values of Anatomic Region Sequence (0008,2218) and Body Part Examined (0018,0015) in the SOP Instances of this Study. Recognized values of Body Part Examined can be transcoded to their coded equivalent value, such as by using the tables of Annex L "Correspondence of Anatomic Region Codes and Body Part Examined Defined Terms" in PS3.16, or
550 by implementation- or site-specific mappings for locally defined terms.

C.YY.1.2.5 Original Attributes Macro

The Original Attributes Sequence (0400,0561) is defined at each of the Study, Series, and Instance levels in the Inventory IOD.

Within the Inventoried Studies Sequence (0400,06x1), i.e., at the Study level, the Original Attributes Sequence
555 (0400,0561) may be used to record the prior values of updated Study, Patient, and Imaging Study Request attributes for the referenced Study, and the provenance of such updates (see Section C.12.1.1.9). Within the Inventoried Series Sequence (0400, 06x2), i.e., at the Series level, and within the Inventoried Instances Sequence (0400, 06x3), i.e., at the Instance level, the Original Attributes Sequence may be used to record the prior values of updated Series and Instance attributes, respectively, and the provenance of such updates (see
560 Section XXXX.6.7.1 "Original Attributes Sequence" in PS3.17).

In the context of an Inventory SOP Instance, the "top level Data Set" in the definition of the Original Attributes Sequence is the top level Data Set of the SOP Instances of the inventoried Study.

The Original Attributes Sequences in the Inventory SOP Instance might not be identical to the Original Attributes Sequence in the individual SOP Instances of the Study, and updates recorded in either location are valid (see
565 Section C.YY.1.4.2).

C.YY.1.2.6 Stored Instance Base URI

Stored Instance Base URI (00gg,0Fx9) is defined at the Study and Series levels in the Inventory IOD.

Stored Instance Base URI contains an [RFC3986] base URI that may be combined with relative path reference URIs for non-DICOM protocol access to SOP Instances of the Study or Series (see Section P.2.1).

570 For any relative path reference URI, the base URI is the one specified in its level, if present, or recursively at the next higher levels, until a Stored Instance Base URI is present with a value.

Notes

1. Stored Instance Base URI is optional, and if not present the Folder and Stored Instance File URIs at the Study level must be complete URIs. If Stored Instance Base URI is present, those other attributes may still provide complete
575 URIs, rather than relative paths to be merged with the Base URI.

2. For example, the Stored Instance Base URI value may be nfs://pacs.exampleinstitution.org/JZ0078555/, which when combined with a Stored Instance File URI ./2.25.916804767294.dcm gives a URI nfs://pacs.exampleinstitution.org/JZ0078555/2.25.916804767294.dcm to access a SOP Instance through the Network File System protocol. The initial "./" in the Stored Instance File URI indicates it is a relative path.

580 ### C.YY.1.2.7 Folder URI and Stored Instance File URI

If all of the stored SOP Instances of the Study, as identified in this Item, are in the DICOM File Format accessible through a non-DICOM protocol, and all the files are catalogued in a single folder (see Section P.1.3), Folder URI (00gg,0FxA) provides the URI for protocol operations on that folder.

If all of the stored SOP Instances of the Study, as identified in this Item, are in a single container format file as
585 specified in Section P.1.2, Stored Instance File URI (00gg,0FxB) provides the URI for accessing that file.

Folder URI and/or Stored Instance File URI may be a relative path reference beginning with the single-dot-segment "./" (see Section P.2.1) . In this case, the URI is relative to the name space of the Stored Instance Base URI (00gg,0Fx9).

See Section C.YY.1.4.2 for requirements on applying metadata from the Inventory SOP Instance to stored SOP
590 Instances of the Study.

Notes

1. The Study may be inventoried in more than one Item (see Section C.YY.1.1.5). Only the SOP Instances identified within this Item (and its subsidiary Inventoried Instances Sequences) need to be included in the folder or container file.

595 2. Stored SOP Instances from other Studies may be present in the folder or container file; the only requirement is that all of the identified SOP Instances of this Study are included.

3. Files other than stored SOP Instances may be present in the container file, e.g., reports in non-DICOM format.

C.YY.1.2.8 Container File Type

Container File Type (00gg,0FxC) identifies the type of container file accessible through the Stored Instance File
600 URI (00gg,0FxB). The Defined Terms are:

ZIP see Section P.1.2.1

TAR see Section P.1.2.2

GZIP see Section P.1.2.3

TARGZIP see Section P.1.2.4

605 ## C.YY.1.2.9 Instance Availability

Instance Availability (0008,0056) is defined at each of the Study, Series, and Instance levels in the Inventory IOD. It specifies how rapidly the Study, Series, or Instance becomes available for transmission after a retrieval request. For a Study or Series, when some Instances are less rapidly available than others, the value is the availability of the least rapidly available.

610 Enumerated Values:

ONLINE the instance is immediately available

NEARLINE the instance needs to be retrieved from relatively slow media such as optical disk or tape, or requires conversion that takes time

OFFLINE the instance needs to be retrieved by manual intervention

615 UNAVAILABLE the instance cannot be retrieved

Note

1. See also PS3.4 Section C.4.1.1.3.2.

2. A Study, Series, or Instance that is UNAVAILABLE may also be marked as Removed from Operational Use with Reason for Removal Code (11yyy1, DCM, "Data not accessible from storage") (see Section C.YY.1.2.2).

620 ## C.YY.1.3 Series Attributes

C.YY.1.3.1 Folder URI and Stored Instance File URI

If all of the stored SOP Instances of the Series are in the DICOM File Format accessible through a non-DICOM protocol, and all the files are catalogued in a single folder (see Section P.1.3), Folder URI (00gg,0FxA) provides the URI for protocol operations on that folder.

625 If all of the stored SOP Instances of the Series are in a single container file format as specified in Section P.1.2, Stored Instance File URI (00gg,0FxB) provides the URI for accessing that file.

Folder URI and/or Stored Instance File URI may be a relative path reference beginning with the single-dot-segment "./" (see Section P.2.1) . In this case, the URI is relative to the name space of the Stored Instance Base URI (00gg,0Fx9) in this Inventoried Series Sequence Item, if present, or otherwise to the Stored Instance Base
630 URI at the Study level.

C.YY.1.3.2 Body Part Examined

Defined Terms for Body Part Examined (0018,0015) are found in Annex L "Correspondence of Anatomic Region Codes and Body Part Examined Defined Terms" in PS3.16.

Recognized values of Body Part Examined can be transcoded to their coded equivalent value, such as by using
635 the tables of Annex L in PS3.16 or by implementation- or site-specific mappings for locally defined terms. Those codes may be added to the Study level attribute Anatomic Regions in Study Code Sequence (0008,0063) (see Section C.YY.1.2.4).

C.YY.1.3.3 Request Attributes Sequence

Request Attributes Sequence (0040,0275) may be used to support Series that are related to multiple Imaging
640 Service Requests.

> Note

> See, for example, [IHE RAD TF-2] *Section 4.6.4.1.2.3 Relationship between Scheduled and Performed Procedure Steps.*

C.YY.1.4 SOP Instance Attributes

645 ### C.YY.1.4.1 File Access Sequence

The File Access Sequence provides one Item for each URI of a non-DICOM protocol through which the SOP Instance in the DICOM File Format may be accessed (see Section P.1.1). A URI relative path reference within this Sequence shall be relative to the Stored Instance Base URI (00gg,0Fx9) specified at the Series level, if present, or otherwise to the Stored Instance Base URI at the Study level.

650 > Note

> A storage application entity may store a SOP Instance on different storage devices (e.g., fast short-term media and slower long-term media), or with different Transfer Syntaxes. The SOP Instance may therefore be accessible through a non-DICOM protocol at multiple URIs.

C.YY.1.4.2 Metadata from Inventory

655 A stored SOP Instance accessible through a non-DICOM protocol shall be conformant to its IOD and to the DICOM File Format. However, some attributes might not have current values (e.g., Patient Name may have been corrected or changed after the Instance was stored). The current values are encoded in the Inventory SOP Instance. For all attributes defined in the IOD of the stored SOP Instance, and for which values are recorded in the Inventory, the values in the Inventory are authoritative.

660 Items of the Original Attributes Sequences (0400,0561) in the Inventory at the Study, Series, and Instance levels (see Section C.YY.1.2.5) are additive to any Items in the Original Attributes Sequence of the stored SOP Instance.

> Notes:

> 1.   SOP Instances accessed through DICOM protocols are expected to have current values in all attributes.

665 > 2.   Attributes that are not defined in the IOD of the stored SOP Instance, but that are in the Inventory IOD, include study level summary attributes, such as Modalities in Study and Number of Study Related Instances, and file access attributes, such as Retrieve AE Title, Stored Instance Base URI, and Stored File URI. These attributes are not to be imputed to the stored SOP Instance.

3. Attributes not defined in the Inventory IOD, but that are defined in the IOD of the stored SOP Instance, may be
670   included in the Inventory SOP Instance, especially if they have been updated, such as change of an SR Concept Name Code Sequence from a SNOMED RT to a SNOMED CT code. (See Section XXXX.2.5 "Additional Data Elements" in PS3.17.) These attributes are to be imputed to the stored SOP Instance.

4. As described in Section C.12.1.1.9, when performing a change or correction to SOP Instance data, the application may add an Item to the Original Attributes Sequence describing the provenance of the change, but preserving any
675   existing Items in the Sequence. Thus, changes to Study, Patient, or Imaging Service Request attributes recorded in the Inventory may have the provenance of those changes recorded in the Original Attributes Sequence at the Study level, changes to Series or Instance attributes may have their provenance recorded in the Original Attributes Sequence at the Series or Instance level, respectively. When those changes are applied to the stored SOP Instances of the Study, all Items in the Original Attributes Sequence in the Inventory at the Study, Series, and
680   Instance levels are added to the Original Attributes Sequence of the stored SOP Instances.

C.YY.1.4.3 Expiration Datetime

A storage application entity may store a SOP Instance in a temporary location (e.g., on fast short-term media). If such storage has a scheduled or expected expiration time, it may be identified in the Expiration Datetime (00gg,0Fy9) attribute. This is not a guaranteed expiration datetime; the SOP Instance may become unavailable
685 at the access URI at an earlier or later time. This Expiration Datetime applies only to this non-DICOM protocol access; SOP Instances may continue to be accessible through other DICOM or non-DICOM access methods.

C.YY.2 Inventory Related Macros

Note

690 See Section 5.5 for the interpretation of the Type column when these Macros are invoked in Normalized IODs.

C.YY.2.1 Scope of Inventory Macro

Table C.YY.2-1 specifies the attributes that define the parameters that constrain the Studies included in the inventory, i.e., the scope of the inventory. An empty attribute means the scope is unconstrained by the value of the corresponding attribute in the Study.

695 The full semantics of the matching process is specified in Section ZZ.2.2.1.1 "Scope of Inventory Sequence" in PS3.4.

Table C.YY.2-1 Scope of Inventory Macro Attributes

Name	Tag	Type	Description
Scope of Inventory Sequence	(00gg,0Fx0)	2	Matching key values for the studies to be included in the inventory. Zero or one Item shall be included in this Sequence. The absence of an Item indicates universal match to all managed Studies. Shall not be extended in a Standard Extended SOP Class (see Section C.YY.2.1.1.1).
>Study DateTime Range	(00gg,0Fy1)	2	Matching datetime range for Study Date and Study Time. See Section C.YY.2.1.1.2.
>Study Update DateTime Range	(00gg,0Fy2)	2	Matching datetime range for Study Update DateTime. See Section C.YY.2.1.1.2.
>Study Description	(0008,1030)	2	Matching key value for Study Description. See Section C.YY.2.1.1.3

>Series Description	(0008,103E)	2	Matching key value for Series Description. See Section C.YY.2.1.1.3
>Modalities in Study	(0008,0061)	2	Matching key values for Modality. See Section C.YY.2.1.1.4
>Study Instance UID List	(00gg,0Fy4)	2	Matching key values for Study Instance UID. See Section C.YY.2.1.1.5
>Patient List Sequence	(00gg,0Fy6)	2	Match key values for Patient ID. See Section C.YY.2.1.1.6 Zero or more Items shall be included in this Sequence.
>>Issuer of Patient ID	(0010,0021)	2	Matching key value for Issuer of Patient ID. See Section C.YY.2.1.1.6
>>Patient ID List	(00gg,0Fy5)	1	Matching key values for Patient ID. See Section C.YY.2.1.1.6

>Attributes Missing Value	(00gg,0FyD)	2	Attributes that match the scope of inventory when the corresponding attribute in the study has no value. Enumerated Values:
			    00080020 – Study Date
			    00080050 – Accession Number
			    00100020 – Patient ID
			See Section C.YY.2.1.1.7

C.YY.2.1.1 Scope of Inventory Macro Attributes

C.YY.2.1.1.1 Scope of Inventory Sequence

The Scope of Inventory Sequence (00gg,0Fx0) specifies attributes parameterizing the selection of Studies to be included in an inventory. A Standard Extended SOP Instance (see Section 3.11 "DICOM Conformance" in PS3.2) shall not include additional Type 3 attributes in the Scope of Inventory Sequence.

Note

Each attribute of this sequence contributes to the semantics of inclusion in the set of inventoried Studies, and an additional attribute would thus change the semantics of the IOD.

C.YY.2.1.1.2 Study DateTime Range and Study Update DateTime Range

Study DateTime Range (00gg,0Fy1) and Study Update DateTime Range (00gg,0Fy2) specify datetime ranges for studies to be included in an inventory. Study DateTime Range is matched to the combination of Study Date (0008,0020) and Study Time (0008,0030), and Study Update DateTime Range is matched to Study Update DateTime (0020,121x).

Each attribute has two values (VM=2), representing the beginning and the end of the time range. Either the first and/or the second value may be empty, in which case the datetime range has an unbounded beginning or end, respectively. The datetime range is inclusive; Studies whose time is at or after the first value and before or after the second value match the range.

If Study DateTime Range is specified for matching, but a Study does not have a value for Study Date, that Study is not included in the inventory.

If Study Update DateTime is specified for matching, but a Study does not have a defined Study Update DateTime available for matching, the datetime range is matched to Study Date and Time as a default.

Note

See discussion of some operational considerations with respect to time matching to Study Update DateTime in Section XXXX.6.2 "Obtaining Incremental Inventories" in PS3.17.

C.YY.2.1.1.3 Study Description and Series Description

Study Description (0008,1030) and Series Description (0008,103E) specify strings to be matched to the
respective attributes in studies to be included in an inventory. The characters "*" and "?" are interpreted as wild
cards for matching, as described in Section C.2.2.2.4. "Wild Card Matching" in PS3.4.

A match of the Series Description to that attribute in any Series of a Study is a successful match for the Study to
be included in an inventory.

C.YY.2.1.1.4 Modalities in Study

Modalities in Study (0008,0061) provides a list of codes to be matched against the Modality (0008,0060)
attribute of the Series of Studies to be included in an inventory. Multiple values may be present, and the
equivalence of each and every value from this list with any Modality value among the Series of a Study is a
successful match.

C.YY.2.1.1.5 Study Instance UID List

Study Instance UID List (00gg,0Fy4) provides a list of UIDs to be matched against the Study Instance UID
(0020,000D) attribute of studies to be included in an inventory. Multiple values may be present.

> Note
>
> The UI VR allows a length of up to 2^{16}-2 bytes, limiting the number of UIDs in the list. Uses that require inventories of
> more Studies may need to initiate multiple inventories.

C.YY.2.1.1.6 Patient List Sequence

Patient List Sequence (00gg,0Fy6) provides a list of patient identifiers whose studies are to be included in an
inventory. Matching must be successful to both the Patient ID (0010,0020) and the Issuer of Patient ID
(0010,0021) attribute from the same Patient List Sequence Item.

Patient ID List (00gg,0Fy5) provides a list of values to be matched against the Patient ID (0010,0020) attribute
of studies to be included in an inventory. Multiple values may be present. A match of any value in the Patient ID
List to the Patient ID attribute is a successful match.

Issuer of Patient ID (0010,0021) is matched to the same attribute for the Study. An absent value for the attribute
in the Patient List Sequence Item is universal matching.

> Note
>
> The LO VR for Patient ID List allows a length of up to 2^{16}-2 bytes, limiting the number of IDs in the list. Uses that require
> inventories of more patients may use multiple Patient List Sequence Items, possibly with the same Issuer of Patient ID
> value.

C.YY.2.1.1.7 Attributes Missing Value

Attributes Missing Value (00gg,0FyD) specifies the tags of attributes that match the scope of inventory when the
corresponding attribute in the Study has an empty value, i.e., studies for which the identified attributes have no
value are to be included in the inventory. Multiple values may be present, and the absence of a value in any of
the attributes identified in this list is a successful match.

Attributes Missing Value shall not have a value of 00080020 (Study Date) if Study DateTime Range (00gg,0Fy1)
has a value. Attributes Missing Value shall not have a value of 0010020 (Patient ID) if Patient List Sequence
(00gg,0Fy6) has a value.

C.YY.2.2 Stored File Access Macro

Table C.YY.2-2 specifies the attributes that describe non-DICOM protocol access to a stored SOP Instance in
the DICOM File Format, possibly contained in a container file.

Table C.YY.2-2 Stored File Access Macro Attributes

Name	Tag	Type	Description
Stored Instance File URI	(00gg,0FxB)	1C	Access URI for file containing the SOP Instance. See Section C.YY.2.2.1.1 Required if referenced SOP instance is in the DICOM File Format, and is accessible through a non-DICOM protocol (see Annex P).
Container File Type	(00gg,0FxC)	1C	Type of container file. See Section C.YY.1.2.8 for Defined Terms. Required if Stored Instance File URI references a container file.
Filename in Container	(00gg,0FxD)	1C	Filename within a container file of the file containing the SOP Instance. See Section C.YY.2.2.1.2 Required if Stored Instance File URI references a container file.
File Offset in Container	(00gg,0FxE)	3	Byte offset (zero-based) within a container file for the start of the SOP Instance file. See Section C.YY.2.2.1.2
Stored Instance Transfer Syntax UID	(00gg,0Fy0)	1	Transfer Syntax of SOP Instance encoded in DICOM File Format. Equal to Transfer Syntax UID (0002,0010) in File Meta Information header.
MAC Algorithm	(0400,0015)	3	The algorithm used for generating a Message Authentication Code. See Table C.12-6a for Defined Terms. [CP2092]
MAC	(0400,0404)	3	Message Authentication Code computed across stored instance file for verification of file integrity. See Section C.YY.2.2.1.3

765

C.YY.2.2.1 Stored File Access Macro Attributes

C.YY.2.2.1.1 Stored Instance File URI

Stored Instance File URI (00gg,0FxB) provides the URI for accessing a file containing the SOP Instance in the DICOM File Format.

770 Stored Instance File URI may be a relative path reference beginning with the single-dot-segment "./" (see Section P.2.1) . In this case, the value is relative to the name space of a Stored Instance Base URI (00gg,0Fx9) specified for the context of this attribute.

C.YY.2.2.1.2 Filename in Container, and File Offset in Container

If the stored SOP Instance is included in a container file as specified in Section P.1.2, Filename in Container
775 (00gg,0FxD) is the filename within that container file of the file containing the SOP Instance. File Offset in Container (00gg,0FxE) is an optional byte offset for the start of the file for the SOP Instance. For TARGZIP container files, this is the offset within the TAR container file after extraction from the GZIP container file.

> Note
>
> Locating a file by filename within a TAR container file may require parsing the entire TAR file. The File Offset in
> 780 Container is an optimization, if known, to quickly locate a specific file.

C.YY.2.2.1.3 MAC

The integrity of a stored SOP Instance file may be verified by a Message Authentication Code (also known as a message digest, hash, or cryptographic checksum) computed across the file. The MAC (0400,0404) value is computed across the entire file as a byte stream, including the Preamble, Meta-Information Header, and Data

785 Set Trailing Padding (see Section 7 "DICOM File Format" in PS3.10). For files stored in container files, the MAC is computed on the file extracted from the container.

Note

This is different from the MAC for Digital Signatures (see Section C.12.1.1.3), which is computed across an enumerated list of attributes within the SOP Instance, not across the entire file.

790

C.YY.2.3 Inventory Reference Macro

Table C.YY.2-3 specifies the attributes that describe a referenced Inventory SOP Instance, and its referenced Inventory SOP Instances in a hierarchical tree.

Note

795 Because each node in the tree of references includes its complete subsidiary tree, it is not possible to have cycles of references (a node being a descendent of itself).

Table C.YY.2-3 Inventory Reference Macro Attributes

Name	Tag	Type	Description
Referenced SOP Class UID	(0008,1150)	1	SOP Class UID of the Inventory SOP Instance
Referenced SOP Instance UID	(0008,1155)	1	SOP Instance UID of the Inventory SOP Instance
Retrieve AE Title	(0008,0054)	1C	AE Title from which this Inventory SOP Instance may be retrieved. Required if Inventory SOP instance is available through Inventory MOVE or Inventory GET SOP Classes
Retrieve URL	(0008,1190)	1C	Base URI of the origin server to retrieve this Inventory SOP Instance. Required if Inventory SOP instance is available through the Non-Patient Instance Service (see Section 12 in PS3.18)
Include *Table C.YY.2-2 "Stored File Access Macro"*			Stored Instance File URI (00gg,0FxB) shall not be a relative path reference as there is no Stored Instance Base URI (00gg,0Fx9) specified for the context of this attribute.
Incorporated Inventory Instance Sequence	(0400,06x0)	1C	Inventory SOP Instances whose content is included by reference. See Section C.YY.1.1.4 Required if referenced Inventory SOP Instance itself includes Inventory SOP Instances by reference. One or more Items shall be included in this Sequence.
>Include *Table C.YY.2-3 "Inventory Reference Macro"*			Recursive inclusion to describe a tree of referenced SOP Instances

800 **C.YY.3 Inventory Creation Module**

Table C.YY.3-1 specifies the attributes that are contained in an Inventory Creation Request/Response.

Table C.YY.3-1 Inventory Creation Module Attributes

Name	Tag	Description
Include *Table C.YY.2-1 Scope of Inventory Macro Attributes*		Parameters for the studies to be included in the inventory
Transaction UID	(0008,1195)	Unique identifier of a transaction initiating the production of an inventory
Transaction Status	(00gg,0FyB)	Processing status for this transaction. (See Section ZZ.2.1.1 in PS3.4) Enumerated Values: PROCESSING – Processing continuing COMPLETE – Processing complete with completed inventory FAILURE – Processing failure CANCELED – Processing canceled by requester PAUSED – Processing paused or suspended
Transaction Status Description	(00gg,0FyC)	Explanation or further detail of Transaction Status
Expiration Datetime	(00gg,0Fy9)	Expected time after which Inventory SOP Instances might no longer be accessed. See Section C.YY.3.1.2
Inventory Purpose	(00gg,0Fx1)	Purpose for which the inventory is to be created
Inventory Level	(00gg,0Fx3)	Information Entity level for records to be included in inventory. Enumerated Values: STUDY – study records SERIES – study and series records INSTANCE – study, series, and instance records
Referenced SOP Class UID	(0008,1150)	SOP Class UID of the produced Inventory SOP Instance
Referenced SOP Instance UID	(0008,1155)	SOP Instance UID of the produced Inventory SOP Instance at the root of the tree of incorporated Inventory SOP Instances (see Section C.YY.1.1.4)
Retrieve AE Title	(0008,0054)	AE Title from which this Inventory SOP Instance may be retrieved.
Retrieve URL	(0008,1190)	Base URI of the origin server to retrieve this Inventory SOP Instance.
Include *Table C.YY.2-2 "Stored File Access Macro"*		See Section C.YY.3.1.1.
Number of Study Records	(0400,06x6)	Number of Studies found and processed into Inventory SOP Instances (as of time of message) that match Scope of Inventory.
Requested Status Interval	(00gg,0Fy7)	Requested interval between Inventory Status notifications, in minutes
Retain Instances	(00gg,0Fy8)	For a Canceled Inventory, whether produced SOP Instances should be retained. Enumerated Values: Y – SOP Instances are to be retained N – SOP Instances may be deleted

C.YY.3.1 Inventory Creation Module Attributes

C.YY.3.1.1 Stored File Access Macro

If the resultant Inventory SOP Instance is stored in the DICOM File Format (see Section 7 "DICOM File Format" in PS3.10) and is accessible through a non-DICOM protocol (see Annex P), the Stored Instance File URI (00gg,0FxB) contains a [RFC3986] URI for non-DICOM protocol access to the Inventory SOP Instance.

Stored Instance File URI may reference a container format file that includes the Inventory SOP Instance.

Note

Storage in a container file (ZIP or GZIP) supports the application of compression to the entre file using operating system applications, rather than compression of the SOP Instance using the Deflated Little Endian Transfer Syntax (see Section A.5 in PS3.5) using DICOM applications.

C.YY.3.1.2 Expiration Datetime

A data management system may manage the storage lifecycle of the Inventory SOP Instances (e.g., deleting older objects after a predetermined time). If there is a scheduled or expected time for SOP Instance deletion, it may be identified in the Expiration Datetime (00gg,0Fy9) attribute. This is not a guaranteed expiration datetime; the SOP Instances may become unavailable at an earlier or later time. This Expiration Datetime applies to all available access methods (both DICOM and non-DICOM protocols).

Add new Inventory Directory Record Type to Section F.3.2.2

Table F.3-3. Directory Information Module Attributes

Attribute Name	Tag	Type	Attribute Description
>Directory Record Type	(0004,1430)	1	Defines a specialized type of Directory Record by reference to its position in the Media Storage Directory Information Model (see Section F.4). Enumerated Values: **PATIENT** … **INVENTORY** **PRIVATE** Privately defined record hierarchy position. Type shall be defined by Private Record UID (0004,1432).

825

Add new Inventory Directory Record Type to Section F.4 and update Figure

Table F.4-1. Relationship Between Directory Records

Directory Record Type	Section	Directory Record Types that may be included in the next lower-level directory Entity
(Root Directory Entity)		PATIENT, HANGING PROTOCOL, … **INVENTORY,** PRIVATE
INVENTORY	**F.5.x**	**PRIVATE**
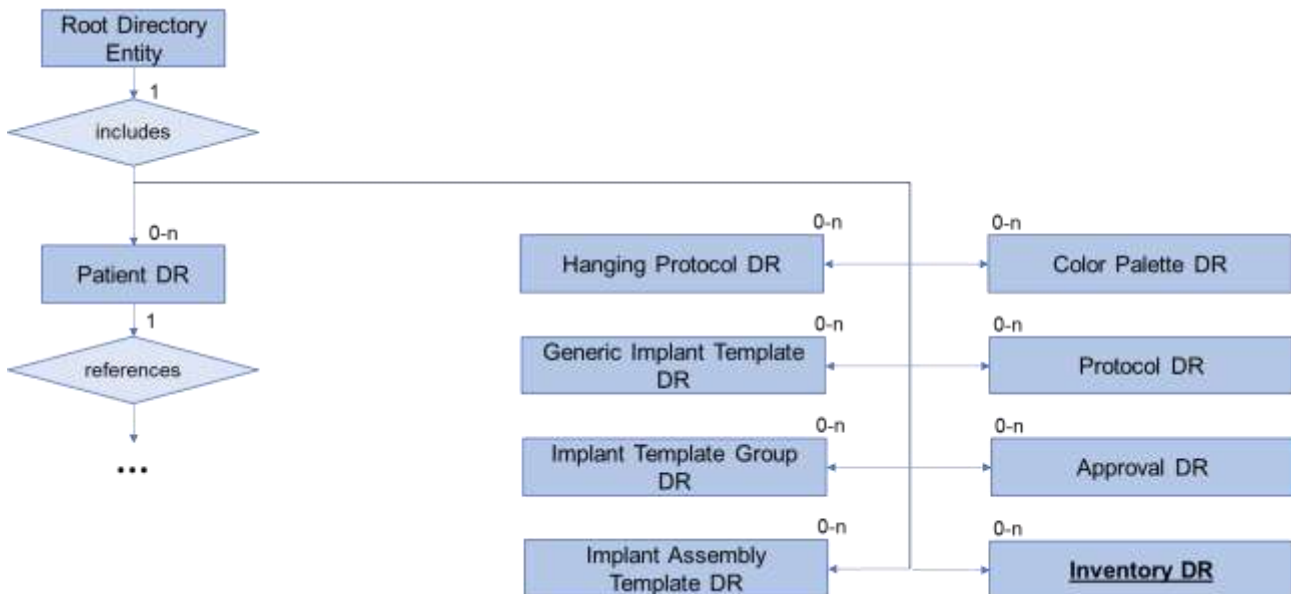
830

Figure F.4-1. Basic Directory IOD Information Model

Add new Inventory Directory Record Definition to Section F.5

F.5.x Inventory Directory Record Definition

835 The Directory Record is based on the specification of Section F.3. It is identified by a Directory Record Type of Value "INVENTORY". Table F.5-x lists the set of keys with their associated Types for such a Directory Record Type. The description of these keys may be found in the Inventory IOD. This Directory Record shall be used to reference an Inventory SOP Instance. This type of Directory Record may reference a Lower-Level Directory Entity that includes one or more Directory Records as defined in Table F.4-1.

840 **Table F.5-x. Inventory Keys**

Attribute Name	Tag	Type	Attribute Description
Specific Character Set	(0008,0005)	1C	Required if an extended or replacement character set is used in one of the keys
Content Date	(0008,0023)	1	
Content Time	(0008,0033)	1	
Transaction UID	(0008,1195)	1	
Inventory Purpose	(00gg,0Fx1)	2	
Inventory Level	(00gg,0Fx3)	1	
Include *Table C.YY.2-1 Scope of Inventory Macro Attributes*			
Inventory Completion Status	(0400,06x5)	1	
Number of Study Records	(0400,06x6)	1	

Add new Annex P Stored File Access Through Non-DICOM Protocols

845
Annex P Stored File Access Through Non-DICOM Protocols (Normative)

The Inventory IOD (see Section A.XX, and specifically the Inventory Module in Section C.YY.1) includes optional attributes providing a URI link to SOP Instances stored in the DICOM File Format and accessible through a non-DICOM file access protocol. Additionally, Inventory SOP Instances themselves may be stored in
850 the DICOM File Format and be accessible through a non-DICOM file access protocol as specified in a URI link in the Inventory Creation Response (see Section B.XX, and the Inventory Creation Module in Section C.YY.3). This Annex specifies constraints on those URI links and their target resources (files).

P.1 FILES AND SETS OF FILES

P.1.1 DICOM File Format

855 Each stored SOP Instance that is a target of an Inventory URI link shall be encoded in accordance with the DICOM File Format (see Section 7 "DICOM File Format" in PS3.10).

Note

Depending on the storage mechanism, this target resource may be denoted a "file" or an "object". The term "file" is used in this specification.

860 ### P.1.2 Container File Formats

SOP Instances in the DICOM File Format may be aggregated into container files in accordance with the mechanisms specified in this Section.

Note

Directly or indirectly executable files within container files are a security risk. They should be avoided when producing a
865 container file, and should be mitigated against when reading a container file. See Section XXXX.5.2 in PS3.17.

P.1.2.1 ZIP

Multiple SOP Instance files may be included in a ZIP container file in accordance with [ISO/IEC 21320-1].

Note

[ISO/IEC 21320-1] is a constraint on the [ZIP] specification, and in particular supports only uncompressed or [RFC1951]
870 DEFLATE compressed files, and does not permit encryption.

P.1.2.2 TAR

Multiple SOP Instance files may be included in a TAR container file (i.e., a file in accordance with the *ustar Interchange Format*) in accordance with [POSIX].

Note

875 "ustar" is the formal name, derived from "Unix Standard Tape ARchive", although TAR is the common term and is used in this specification. The specification for *ustar Interchange Format* is in the [POSIX] section "pax - portable archive interchange".

P.1.2.3 GZIP

A single SOP Instance file may be included in a GZIP container file in accordance with [RFC1952].

880   Note

The GZIP format supports compression of the contained file using [RFC1951] DEFLATE compression.

P.1.2.4 TARGZIP

A TAR container file in accordance with [POSIX] may be contained within a GZIP container file in accordance with [RFC1952].

885   Note

The TARGZIP format supports compression of the multiple SOP Instances in the TAR.

P.1.3 Folders for Sets of Files

The file storage mechanism may support data structures that manage references to sets of files in a folder.

Note

890   Depending on the storage mechanism, these structures may alternately be denoted "directories". The term "folder" is used in this specification due to potential ambiguities with the term "directory".

The storage of multiple SOP Instance files may be managed such that all SOP Instances in a Study or in a Series are contained within a single folder. Within that folder, some or all of the files may be included in a ZIP, TAR, GZIP, or TARGZIP container file.

895 ## P.2 ACCESS PROTOCOLS

P.2.1 URI Format

The URI for non-DICOM protocol access to stored SOP Instances shall be in accordance with [RFC3986].

If the URI is split between two data elements, a base URI and a relative path, the merge of those data elements in accordance with [RFC3986] Section 5.2.3 shall form the conformant URI. The split shall be at a path segment
900   boundary, and if a slash ("/") character occurs at that boundary, it shall be placed in the base URI data element only. The second data element comprising the relative-path shall begin with a single-dot-segment "./".

Dot-segments shall not otherwise be used.

P.2.2 Protocol

The non-DICOM file access protocol used in the Inventory URI link is not constrained by this Standard. The
905   Conformance Statement for the implementation shall specify the protocol(s) in its description of conformance to the Inventory SOP Class. Common protocols are listed in Table P.2-1.

Note

1. Conformance specification may be facilitated by reference to IANA-registered schemes
    https://www.iana.org/assignments/uri-schemes/

910   2. Several protocols are layered on HTTPS. While the specific protocol should be specified in the Conformance Statement, the only feature required is the ability to read an object, which is generally implemented simply as an HTTP GET in such protocols.

Table P.2-1. Common Non-DICOM File Access Protocols (Informative)

915

IANA-registered Scheme	Protocol	Further Specification
nfs	NFS	RFC7530
smb	SMB	https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-smb2
https	HTTP GET	Includes various cloud storage implementations

DICOM PS3.4: Service Class Specifications

Add Inventory to Annex GG Non-Patient Object Storage Service Class

GG.3 SOP CLASSES

920 The application-level services addressed by the Non-Patient Object Storage Service Class definition are specified in the SOP Classes specified in Table GG.3-1.

Table GG.3-1. Standard SOP Classes

SOP Class Name	SOP Class UID	IOD Specification (defined in PS3.3)
Hanging Protocol Storage	1.2.840.10008.5.1.4.38.1	Hanging Protocol IOD
Color Palette Storage	1.2.840.10008.5.1.4.39.1	Color Palette IOD
Generic Implant Template Storage	1.2.840.10008.5.1.4.43.1	Generic Implant Template IOD
Implant Assembly Template Storage	1.2.840.10008.5.1.4.44.1	Implant Assembly Template IOD
Implant Template Group Storage	1.2.840.10008.5.1.4.45.1	Implant Template Group IOD
CT Defined Procedure Protocol Storage	1.2.840.10008.5.1.4.1.1.200.1	CT Defined Procedure Protocol IOD
Protocol Approval Storage	1.2.840.10008.5.1.4.1.1.200.3	Protocol Approval IOD
Inventory Storage	**1.2.840.10008.5.1.4.xxuid.1**	**Inventory IOD**

925 *Add new section for Inventory requirements to Section GG.6 Application Behavior*

GG.6.x Inventory Storage SOP Class

GG.6.x.1 Instance Creator

An implementation that conforms to the Inventory Storage SOP Class as an SCU and is a SOP Instance creator shall state in its Conformance Statement:

930 • The mechanisms by which creation of Inventory SOP Instances is initiated.

• Policies/parameters for splitting a logical inventory across multiple Inventory SOP Instances.

• Any non-DICOM mechanisms used to access Inventoried SOP Instances that may be specified by a Stored Instance URI, including protocols (such as NFS, SMB, or HTTP), and use of folders and/or container files (ZIP, TAR, TARGZIP) for the SOP Instances.

935 • Policies/parameters for duration of retention, or for deletion, of created Inventory SOP Instances.

Add Inventory Query/Retrieve Service Class

XX Inventory Query/Retrieve Service Class

XX.1 OVERVIEW

XX.1.1 Scope

The Inventory Query/Retrieve Service Class defines an application-level class-of-service that facilitates discovery of and access to Inventory composite objects.

XX.1.2 Conventions

See Conventions for the Basic Worklist Management Service (see Section K.1.2).

XX.1.3 Service Definition

The SOP Classes of the Inventory Query/Retrieve Service Class are each composed of an Information Model and a DIMSE-C Service Group.

In order to serve as an SCP of the Inventory Query/Retrieve Service Class, a DICOM AE possesses information about the Attributes of a number of Inventory composite SOP Instances. The information is organized into an Inventory Query/Retrieve Information Model.

Two peer DICOM AEs implement a SOP Class of the Inventory Query/Retrieve Service Class with one serving in the SCU role and one serving in the SCP role. SOP Classes of the Inventory Query/Retrieve Service Class are implemented using the DIMSE-C C-FIND, C-MOVE and C-GET services as defined in PS3.7.

The semantics of the C-FIND service are the same as those defined in the Service Definition of the Basic Worklist Management Service Class (see Section K.1.4).

The semantics of the C-MOVE and C-GET services are the same as those defined in the Service Definition of the Query/Retrieve Service Class (see Section C.1.4), with the exception that there is only one level of retrieval.

XX.2 INVENTORY Q/R INFORMATION MODEL

The Inventory Query/Retrieve (Q/R) Information Model is based on an Entity-Relationship Model Definition and a Key Attributes Definition analogous to those defined in the Worklist Information Model Definition of the Basic Worklist Management Service (see Section K.2)

XX.2.1 E-R Model

The Inventory Q/R Information Model is a single level entity:

• Inventory Information Entity
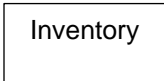
Inventory

Figure XX.2-1. Inventory Q/R Information Model E-R Diagram

While the Inventory Composite SOP Instances that are the subject of this Service Class contain Attributes associated with the several IEs (see Section 7.13.x in PS3.3), the Inventory Q/R Information Model operates only on the Attributes of the Inventory IE of the Inventory IOD.

970 ## XX.2.2 Inventory Q/R Information Model Attributes

Table XX.2-1 defines the Attributes of the Inventory Q/R Information Model.

Table XX.2-1. Attributes for the Inventory Q/R Information Model

Name	Tag	Matching Key Type	Return Key Type	Remark / Return Key Type
SOP Class UID	(0008,0016)	R	1	
SOP Instance UID	(0008,0018)	U	1	.
Transaction UID	(0008,1195)	O	1C	Required if Inventory SOP Instance was created by a Inventory Creation SOP Class transaction, may be present otherwise.
Content Date	(0008,0023)	R	1	
Content Time	(0008,0033)		1	
Scope of Inventory Sequence	(00gg,0Fx0)		2	
>Study DateTime Range	(00gg,0Fy1)		2	
>Study Update DateTime Range	(00gg,0Fy2)		2	
>Study Description	(0008,1030)		2	
>Series Description	(0008,103E)		2	
>Modalities in Study	(0008,0061)		2	
>Study Instance UID List	(00gg,0Fy4)		2	
>Patient List Sequence	(00gg,0Fy6)		2	
>>Issuer of Patient ID	(0010,0021)		2	
>>Patient ID List	(00gg,0Fy5)		1	
>Attributes Missing Value	(00gg,0FyD)		2	
Inventory Purpose	(00gg,0Fx1)		1	
Inventory Instance Description	(00gg,0Fx2)		2	
Inventory Level	(00gg,0Fx3)		1	
Inventory Completion Status	(0400,06x5)		1	
Number of Study Records	(0400,06x6)		1	

975   Note
1. See also the additional returned attributes in the Response Identifier specified in Section XX.4.2.2.1.
2. Further definition of these attributes is specified in Section C.YY.1 in PS3.3

XX.3 DIMSE-C SERVICE GROUPS

980 ## XX.3.1 C-FIND Operation

See the C-FIND Operation definition for the Basic Worklist Management Service Class (Section K.4.1), and substitute "Inventory" for "Worklist". The "Worklist" Search Method shall be used.

The C-FIND is to be performed against the Inventory Q/R Information Model (Section XX.2). The Key Attributes and values allowable for the query are defined in the Inventory Q/R Information Model.

985 ## XX.3.2 C-MOVE Operation

See the C-MOVE Operation definition for the Query/Retrieve Service Class (Section C.4.2). No Extended
Behavior or Relational-Retrieve is defined for the Inventory Query/Retrieve Service Class.

Query/Retrieve Level (0008,0052) is not relevant to the Inventory Query/Retrieve Service Class, and therefore
shall not be present in the Identifier. The Unique Key Attribute of the Identifier is the SOP Instance UID
990 (0008,0018). The SCU shall supply one UID or a list of UIDs.

> Note
>
> More than one Inventory SOP Instance may be retrieved, using List of UID matching.

XX.3.3 C-GET Operation

995 See the C-GET Operation definition for the Query/Retrieve Service Class (Section C.4.3). No Extended
Behavior or Relational-Retrieve is defined for the Inventory Query/Retrieve Service Class.

Query/Retrieve Level (0008,0052) is not relevant to the Inventory Query/Retrieve Service Class, and therefore
shall not be present in the Identifier. The Unique Key Attribute of the Identifier is the SOP Instance UID
(0008,0018). The SCU shall supply one UID or a list of UIDs.

1000 > Note
>
> More than one Inventory SOP Instance may be retrieved, using List of UID matching.

XX.4 SOP CLASS DEFINITIONS

XX.4.1 SOP Classes

1005 The SOP Classes of the Inventory Query/Retrieve Service Class identify the Inventory Q/R Information Model,
and the DIMSE-C operations supported. The following Standard SOP Classes are identified:

Table XX.4-1. Inventory Q/R Service SOP Classes

SOP Class Name	SOP Class UID
Inventory FIND	1.2.840.10008.5.1.4.xxuid.2
Inventory MOVE	1.2.840.10008.5.1.4.xxuid.3
Inventory GET	1.2.840.10008.5.1.4.xxuid.4

1010 ## XX.4.2 Conformance Requirements

An implementation may conform to one of the Inventory Query/Retrieve Service SOP Classes as an SCU, SCP
or both. The Conformance Statement shall be in the format defined in PS3.2.

XX.4.2.1 SCU Conformance

XX.4.2.1.1 C-FIND SCU Conformance

1015 An implementation that conforms to the Inventory FIND SOP Class as an SCU shall support queries against the
Inventory Q/R Information Model using the C-FIND SCU Behavior described for the Basic Worklist Management
Service Class (see Section K.4.1.2 and Section XX.3.1).

XX.4.2.1.2 C-MOVE SCU Conformance

An implementation that conforms to the Inventory MOVE SOP Class as an SCU shall support transfers against
1020 the Inventory Q/R Information Model using the C-MOVE SCU baseline behavior described for the
Query/Retrieve Service Class (see Section C.4.2.2.1 and Section XX.3.2).

XX.4.2.1.3 C-GET SCU Conformance

An implementation that conforms to the Inventory GET SOP Class as an SCU shall support transfers against the Inventory Q/R Information Model using the C-GET SCU baseline behavior described for the Query/Retrieve
1025 Service Class (see Section C.4.3.2.1 and Section XX.3.3).

XX.4.2.2 SCP Conformance

XX.4.2.2.1 C-FIND SCP Conformance

An implementation that conforms to the Inventory FIND SOP Class as an SCP shall support queries against the Inventory Q/R Information Model using the C-FIND SCP Behavior described for the Basic Worklist Management
1030 Service Class (see Section K.4.1.3 and Section XX.3.1).

The implementation shall support the additional Attributes in the Response Identifier as specified in Table XX.4-2 to support character set specification and access to Inventory SOP Instances.

> Note
>
> Compare Section C.4.1.1.3.2 "Response Identifier Structure" for the similar specification of non-key attributes that are to
1035 > be returned in the Query/Retrieve C-FIND SOP Classes to support character set specification and access to the
> referenced SOP Instances.

Table XX.4-2. Inventory FIND additional returned attributes

Attribute	Tag	Type	Remark
Specific Character Set	(0008,0005)	1C	Required if expanded or replacement character sets are used in any attributes of the Query response.
Retrieve AE Title	(0008,0054)	1C	Required if Inventory SOP instance is available through Inventory MOVE or Inventory GET SOP Classes
Retrieve URL	(0008,1190)	1C	Required if Inventory SOP instance is available through the Non-Patient Instance Service (see Section 12 in PS3.18)
Stored Instance File URI	(00gg,0FxB)	1C	Required if Inventory SOP instance is in the DICOM File Format, and is accessible through a non-DICOM protocol (see Annex P in PS3.3).
Container File Type	(00gg,0FxC)	1C	Required if Stored Instance File URI is present and references a container format file.
Filename in Container	(00gg,0FxD)	1C	Required if Stored Instance File URI is present and references a container format file.
Stored Instance Transfer Syntax UID	(00gg,0Fy0)	1C	Required if Stored Instance File URI is present
MAC Algorithm	(0400,0015)	3	
MAC	(0400,0404)	3	

1040 > Note
>
> Further definition of these attributes is specified in Section C.YY.3 in PS3.3

XX.4.2.2.2 C-MOVE SCP Conformance

An implementation that conforms to the Inventory MOVE Class as an SCP shall support transfers against the Inventory Q/R Information Model using the C-MOVE SCP baseline behavior described for the Query/Retrieve
1045   Service Class (see Section C.4.2.3.1). The implementation shall support the Inventory Storage SOP Class (see Section GG) for the C-STORE sub-operations generated by the C-MOVE.

XX.4.2.2.3 C-GET SCP Conformance

An implementation that conforms to the Inventory GET SOP Class as an SCP shall support transfers against the Inventory Q/R Information Model using the C-GET SCP baseline behavior described for the Query/Retrieve
1050   Service Class (see Section C.4.3.3.1). The implementation shall support the Inventory Storage SOP Class (see Section GG) for the C-STORE sub-operations generated by the C-GET.

Add new annex for Inventory Creation SOP Class within a Storage Management Service Class

Annex ZZ Storage Management Service Class

The Storage Management Service Class defines an application-level class-of-service that facilitates peer-to-
peer controls for management of persistent storage of Composite SOP Instances. The Service Class allows
asynchronous operations between the Service Class User (SCU) and the Service Class Provider (SCP).

ZZ.1 OVERVIEW

ZZ.1.1 Use Cases

DICOM supports all manner of peer-to-peer interactions for systems within the biomedical imaging domain. In
many enterprises, one or more systems are responsible for long-term management of stored SOP Instances.
This Service Class supports the interoperability use cases associated with such long-term storage management.

ZZ.1.2 SOP Classes

Each SOP Class of the Storage Management Service Class is formed from a combination of a common DIMSE
Service Group and a specific Information Object Definition.

ZZ.1.2.1 DIMSE Service Group

The DIMSE-N Services applicable to all SOP Classes of the Storage Management Service Class are shown in
Table ZZ.1-1.

Table ZZ.1-1. DIMSE Service Group Applicable to the Storage Management Service Class

DICOM Message Service Element	Usage SCU/SCP
N-ACTION	M/M
N-EVENT-REPORT	M/M

The DIMSE-N Services and Protocol are specified in Section 10 in PS3.7. Additional constraints on these
services, such as specific action and event types, are specified for each SOP Class.

ZZ.1.2.2 Information Object Definitions

The SOP Classes of the Storage Management Service Class are defined using the IODs specified in
Table ZZ.1-2.

Table ZZ.1-2. Storage Management Service SOP Classes

SOP Class Name	SOP Class UID	IOD Specification (defined in PS3.3)
Inventory Creation	1.2.840.10008.5.1.4.xxuid.5	Inventory Creation IOD (Section B.XX in PS3.3)

Additional constraints on these IODs, such as specific attributes required for the different action types, are
specified for each SOP Class.

1080 ## ZZ.1.3 Service Protocol

ZZ.1.3.1 Association Negotiation

Association establishment is the first phase of any instance of communication between peer DICOM AEs. The Association negotiation rules as specified in Annex D in PS3.7 are used to negotiate the supported SOP Classes and peer AE roles.

1085 Note

Implementations may restrict Association establishment subject to exchange of security related information, such as application identity and authorization, either within DICOM Association negotiation or outside the scope of the DICOM protocol. See Section XXXX.5 "Security Considerations" in PS3.17.

Support for the SCP/SCU role selection negotiation is mandatory. The SOP Class Extended Negotiation is not
1090 defined for this Service Class.

The SCU will open an Association when it desires to request a storage management operation by the SCP.

The SCP will typically open an Association when it is reporting status, or has completed the requested operation or reached some other termination condition, such as a failure. This Association establishment includes negotiation of SCP/SCU role.

1095 Note

1. The SCP may attempt to issue the N-EVENT-REPORT on the same Association as the N-ACTION, but this operation may fail because the SCU is free to release at any time the Association on which it sent the N-ACTION-Request.

2. As DICOM defaults the association requestor to the SCU role, the SCP (i.e., the association requester)
1100 negotiates an SCP role using the SCU/SCP role negotiation (see Section D.3.3.4. "SCP/SCU Role Selection Negotiation" in PS3.7).

3. When responding on a different Association, the SCP must use the same AE Title as it used on the original Association, because the DICOM Standard defines a Service between two peer applications, each identified by an AE Title. Thus, the SCP should be consistently identified for all transactions with a particular peer in a SOP
1105 Class.

ZZ.1.3.2 Operations and Notifications

Following Association establishment, peer-to-peer communication between the SCU and SCP uses the DIMSE N-ACTION and N-EVENT-REPORT (see Section 10.1 "DIMSE-N Services" in PS3.7).

The N-ACTION and N-EVENT-REPORT primitives shall contain the well-known Storage Management SOP
1110 Instance UID "1.2.840.10008.1.20.1.x" in their Requested SOP Instance UID and Affected SOP Instance UID parameters.

Note

In the usage described here, there is no explicit creation of a SOP Instance (using the DIMSE N-CREATE) upon which an N-ACTION may operate. Instead, the N-ACTION operates upon a constant well-known SOP Instance. This
1115 SOP Instance is conceptually created during start-up of each Storage Management Service Class SCP Application.

The SCP requests a storage management operation using the N-ACTION-Request primitive of one of the Storage Management Service SOP Classes. The request includes a Transaction UID for tracking purposes.

If the SCP accepts the N-ACTION request for processing, it sends a Success N-ACTION Response Status Code. If it does not accept the N-ACTION request for processing, it sends an Error N-ACTION Response Status
1120 Code. The actions taken by the SCU upon receiving the status is beyond the scope of this Standard.

At any time after receipt of the N-ACTION-Response, the SCU may release the association on which it sent the N-ACTION-Request.

The SCP notifies the SCU using the N-EVENT-REPORT primitive when it has completed the requested operation or reached some other termination condition, such as a failure or a time-out. The notification includes the Transaction UID of the request. Upon completion or termination, the Transaction UID is no longer active and shall not be reused for other transactions.

ZZ.2 INVENTORY CREATION SOP CLASS

ZZ.2.1 Overview

The Inventory Creation SOP Class allows an SCU Application Entity to request a peer SCP Application Entity to produce an inventory of the studies, series, and instances that it manages. The request for the inventory may include a set of matching keys to limit the scope of the inventory to studies with corresponding attribute values. The inventory produced is encoded in one or more SOP Instances of the Inventory SOP Class (see Annex GG). Success or failure of inventory production is indicated by a notification from the SCP to the SCU. The SCU may then access the Inventory SOP Instances by a mechanism supported by the SCP, including Query/Retrieve (see Annex XX), Non-Patient Instance Services (see PS3.18 Section 12), or a non-DICOM file access protocol. (See additional explanatory information in Annex XXXX in PS3.17.)

ZZ.2.1.1 Inventory Production Statuses

The SOP Class defines five statuses for the inventory production with associated status transitions (see Figure ZZ.2-1). Some of the transitions are associated with receipt and processing of N-ACTION requests (labeled arrows in the figure).

.
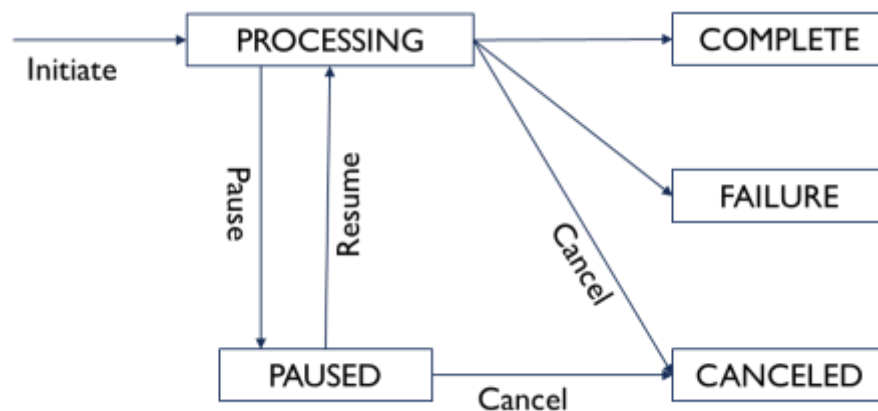
Figure ZZ.2-1 Inventory Production Statuses and Status Transitions

After the Initiate request, production is in the PROCESSING status. Production will proceed and reach one of two terminal statuses without further N-ACTION operations – COMPLETE or FAILURE.

Production may be PAUSED, either due to an N-ACTION request or due to conditions internal to the SCP, for example temporary shortage of resources. Production may be returned to PROCESSING status by the Resume request or SCP internal actions.

From either the PROCESSING or PAUSED status the production may be changed to the CANCELED terminal status by the Cancel request or SCP internal actions.

ZZ.2.2 Operations

ZZ.2.2.1 Action Information

The DICOM AEs that claim conformance to this SOP Class as an SCU and/or an SCP shall support the Action Types and Action Information as specified in Table ZZ.2-1. Descriptions of the attributes of the Action Information are specified in Section B.XX "Inventory Creation IOD" in PS3.3.

Table ZZ.2-1. Inventory Creation Requests - Action Information

Action Type Name	Action Type ID	Attribute Name	Tag	Requirement Type SCU/SCP
Initiate	11	Transaction UID	(0008,1195)	1/1
		Scope of Inventory Sequence	(00gg,0Fx0)	2/2
		>Study DateTime Range	(00gg,0Fy1)	2/2
		>Study Update DateTime Range	(00gg,0Fy2)	2/2
		>Study Description	(0008,1030)	2/2
		>Series Description	(0008,103E)	2/2
		>Modalities in Study	(0008,0061)	2/2
		>Study Instance UID List	(00gg,0Fy4)	2/2
		>Patient List Sequence	(00gg,0Fy6)	2/2
		>>Issuer of Patient ID	(0010,0021)	2/2
		>>Patient ID List	(00gg,0Fy5)	1/1
		>Attributes Missing Value	(00gg,0FyD)	2/2
		Inventory Purpose	(00gg,0Fx1)	3/1
		Inventory Level	(00gg,0Fx3)	1/1
		Requested Status Interval	(00gg,0Fy7)	3/1
Request Status	12	Transaction UID	(0008,1195)	1/1
		Requested Status Interval	(00gg,0Fy7)	3/1
Cancel	13	Transaction UID	(0008,1195)	1/1
		Retain Instances	(00gg,0Fy8)	1/1
Pause	14	Transaction UID	(0008,1195)	1/1
Resume	15	Transaction UID	(0008,1195)	1/1

.

ZZ.2.2.1.1 Scope of Inventory Sequence

The Scope of Inventory Sequence (00gg,0Fx0) specifies constraints on Studies to be included in the inventory.
See the attribute descriptions for the Scope of Inventory Sequence in Section C.YY.2.1 in PS3.3.

The attributes are similar to Key Attributes used for attribute matching in the Query/Retrieve Service Class (see Section C.2.2.2). The attribute values provided in the N-ACTION shall be matched against the values of attributes of Studies managed by the SCP, in accordance with the requirements of this section, with matching Studies to be included in the produced inventory.

Study DateTime Range (00gg,0Fy1) and Study Update DateTime Range (00gg,0Fy2) specify a datetime range for studies to be included in an inventory. Study DateTime Range is matched to the combination of Study Date (0008,0020) and Study Time (0008,0030), and Study Update DateTime Range to Study Update DateTime (0020,121x). Each attribute has two values (VM=2), representing the beginning and the end of the time range. Either the first and/or the second value may be empty, in which case the datetime range has an unbounded
beginning or end, respectively. The datetime range is inclusive; Studies whose time is at or after the first value and before or after the second value match the range. If Study DateTime Range is specified for matching, but a Study does not have a value for Study Date, that Study is not included in the inventory. If Study Update DateTime is specified for matching, but a Study does not have a defined Study Update DateTime available for matching, the datetime range is matched to Study Date and Time as a default.

1175 Study Description (0008,1030) and Series Description (0008,103E) specify strings to be matched to the respective attributes in studies to be included in an inventory. The characters "*" and "?" are interpreted as wild cards for matching, as described in Section C.2.2.2.4. A match of the Series Description to that attribute in any Series of a Study is a successful match for the Study to be included in an inventory.

Modalities in Study (0008,0061) provides a list of values to be matched against the Modality (0008,0060)
1180 attribute of Series within Studies to be included in an inventory. Multiple values may be present, and the equivalence of each and every value from this list with any Modality value among the Series of a Study is a successful match.

Notes

1. This is a logical 'AND' function for all values in this attribute; a value of CT\PT will match only Studies with both CT
1185 Series and PT Series. A value of MR will match all Studies with at least one MR Series among all the Study's Series, which may have other modalities as well (PR, SR, etc.).

2. There is no equivalent specification in Section C.2.2.2 for matching multiple values in a Request Identifier against multiple values in the equivalent attribute in the data set.

Study Instance UID List (00gg,0Fy4) provides a list of UIDs to be matched against the Study Instance UID
1190 (0020,000D) attribute of studies to be included in an inventory. Multiple values may be present. Matching is performed as described for List of UID Matching in Section C.2.2.2.2., i.e., the equivalence of the Study Instance UID with any value in the Study Instance UID List is a successful match.

Note

Studies are matched to the Study Instance UID List, not vice versa. That is, it is not a failure for there to be no matching
1195 Study for a value in the Study Instance UID List.

Patient List Sequence (00gg,0Fy6) provides a list of patient identifiers whose studies are to be included in an inventory. Matching must be successful to both the Patient ID (0010,0020) and the Issuer of Patient ID (0010,0021) attribute from the same Patient List Sequence Item, and wild cards shall not be used. Patient ID List (00gg,0Fy5) provides a list of values to be matched against the Patient ID (0010,0020) attribute of studies to
1200 be included in an inventory. Multiple values may be present. A match of any value in the Patient ID List to the Patient ID attribute is a successful match. Issuer of Patient ID (0010,0021) is matched to the same attribute for the Study. An absent value for the attribute in the Patient List Sequence Item is universal matching.

Attributes Missing Value (00gg,0FyD) specifies the tags of attributes that match the scope of inventory when the corresponding attribute in the Study has no value, i.e., studies for which the identified attributes have no value
1205 are to be included in the inventory. Multiple values may be present, and the absence of a value in any of the attributes identified in this list is a successful match.

Notes

1. This is a logical 'OR' function for all values in this attribute; a value of 00080050 00100020 will match Studies with either Accession Number or Patient ID empty.

1210 2. There is no equivalent specification in Section C.2.2.2 for matching to empty values in Study attributes.

An empty value for Modalities in Study, Study Instance UID List, or Patient ID List is interpreted as Universal Matching, as described in Section C.2.2.2.3. The Scope of Inventory Sequence may be zero-length to indicate universal match to all managed Studies (i.e., a complete inventory).

Studies that match all specified keys shall be included in the produced inventory.

1215 **ZZ.2.2.1.2 Inventory Level**

Inventory Level (00gg,0Fx3) specifies the types of records requested for inclusion in the produced Inventory SOP Instances. Inventories may be requested with Study level records, with Study and Series level records, or with Study, Series, and Instance level records.

ZZ.2.2.2 Service Class User Behavior

1220 The SCU shall use the N-ACTION primitive containing the well-known Storage Management SOP Instance UID (defined in Section ZZ.1.3.2) in its Requested SOP Instance UID parameter.

The SCU shall supply the Transaction UID (0008,1195) to uniquely identify each Storage Management Request. The value of the Transaction UID will be included by the SCP in N-EVENT-REPORT primitives (see Section ZZ.2.3.1). Use of the Transaction UID Attribute allows the SCU to match requests and results that may
1225 occur over the same or different Associations.

The SCU shall use the N-ACTION primitive with Action Type "Initiate" (Action Type ID = 11) to request the SCP to produce an inventory of instances managed by the SCP. The SCU shall supply the Inventory Level (00gg,0Fx3) to indicate whether the produced inventory should include records for series and instances. The SCU shall supply the Scope of Inventory Sequence (00gg,0Fx0) to specify constraints on Studies to be included
1230 in the inventory. The SCU may request the SCP to regularly report on the progress of producing the inventory associated with the Transaction UID at the interval specified by Requested Status Interval (00gg,0Fy7).

The SCU shall use the N-ACTION primitive with Action Type "Request Status" (Action Type ID = 12) to request the SCP to report on the progress of producing the inventory associated with the Transaction UID. The SCU may request the SCP to change the interval for progress reports as specified by Requested Status Interval.

1235 Note

Status of inventory production, including failure, will be signaled by the SCP via the N-EVENT-REPORT primitive.

The SCU shall use the N-ACTION primitive with Action Type "Cancel" (Action Type ID = 13) to request the SCP to stop producing the inventory associated with the Transaction UID.

The SCU shall use the N-ACTION primitive with Action Type "Pause" (Action Type ID = 14) to request the SCP
1240 to pause production of the inventory associated with the Transaction UID. It shall use the N-ACTION primitive with Action Type "Resume" (Action Type ID = 15) to request the SCP to resume producing the inventory from a PAUSED status.

ZZ.2.2.3 Service Class Provider Behavior

Upon receipt of the N-ACTION request, the SCP shall return, via the N-ACTION response primitive, the N-
1245 ACTION Response Status Code applicable to the associated request. A success status conveys that the SCP has successfully received the request. An error status conveys that the SCP is not processing the request.

If the SCP conveys a success status for Action Type "Initiate" (Action Type ID = 11), it shall initiate production of an inventory and periodic status reporting in accordance with the Scope of Inventory and Inventory Level provided in the N-ACTION request. Any value provided by the SCU for Inventory Purpose (00gg,0Fx1) shall be
1250 included in the produced Inventory SOP Instances. The SCP shall produce at least one Inventory SOP Instance upon reaching the COMPLETE status.

If the SCP encounters a failure when producing the inventory, it may attempt to complete the encoding into an Inventory SOP Instance of Study records already obtained, and set Inventory Completion Status (0400,06x5) to "FAILURE".

1255 If the SCP conveys a success status for Action Type "Request Status" (Action Type ID = 12), it shall initiate sending an N-EVENT-REPORT with Event Type "Inventory Status" (Event Type ID = 12).

Note

The N-EVENT-REPORT may be sent on the same or a different Association as the N-ACTION request (see Section ZZ.1.3.1).

1260 If the SCP conveys a success status for Action Type "Cancel" (Action Type ID = 13), it shall stop production of the inventory associated with the Transaction UID. If the Cancel request includes Retain Instances (00gg,0Fy8) value "N", the SCP may delete any Inventory SOP Instances associated with the Transaction UID. If the Cancel request includes Retain Instances (00gg,0Fy8) value "Y", the SCP shall complete the encoding into an Inventory SOP Instance of Study records already obtained, and set Inventory Completion Status (0400,06x5) to

1265 "CANCELED". Whether Study records partially complete at the time of the Cancel request are completed and included in the Inventory SOP Instance is implementation specific. Partial Study records shall not be included in an Inventory.

If the SCP conveys a success status for Action Type "Pause" (Action Type ID = 14), it shall pause production of the inventory associated with the Transaction UID. The SCP may complete any processing necessary to reach

1270 a stable state for pausing the inventory production that would allow resumption.

If the SCP conveys a success status for Action Type "Resume" (Action Type ID = 15), it shall resume production of the inventory associated with the Transaction UID.

Note

The SCP may return an error status Refused: Not authorized (0124H) or Resource Limitation (0213H) if the
1275 conditions that caused the Pause have not been corrected and the production is not resumed.

No SOP Class specific status values are defined for the N-ACTION Service. See Section 10.1.4.1.10 "Status" in PS3.7 for general response status codes.

ZZ.2.3 Notifications

ZZ.2.3.1 Event Information

1280 The DICOM AEs that claim conformance to this SOP Class as an SCU and/or an SCP shall support the Event Types and Event Information as specified in Table ZZ.2-2. Descriptions of the attributes of the Event Information are specified in Section B.XX "Inventory Creation IOD" in PS3.3.

Table ZZ.2-2. Inventory Creation Responses - Event Information

Event Type Name	Event Type ID	Attribute Name	Tag	Requirement Type SCU/SCP
Inventory Terminated with Instances	11	Transaction UID	(0008,1195)	1/1
		Transaction Status	(00gg,0FyB)	-/1
		Transaction Status Comment	(00gg,0FyC)	3/3
		Referenced SOP Class UID	(0008,1150)	1/1
		Referenced SOP Instance UID	(0008,1155)	1/1
		Retrieve AE Title	(0008,0054)	3/1C See ZZ.2.3.1.1
		Retrieve URL	(0008,1190)	3/1C See ZZ.2.3.1.1
		Stored Instance File URI	(00gg,0FxB)	3/3
		Container File Type	(00gg,0FxC)	3/1C See ZZ.2.3.1.1
		Filename in Container	(00gg,0FxD)	3/1C See ZZ.2.3.1.1

Event Type Name	Event Type ID	Attribute Name	Tag	Requirement Type SCU/SCP
		Stored Instance Transfer Syntax UID	(00gg,0Fy0)	3/1C See ZZ.2.3.1.1
		MAC Algorithm	(0400,0015)	-/3
		MAC	(0400,0404)	-/3
		Expiration Datetime	(00gg,0Fy9)	-/3
		Number of Study Records	(0400,06x6)	-/1
Inventory Status	12	Transaction UID	(0008,1195)	1/1
		Transaction Status	(00gg,0FyB)	-/1
		Transaction Status Comment	(00gg,0FyC)	3/3
		Number of Study Records	(0400,06x6)	-/1
Inventory Terminated without Instances	13	Transaction UID	(0008,1195)	1/1
		Transaction Status	(00gg,0FyB)	-/1
		Transaction Status Comment	(00gg,0FyC)	3/3

1285

ZZ.2.3.1.1 Inventory Terminated with Instances

Within the Inventory Terminated with Instances Event Information, either Retrieve AE Title (0008,0054) or Retrieve URL (0008,1190), or both, shall be present.

Stored Instance File URI (00gg,0FxB) provides a non-DICOM Protocol access method (see Annex P in PS3.3)
1290 to a produced Inventory SOP Instance encoded in accordance with the DICOM File Format (see Section 7 "DICOM File Format" in PS3.10). Stored Instance Transfer Syntax UID (00gg,0Fy0) shall be present if Stored Instance File URI is present. Container File Type (00gg,0FxC) and Filename in Container (00gg,0FxD) shall be present if Stored Instance File URI references a container format file.

ZZ.2.3.2 Service Class Provider Behavior

1295 The SCP shall use the N-EVENT-REPORT primitive containing the well-known Storage Management SOP Instance UID (defined in Section ZZ.1.3.2) in its Affected SOP Instance UID parameter. The SCP shall supply the Transaction UID (0008,1195) corresponding to the Inventory Creation N-ACTION Initiate request.

At any time during the creation of the inventory, the SCP may send an N-EVENT-REPORT primitive with Event Type "Inventory Status" (Event Type ID = 12), with an indication of process status in the Transaction Status
1300 Attribute. The SCP shall send an N-EVENT-REPORT "Inventory Status" if production of the inventory changes status. The SCP shall send an N-EVENT-REPORT "Inventory Status" in response to an N-ACTION "Request Status".

When production of the inventory reaches a terminal status (COMPLETE, FAILURE, or CANCELED) and Inventory SOP Instances are available, the SCP shall send an N-EVENT-REPORT primitive with Event Type
1305 "Inventory Terminated with Instances" (Event Type ID = 11). The Event Information shall include attributes specifying at least one method of accessing the root SOP Instance of the inventory, i.e., it shall include at least one of the attributes Retrieve AE Title (0008,0054) or Retrieve URL (0008,1190), and may optionally include Stored Instance File URI (00gg,0FxB). If Stored Instance File URI is present, then the Stored Instance Transfer Syntax UID (00gg,0Fy0) shall be present. Further, if Stored Instance File URI references a container format file,
1310 Container File Type (00gg,0FxC) and Filename in Container (00gg,0FxD) shall be present. The Event Information may also include an Expiration Datetime (00gg,0Fy9) attribute to indicate the expected time until which the Inventory SOP Instance is available for access by the SCU.

If production of the inventory reaches a terminal status (FAILURE or CANCELED) and Inventory SOP Instances are not available, the SCP shall send an N-EVENT-REPORT primitive with Event Type "Inventory Terminated without Instances" (Event Type ID = 13).

ZZ.2.3.3 Service Class User Behavior

Upon receipt of the N-EVENT-REPORT notification, the SCU shall return, via the N-EVENT-REPORT response primitive, the N-EVENT-REPORT Response Status Code applicable to its processing of the associated notification. A success status conveys that the SCU has successfully received the notification. An error status conveys that the SCU is not processing the notification.

> Note
>
> The action of the SCP upon an error status for the N-EVENT-REPORT is outside the scope of this Standard.

ZZ.2.4 Conformance

An implementation may claim conformance to this SOP Class as an SCU, SCP or both. The Conformance Statement shall be in the format defined in PS3.2.

ZZ.2.4.1 SCU Conformance

An implementation that is conformant to the Inventory Creation SOP Class as an SCU shall meet conformance requirements for

• the operations and actions that it invokes

• the notifications that it receives.

ZZ.2.4.1.1 Operations

The SCU shall document in the SCU Operations Statement

• the behavior and actions that cause the SCU to generate an N-ACTION primitive (Initiate, Request Status, Cancel, Pause, or Resume)

• the behavior and actions taken by the SCU upon receiving an N-ACTION error status

ZZ.2.4.1.2 Notifications

The SCU shall document in the SCU Notifications Statement

• the behavior and actions taken by the SCU upon receiving an N-EVENT-REPORT primitive (Inventory Terminated with Instances, Inventory Status, or Inventory Terminated without Instances).

ZZ.2.4.2 SCP Conformance

An implementation that is conformant to the Inventory Creation SOP Class as an SCP shall meet conformance requirements for

• the operations and actions that it performs

• the notifications that it generates.

ZZ.2.4.2.1 Operations

The SCP shall document in the SCP Operations Statement

• the behavior and actions of the SCP upon receiving the N-ACTION primitive (Initiate, Request Status, Cancel, Pause, or Resume).

- the range of values that it supports for the Requested Status Interval (00gg,0Fy7), and the default value for periodic reporting of status if that attribute is not provided by the SCU

- support for the Study Update DateTime attribute in managed Studies for matching on Study Update DateTime Range

    Note

    Conformance requirements for the SCP with respect to the created Inventory SOP Instances are detailed in Section GG.6.x.1.

ZZ.2.4.2.2 Notifications

The SCP shall document in the SCP Notifications Statement

- the behavior and actions that cause the SCP to generate an N-EVENT-REPORT primitive (Inventory Terminated with Instances, Inventory Status, or Inventory Terminated without Instances), and in particular the conditions that cause a transition to the PAUSED or FAILURE processing status and the generation of the associated N-EVENT-REPORT.

- the supported access mechanisms to Inventory SOP Instances referenced in the Inventory Terminated with Instances Notification, including the specific non-DICOM protocols (if any) referenced in the Stored Instance File URI

DICOM PS 3.6: Data Dictionary

Add new data elements to Section 6 Registry of DICOM Data Elements

Table 6-1. Registry of DICOM Data Elements

Tag	Name	Keyword	VR	VM	
(00gg,0Fx0)	Scope of Inventory Sequence	ScopeOfInventorySequence	SQ	1	
(00gg,0Fx1)	Inventory Purpose	InventoryPurpose	LO	1	
(00gg,0Fx2)	Inventory Instance Description	InventoryInstanceDescription	LO	1	
(00gg,0Fx3)	Inventory Level	InventoryContentLevel	CS	1	
(00gg,0Fx5)	Item Inventory DateTime	ItemInventoryDateTime	DT	1	
(00gg,0Fx6)	Removed from Operational Use	RemovedFromOperationalUse	CS	1	
(00gg,0Fx7)	Reason for Removal Code Sequence	ReasonForRemovalCodeSequence	SQ	1	
(00gg,0Fx9)	Stored Instance Base URI	StoredInstanceBaseURI	UR	1	
(00gg,0FxA)	Folder URI	FolderURI	UR	1	
(00gg,0FxB)	Stored Instance File URI	ContainerFileURI	UR	1	
(00gg,0FxC)	Container File Type	ContainerFileType	CS	1	
(00gg,0FxD)	Filename in Container	FilenameInContainer	UR	1	
(00gg,0FxE)	File Offset in Container	FileOffsetInContainer	UL	1	
(00gg,0Fy0)	Stored Instance Transfer Syntax UID	StoredInstanceTransferSyntaxUID	UI	1	
(00gg,0Fy1)	Study DateTime Range	StudyDateTimeRange	DT	2	
(00gg,0Fy2)	Study Update DateTime Range	StudyUpdateDateTimeRange	DT	2	
(00gg,0Fy4)	Study Instance UID List	StudyInstanceUIDList	UI	1-n	
(00gg,0Fy5)	Patient ID List	PatientIDList	LO	1-n	
(00gg,0Fy6)	Patient List Sequence	PatientListSequence	SQ	1	
(00gg,0Fy7)	Requested Status Interval	RequestedStatusInterval	US	1	
(00gg,0Fy8)	Retain Instances	RetainInstances	CS	1	
(00gg,0Fy9)	Expiration DateTime	ExpirationDateTime	DT	1	
(00gg,0FyB)	Transaction Status	TransactionStatus	CS	1	
(00gg,0FyC)	Transaction Status Comment	TransactionStatusComment	LT	1	
(00gg,0FyD)	Attributes Missing Value	AttributesMissingValue	AT	1-n	
(00gg,0FyF)	File Access Sequence	FileAccessSequence	SQ	1	
(0020,121x)	Study Update DateTime	StudyUpdateDateTime	DT	1	
(0400,06x0)	Incorporated Inventory Instance Sequence	IncorporatedInventoryInstanceSequence	SQ	1	
(0400,06x1)	Inventoried Studies Sequence	InventoriedStudiesSequence	SQ	1	
(0400,06x2)	Inventoried Series Sequence	InventoriedSeriesSequence	SQ	1	
(0400,06x3)	Inventoried Instances Sequence	InventoriedInstancesSequence	SQ	1	

Tag	Name	Keyword	VR	VM	
(0400,06x5)	Inventory Completion Status	InventoryCompletionStatus	CS	1	
(0400,06x6)	Number of Study Records	NumberofStudyRecords	UL	1	

1370

Add new UIDs to Annex A Registry of DICOM Unique Identifiers (UIDs)

Table A-1. UID Values

UID Value	UID Name	UID Type	Part
1.2.840.10008.5.1.4.xxuid.1	Inventory Storage	SOP Class	PS3.4
1.2.840.10008.5.1.4.xxuid.2	Inventory FIND	SOP Class	PS3.4
1.2.840.10008.5.1.4.xxuid.3	Inventory MOVE	SOP Class	PS3.4
1.2.840.10008.5.1.4.xxuid.4	Inventory GET	SOP Class	PS3.4
1.2.840.10008.5.1.4.xxuid.5	Inventory Creation	SOP Class	PS3.4
1.2.840.10008.1.20.1.x	Storage Management SOP Instance	Well-known Instance	PS3.4

1375   **Table A-3. Context Group UID Values**

Context UID	Context Identifier	Context Group Name	Comment
1.2.840.10008.6.1.nnn	CID 7xxx	Reason for Removal from Operational Use	

1380

DICOM PS 3.16: Content Mapping Resource

Add Context Group to Annex B DCMR Context Groups

CID 7XXX REASON FOR REMOVAL FROM OPERATIONAL USE

Type: Extensible

Version: 2021mmdd

UID: 1.2.840.10008.6.1.nnn

Table CID 7xxx. Reason for Removal from Operational Use

Coding Scheme Designator	Code Value	Code Meaning
DCM	113001	Rejected for Quality Reasons
DCM	113037	Rejected for Patient Safety Reasons
DCM	113038	Incorrect Modality Worklist Entry
DCM	113039	Data Retention Policy Expired
DCM	113680	Quality Control Intent
DCM	11yyy1	Data not accessible from storage
DCM	11yyy2	Processed Key Object Selection

Add terminology concepts to Annex D DCMR Context Groups

Table D-1. DICOM Controlled Terminology Definitions

Code Value	Code Meaning	Definition	Notes
11yyy1	Data not accessible from storage	Requested data cannot be read from storage media, e.g., data physically deleted, or storage media no longer available	
11yyy2	Processed Key Object Selection	A Key Object Selection SOP Instance, used for referenced object change management, that has had its directives processed and is no longer to be used (see [IHE RAD TF-1] *Image Object Change Management Profile*)	

DICOM PS 3.17: Explanatory Information

Add to Section 2 Normative References

2.2 OTHER REFERENCES

1400 [IHE RAD TF-1] IHE International. *IHE Radiology (RAD) Technical Framework, Volume 1 – Integration Profiles*
https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_TF_Vol1.pdf

[IHE RAD TF-2] IHE International. *IHE Radiology (RAD) Technical Framework, Volume 2 – Transactions*
https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_TF_Vol2.pdf

[RFC7233] IETF. June 2014. *Hypertext Transfer Protocol (HTTP/1.1): Range Requests*
1405 http://tools.ietf.org/html/rfc7233 .

Add explanatory Annex

Annex XXXX Inventories (Informative)

XXXX.1 THE DICOM DATA MANAGEMENT ENVIRONMENT

1410 DICOM data in a healthcare organization is generally managed in a Picture Archiving and Communications System (PACS), which supports a repository of current and historical studies, access to those studies through DICOM standard interfaces, and often workflow management for production and interpretation of studies. Historical images are now routinely retained "forever", and data set sizes are increasing with 3D/4D and multimodality studies. Repositories in many institutions now store over a billion instances with data volumes
1415 over one petabyte. Enterprise-scale management tools and data are required, including interoperability features that operate at large scales.

One important feature supporting repository management is the ability to obtain an inventory of the repository contents in a standard format.

XXXX.2 THE INVENTORY IOD

1420 ### XXXX.2.1 Overview

The Inventory Information Object Definition (Section A.XX "Inventory IOD" in PS3.3) specifies a structure capable of encoding an inventory of all studies, series, and instances in a repository. The IOD is structured hierarchically using sequence attributes – within the inventory is a sequence of study records, within each of which is a sequence of series records, and within each of those is a sequence of instance records. Each
1425 "record" is a set of key attributes describing the studies, series, and instance entities in a repository, and the mechanisms for accessing the stored SOP Instances.

The IOD entity-relationship model is shown in Figure XXXX.2-1. The Inventory is created by an identified piece of Equipment. The content of the Inventory follows the Study Root Query/Retrieve Information Model (see PS3.4 Section C.6.2.1), with Patient and Imaging Service Request information treated as attributes of the Study. The

1430 Imaging Service Request information entity is not explicitly modeled in other Composite IOD E-R models, but it is specifically identified here as its attributes, such as Accession Number, are typically important in repository management.

> Note
>
> 1435 There is a potentially complex relationship between the Study and Imaging Service Requests in the real world (e.g., see [IHE RAD TF-2] *Section 4.6.4.1.2.3 Relationship between Scheduled and Performed Procedure Steps*). However, the Inventory information model follows the basic Study information model and supports only a single Accession Number representing an Imaging Service Request (see Section C.7.2.1 in PS3.3). Note that if a Study has multiple associated Imaging Service Requests, the request attributes may be encoded at the Series level in the Request Attributes Sequence. The Inventory IOD includes the Request Attributes Sequence to support this use.
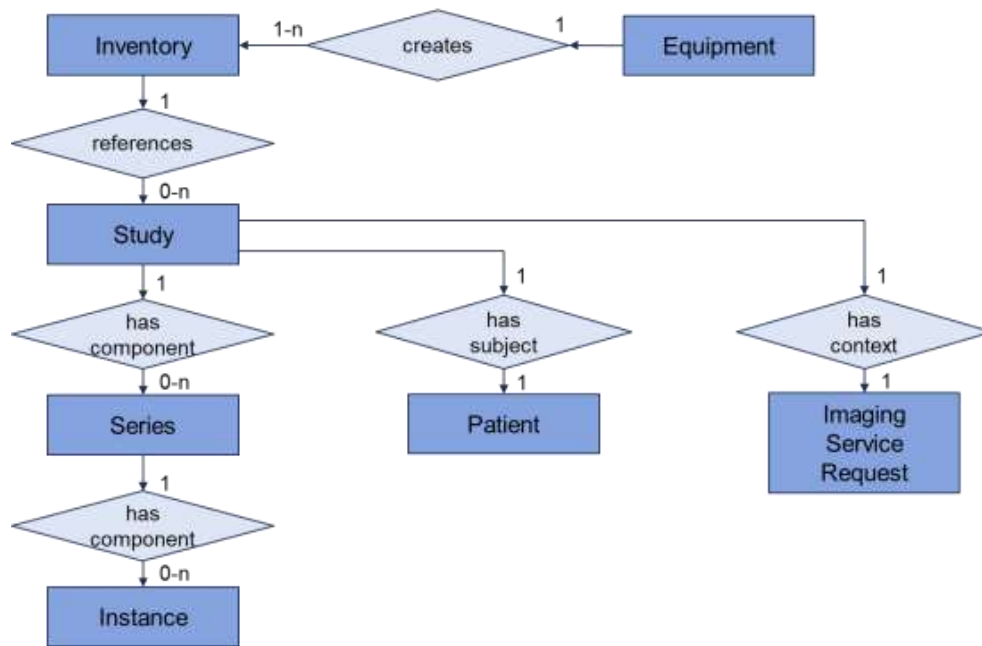
1440

Figure XXXX.2-1. Inventory Information Model E-R Diagram
(reproduced from Figure 7.13.x-1 in PS3.3)

These IOD Information Entities include all the required attributes specified for Query processing by the storage system (see Section C.6 in PS3.4). An Inventory is thus a standard DICOM representation of the key content of 1445 a storage system database for DICOM SOP Instances in the repository. Other aspects of such databases, such as data for workflow queues, are out of scope of the Inventory IOD.

The IOD allows inventories at three Inventory Levels – with only study records, with study and series records, or with study, series, and instance records. While most uses will require inventories with instance level records, production of a study or series level inventory may be significantly faster and may be sufficient for some uses.

1450 ## XXXX.2.2 Scope of Inventory

The Inventory IOD supports inventories of subsets of the repository based on a select number of key Study attributes that specify the scope of the inventory. Those key attributes (see Section C.YY.2.1 "Scope of Inventory Macro" in PS3.3) include Study Date and Time, Study Update Datetime, Modalities in Study, Study UID List, Patient ID List, and Attributes Missing Value. The scope is also implicitly limited to records available at 1455 the Content Date and Time when processing of the inventory began (see Section C.YY.1.1.1 in PS3.3), although records received or updated during inventory production may be included, and studies deleted during production might not be included.

A DICOM-conformant Inventory includes all studies that match the scope of the inventory, and their component series and instances, that are referenced by the storage system database. This includes studies, series, and
1460 instances that are marked in the database as removed from operational use, also commonly denoted as 'soft deleted', 'deprecated' or 'hidden' (see Section C.YY.1.2.2 "Removed from Operational Use" in PS3.3). Studies, Series, or Instances might be marked removed from operational use by actions associated with the processing of specific Key Object Selection SOP Instances, e.g., in accordance with [IHE RAD TF-1] *Image Object Change Management Integration Profile (IOCM).*

1465 Only the key attributes defined in the Standard for the scope of the inventory are allowed for Study selection. It is not permitted to extend the defined attributes of the Scope of Inventory Sequence to select Studies based on other attributes. However, an application that produces inventories might provide multiple Application Entities that implicitly scope to an implementation-specific subset of the repository (see Section XXXX.6.5).

XXXX.2.3 Inventory Instance Tree

1470 With a billion instances in a repository, the inventory itself may be on the order of 300 GB (i.e., $> 2^{38}$ bytes) in size. Producing and processing an inventory of such size may exceed some resource constraints of the creator and/or user application (such as 32-bit indices). The content of an inventory may therefore need to be divided into more than one SOP Instance.

> Note

1475 Because an inventory object has relatively low information entropy, compression of the inventory object may substantially decrease its size. Such compression may be applied to the Inventory SOP Instance using the Deflated Little Endian Transfer Syntax (see Section A.5 in PS3.5), or if the Inventory is stored in a DICOM File Format, the entire file may be compressed (e.g., using ZIP or GZIP). However, generally the instance needs to be fully constructed before it is compressed, or fully uncompressed before it is processed, and inventory applications need to be designed
1480 for the full potential size.

Very large repositories may also be partitioned or distributed into (semi-)independent subsystems. Production of an inventory for such distributed subsystems may be performed by parallel processes, which would be facilitated by each process producing a separate Inventory SOP Instance.

The Inventory IOD supports such cases of multiple SOP Instances comprising a single logical inventory. The
1485 IOD supports one SOP Instance incorporating the content of one or more others *by reference*. The IOD thus has the structure shown schematically in Figure XXXX.2-2. An Inventory SOP Instance may contain links to other Inventory SOP Instances whose content is incorporated by reference, or may contain inventoried study records, or both. A set of incorporated SOP Instances form a tree structure, with one SOP Instance at the root.
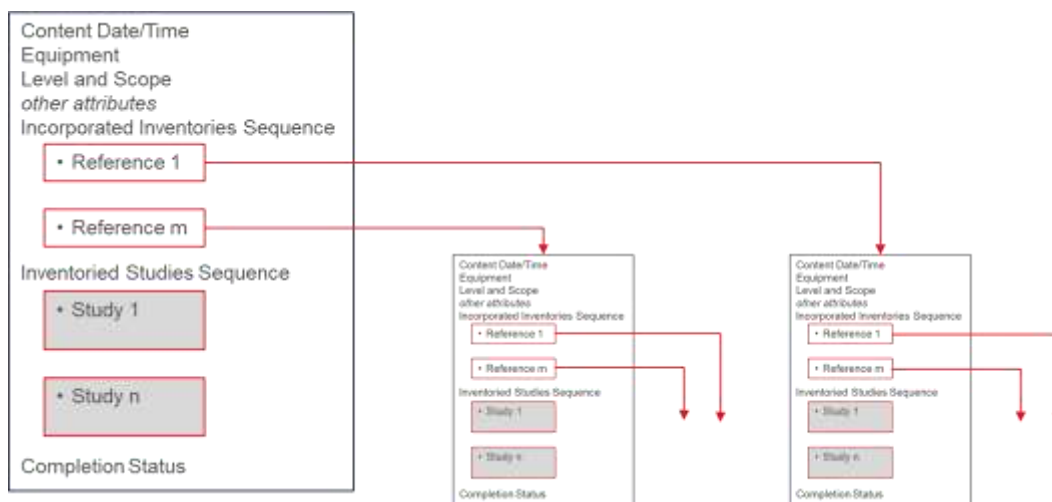

1490 **Figure XXXX.2-2. Inventory IOD Schematic Structure**
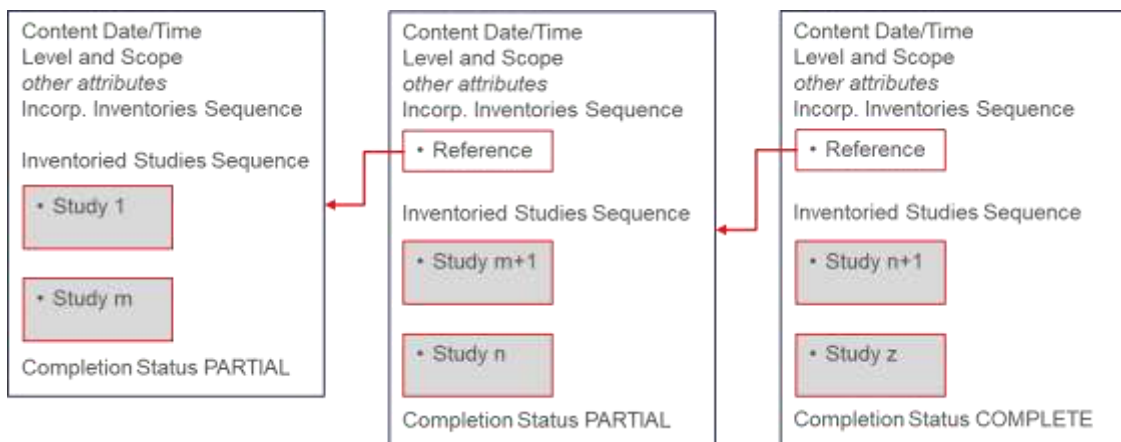
XXXX.2.3.1 Scope and Completion Status

Within any tree (or subtree) of Inventory SOP Instances, the root node specifies the Scope of Inventory and Completion Status for the entire tree, regardless of those attributes in subsidiary referenced objects. As will be seen in the examples, this is true regardless of the process used to create the inventory, whether with new objects, or with reference to previously created objects. The root object is the last SOP Instance to be completed in a tree, and it thus contains the final Completion Status for the tree and its Scope of Inventory. Any Completion Status value other than COMPLETE implies that the defined Scope of Inventory is not satisfied with this object as the root of the tree.

It is the responsibility of the creator of the root object for a tree to ensure that the Completion Status value accurately describes the content of the tree relative to the Scope of Inventory at the Content Date and Time for the system identified in the General Equipment Module.

XXXX.2.3.2 Examples

XXXX.2.3.2.1 Serial Production

As an example of how this tree structure might be used, consider an application producing a large inventory. It creates an Inventory SOP Instance, and begins filling it with inventoried study records. At some point, it reaches a size constraint, completes that object, and begins creation of a second object. That second object includes a link to the first one, and the application fills it with study records until it, too, reaches its limit. The process repeats with a third object, and so on until the inventory is complete (see Figure XXXX.2-3). The last object becomes the root of the tree of the complete inventory.

Figure XXXX.2-3. Serial production example

Note that in the first and second objects, the Scope of Inventory will be the same as in the final object, but the Completion Status of PARTIAL indicates that the sets of inventoried studies in their subtrees do not fulfil that Scope. (The subtree of the first object is just itself, the subtree of the second object is itself and the first, etc.)

XXXX.2.3.2.2 Baseline and increment

A special case of serial production is worth noting. A baseline inventory can be updated to current values by creating an Inventory SOP Instance with the incremental updates (new and changed study records) that includes the baseline Inventory SOP Instance by reference. The IOD allows a study to appear more than once in the tree of Inventory SOP Instances, and reconciliation of those records is facilitated by each appearance being tagged with its datetime of extraction from the database. Note that if an updated study is included in the incremental change inventory object, the full study record as known at that time needs to be encoded (not just records for new or changed series or instances in the study).
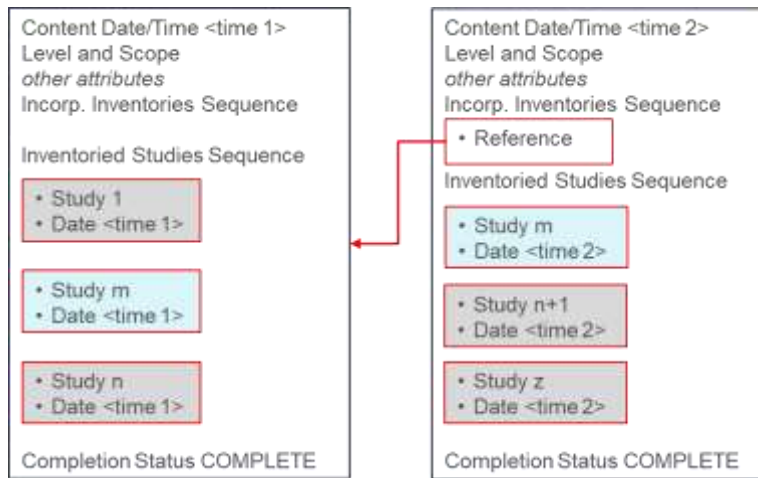
Figure XXXX.2-4. Baseline with incremental update

1525 In this example, the Scope and Equipment for each object are the same, but the objects differ by their Content Date and Time. It is the responsibility of the creator to ensure that study records in the incorporated inventory object are either current as of the Content Date / Time, even though their time of extraction precedes that datetime, or they are superseded by a current study record in the incremental update inventory object.

XXXX.2.3.2.3 Parallel Production

1530 As another example, consider an application producing an inventory in parallel across several independent federated storage subsystems. It tasks each subsystem to produce an Inventory SOP Instance, and itself produces a SOP Instance that links to each of the subsystem inventories (see Figure XXXX.2-4). Note that the Scope of Inventory will be the same for all objects, but the Equipment identifiers will differ.

1535 **Figure XXXX.2-5. Federated or parallel production example**

XXXX.2.3.2.4 Arbitrary tree structure

Combining these concepts, each of the parallel subsystems may produce an inventory which is itself a tree of Inventory SOP Instances. Each of those subtrees may follow the structures of either parallel or serial production. In general, the IOD supports an arbitrary tree structure (see Figure XXXX.2-5), where each node is the root of a
1540   subtree or a terminal leaf.
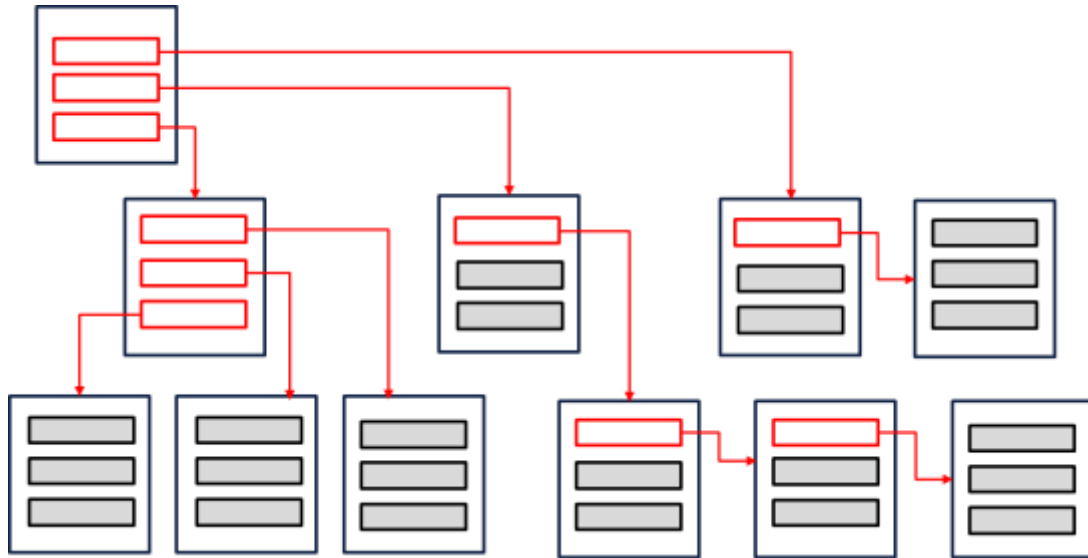
Figure XXXX.2-6. Arbitrary tree structure example
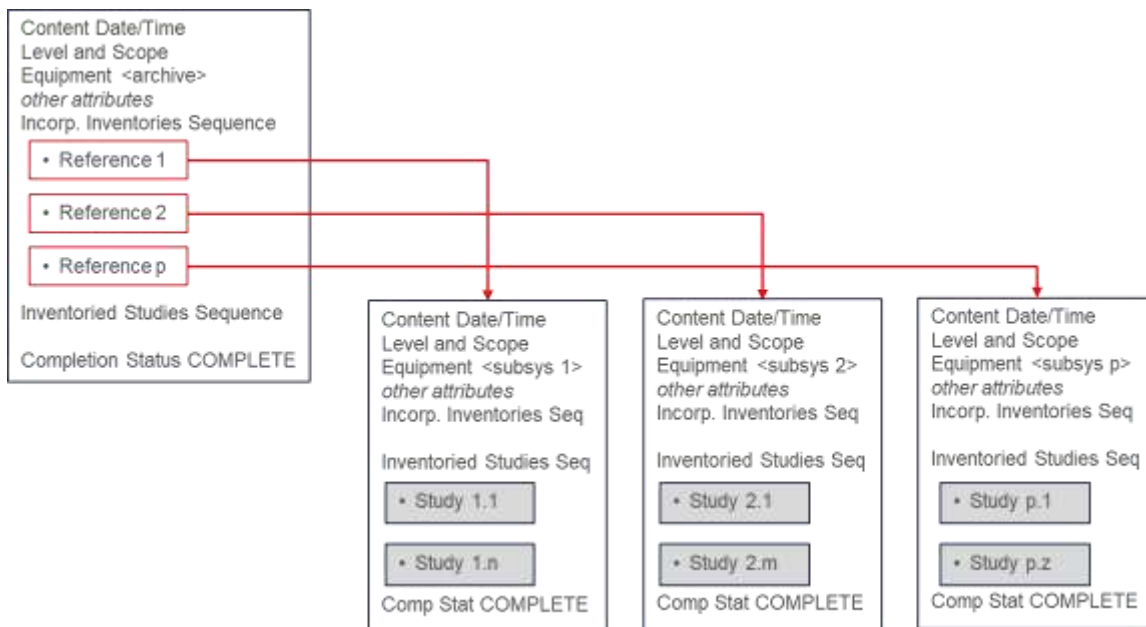
XXXX.2.3.2.5 Empty inventory

A storage system may be tasked with producing an inventory, but for which there are no stored studies that
1545   match the requested scope. For instance, an organization may be producing an inventory of all nuclear medicine studies, and requests each of its several PACS and VNAs to create an inventory with Modalities in Study = NM. A storage system that doesn't have any NM studies will create an empty inventory object, which affirmatively declares that that system does not have any matching studies as of the Content Date / Time.


```
Content Date/Time <date>
Equipment <equipment>
Scope <specified modality>
other attributes
Incorporated Inventories Sequence
Inventoried Studies Sequence
Number of Study Records <0>
Completion Status COMPLETE
```
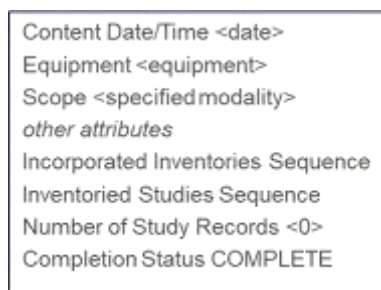
1550   **Figure XXXX.2-7. Empty inventory example**

XXXX.2.4 Access Mechanisms for Repository Data

The Inventory IOD supports the recording of available access mechanisms for each repository stored instance – DIMSE Query/Retrieve (Annex C in PS3.4), web services-based Studies Service Section 10 in PS3.18), and perhaps multiple non-DICOM direct file access protocols (Annex P in PS3.3). Either the access point for DIMSE
1555   (AE Title) or web (origin server address), or both, must be provided for an Inventory SOP Instance; the non-DICOM protocols are optional.

The IOD does not limit the non-DICOM protocol used, which is specified by the scheme of a URI link to the stored instance. Applications that intend to use direct file access may need to be adapted to use the protocol

specified by the Inventory creator. Not all capabilities of the access mechanism may be evident from the URI
1560 scheme, e.g., HTTP is used with several different cloud-based storage protocols which differ in the ways they
use HTTP headers. The specifics of the protocol must be conveyed in the Conformance Statement or other
documentation, rather than in the attributes of the IOD.

The target resource of a non-DICOM protocol must be a stored instance in the DICOM File Format (specified in
Section 7 in PS3.10), and may be further contained within a defined container file structure – ZIP, TAR, GZIP, or
1565 TARGZIP.

Many repository applications do not keep the stored instances updated with metadata that may change over
time (e.g., patient name and ID). Applications that use direct file access are required to use the current correct
metadata, as recorded in the Inventory SOP Instance, rather than the metadata in the stored files (see Section
XXXX.6.7).

1570 Although the Inventory identifies the available access mechanisms for repository stored instances, there will be
security features associated with those access mechanisms that are outside the scope of the IOD specification,
and that will need to be implemented in applications that use the Inventory (see Section XXXX.5).

XXXX.2.5 Additional Data Elements

The Inventory IOD, like all DICOM IODs, may be extended by the addition of optional attributes that do not
1575 impact the semantics of the basic IOD. This is denoted Standard Extended Conformance (see Section 3.11
"DICOM Conformance" in PS3.2).

While the IOD identifies many optional attributes that might be managed in a repository database, the creator of
an inventory is allowed to use such Standard Extended Conformance to include any additional data elements
that it manages. This may support additional use cases for the Inventory SOP Instances, or may provide direct
1580 database record keys in Private Data Elements for implementation-specific processing.

> Note
>
> For example, the repository database may support at the Instance level the Content Label and Content Description to
> support queries against Presentation States in accordance with the [IHE RAD TF-1] *Consistent Presentation of Images
> Profile,* or the Template ID and Concept Name Code Sequence to support queries against Structured Report SOP
1585 > Instances in accordance with the [IHE RAD TF-1] *Evidence Documents Profile.*

XXXX.2.6 Producer vs. Consumer Implementation

In all interoperability design, there is a tradeoff between ease of implementation for the producer of information
versus the consumer of that information. By adding constraints on the message content to which the producer
must adhere, the processing requirements for the receiver might be simplified. Fewer constraints on the
1590 producer means the consumer must account for more variability in the exchanged data.

In the design of the Inventory IOD, a policy was chosen to simplify the production of the SOP Instances, even at
the risk of complicating the implementation of the consumer. The goal is to allow the producer of the inventory to
simply report what it has, without substantial additional processing. For example, in a repository that might
distribute the SOP Instances of a study across multiple subsystems, each subsystem can report on the SOP
1595 Instances that it knows about, and there is no requirement for the producer of the combined inventory to
consolidate or reconcile those different records. For the migration and consolidation use case (see Section
XXXX.4.1), the consumer of the inventory will typically need to perform substantial reconciliation activities, which
do not need to be replicated in the producer.

This policy can also be seen in the approach to repository data that has been removed from operational use
1600 (deprecated, soft-deleted, or hidden). As there has been not an established DICOM standard approach to this
type of data, storage system implementations take a variety of approaches. The Inventory IOD does not attempt
to introduce a single way of managing such data. Rather, the system can simply report a removal status at the
level(s) at which it manages that status, be it study, series, or instance, with an optional reason code if it has
one. If the removal was due to a directive in a Key Object Selection SOP Instance, e.g., in accordance with the

1605 [IHE RAD TF-1] *IOCM Profile*, the Inventory IOD makes no assumption about the presence or status of that KOS object; the system simply reports whether it is stored in the repository.

XXXX.3 RELATED SERVICES

All Inventory-related network services will have associated security features that will need to be implemented in applications that use those services (see Section XXXX.5).

1610 ### XXXX.3.1 Inventory Storage and Query/Retrieve

The Inventory IOD is defined in the category of non-patient-related composite objects. As such, its basic SOP Class for DICOM network transfer is specified in the Non-Patient Object Storage Service Class (Annex GG in PS3.4 and Section 9.1.1 "C-STORE service" in PS3.7). Inventory objects may also be transferred using DICOM Media Interchange (Annex I in PS3.4 and PS3.10).

1615 Query/Retrieve of Inventory SOP Instances is specified in the Inventory Query/Retrieve Service Class (Annex XX in PS3.4). Query/Retrieve of Inventory SOP Instances uses the same C-FIND, C-MOVE, and C-GET DIMSE services as other Query/Retrieve Service Classes.

> Note
>
> Be careful to distinguish between Query/Retrieve of Inventory SOP Instances and Query/Retrieve of the SOP
1620 > Instances in the repository that are referenced in the Inventory.

Inventory Query returns key information about available Inventory SOP Instances, including Content Date and Time, Scope of Inventory, and completion status. This allows the Query SCU to obtain a list of available Inventory objects and determine whether any of them meet the SCU's needs, rather than initiating creation of a new inventory.

1625 Inventory SOP Instances may also be exchanged using DICOM web-based (HTTP) services. The equivalent of the Storage and Query/Retrieve Services is specified for the web through the Non-Patient Instance Services (see Section 12 in PS3.18).

Due to the potentially very large size of Inventory SOP Instances, the creator may make them available through a non-DICOM file access protocol. Such a protocol may allow interactive reading of files, rather than transfer as
1630 a whole to the destination system (see Section XXXX.6.4).

XXXX.3.2 Inventory Creation Service

Creation of an inventory may be initiated by a transaction of the Inventory Creation SOP Class (see Section ZZ.2 in PS3.4). The initiation action for the inventory specifies the requested scope of the inventory.

The Inventory Creation SOP Class is in many ways similar to the Study Root Query/Retrieve Information Model
1635 – FIND SOP Class (see Annex C in PS3.4). In both cases, the SCU requests a list of Studies, managed by the SCP, that match key values. However, the Study Root FIND is designed for routine daily department operational workflows, intended to identify a relatively few studies at a time. The Inventory Creation SOP Class is designed for operations on a repository as a whole, potentially identifying every study in a repository.

Thus, the Study Root FIND operates synchronously (i.e., the query and response occur on the same
1640 Association), in the expectation that the process can be completed in a short time to support critical end-user needs. In particular, it is expected that this SOP Class may be supporting a time-critical human user interface, e.g., a radiologist needing information to read a study. Therefore, the Study Root FIND SOP Class allows the SCP to limit the number of query responses to constrain the demand on processing and network resources, and return responses expeditiously.

1645 In contrast, the Inventory Creation SOP Class operates asynchronously, as production of an enterprise-scale inventory of billions of objects may take considerable time (potentially many days). As an asynchronous process, multiple approaches are available to the SCP to manage the resource demands for inventory

production across a longer time scale and with non-critical priority. The results are stored in information objects that can be accessed asynchronously at the convenience of the SCU.

1650 The mechanisms of the Inventory Creation SOP Class are similar to those of the Storage Commitment SOP Class (see Section CC). The SCU sends a request for the service to the SCP in an N-ACTION message, and the SCP asynchronously reports back status or completion using an N-EVENT-REPORT message.

The Inventory Creation SOP Class provides for regular reports on the status of inventory creation, at an interval specified by the SCU (see Section ZZ.2.2.2 in PS3.4). This allows the SCU to ensure that the operation has not
1655 stalled. For example, such reporting might be desired for each 5% of process progress, and for an inventory that is expected to complete in one day, status reporting could be requested for 30-minute intervals. The SOP Class also allows the SCU to request a status report update at any time.

The Inventory Creation SOP Class allows production of an inventory to be paused and resumed. A pause may occur when resources necessary for inventory production (database processing cycles, disk storage space, etc.)
1660 become temporarily unavailable, or when resource usage has reached a pre-set limit. For example, a system that allows a research application to create an inventory might limit the initial result to some maximum number of studies, and then pause for confirmation before proceeding. It is expected that some human intervention may be required before resuming inventory production.

The Inventory Creation SOP Class does not define conformance to subsets of its features. An application that
1665 implements the Inventory Creation SOP Class must therefore support matching on all of the defined Scope of Inventory keys and Content Levels.

Note that the Inventory Creation SOP Class does not use the Inventory IOD (Section A.XX "Inventory IOD" in PS3.3), but rather the Inventory Creation IOD (Section B.XX "Inventory Creation IOD" in PS3.3), which consists of the controls and statuses for production of an inventory. However, both the Inventory IOD and the Inventory
1670 Creation IOD use many of the same attributes, including the Scope of Inventory Sequence.

XXXX.3.3 Separability of Services

Each defined SOP Class is a separate DICOM Conformance claim for an implementation. Generally, an implementation may implement any of the Inventory-related services without implementing others.

Thus, a producer of inventories may choose any method for exchange of Inventory instances. It could support
1675 DIMSE Inventory STORE (with or without Inventory MOVE or Inventory GET), DICOM web services Retrieve, or DICOM Media exchange, and may additionally support a non-DICOM file access protocol. However, as all DICOM Conformance is to SOP Classes, an implementation cannot claim DICOM Conformance just to the Inventory IOD; it needs to claim conformance to at least one SOP Class that exchanges the Inventory SOP Instances.

1680 Identification and location of Inventory instances may be supported by the Inventory FIND SOP Class or the equivalent DICOM web search service, or may be done by non-DICOM means (e.g., email notification of inventory UIDs or filenames to a client). Similarly, an application may produce inventories under control of its local administrative user interface, and is not required to implement the Inventory Creation SOP Class for remote clients. However, if the producer does implement the Inventory Creation SOP Class, it must also
1685 implement a DICOM method for accessing the produced Inventory instances.

Figure XXXX.3-1 illustrates the relationships of the Inventory-related services to the information definitions.
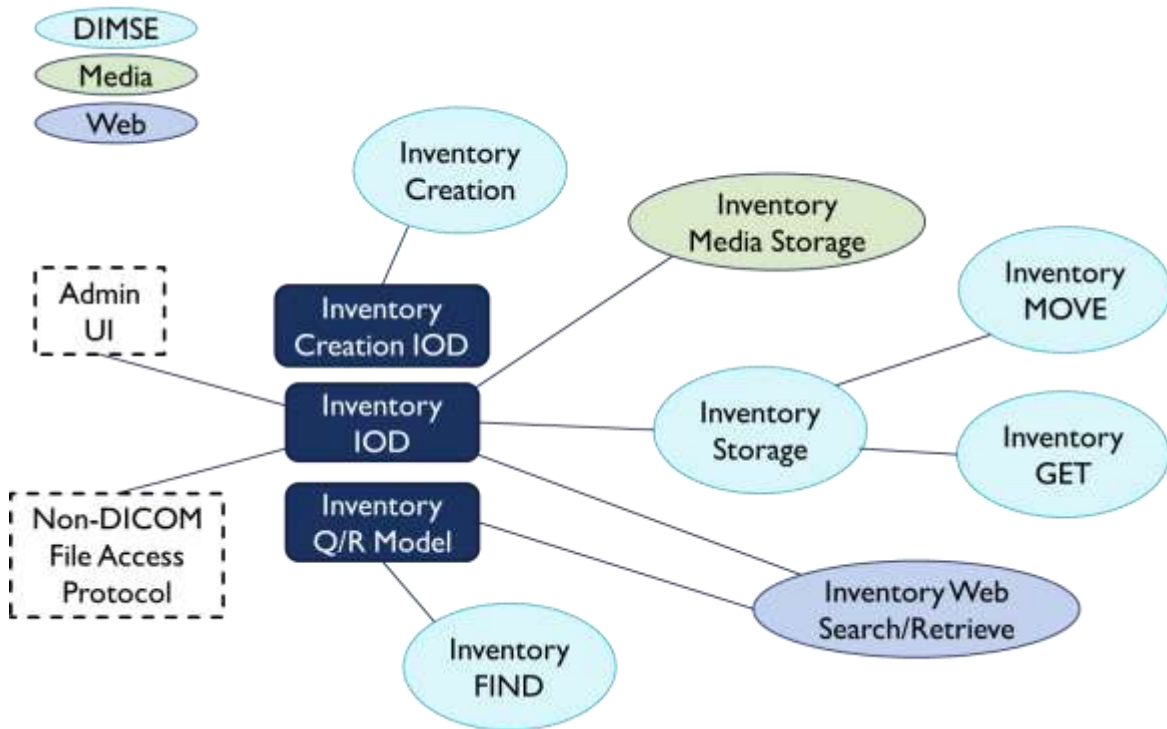
Figure XXXX.3-1 Inventory-related information definitions and services

As noted above, an application implementing the Inventory Creation SOP Class must support matching on all of
1690 the defined Scope of Inventory keys and Content Levels. However, an application that does not support the full
set of those capabilities, e.g., it can only create inventories at the Study level, may create restricted Inventory
SOP Instances under control of its administrative user interface. Such an application may claim conformance if
those SOP Instances are conformant to the IOD and it exchanges them using the Inventory Storage SOP Class.

XXXX.4 USE CASES

1695 ### XXXX.4.1 Migration and Consolidation

A use case of increasing significance is wholesale transfer of large DICOM repositories from one image
management system to another, denoted migration. As a regular part of managing IT obsolescence, users may
replace their image management system after about 12-15 years, often with change of vendor and underlying
hardware. Replacement requires migrating historical data to the new system. Similar transfer needs arise when
1700 healthcare institutions merge previously disparate systems into an enterprise image management system; the
repositories from the old systems need to be migrated.

The process of migration involves multiple phases or steps, of which an early task is obtaining an inventory of
the source repository. This step is directly addressed by the Inventory IOD and related Services. Additional
steps may include data reconciliation between the source repository and the databases of the radiology
1705 information system (RIS), electronic medical record system (EMR), hospital information system (HIS), and/or
master patient index (MPI).

A subsequent step in migration is extracting the DICOM data from the source system and transferring it to the
destination system. There are two significant challenges with this data movement. First is the volume of data to
be migrated, which as noted above may be a petabyte or more. Second, migration often occurs when either the
1710 source system or the destination, or both, are in clinical operation. Systems designed and configured to handle
the throughput of regular operations might not have capacity for the additional massive input/output
requirements of migration.

The Inventory IOD indirectly supports this data movement. Many repositories store their DICOM data in the DICOM File Format (as defined in Section 7 in PS3.10), and can provide a non-DICOM direct file access
1715 protocol. By bypassing the DICOM protocol processing to access these files, significantly higher transfer rates can often be achieved, and there may be less impact on the resources required to support ongoing clinical operations. The Inventory IOD allows identification of an available non-DICOM file access protocol for each SOP Instance in the repository.

Non-DICOM file movement may be further streamlined if the SOP Instances of a Study or Series are combined
1720 into a single container file (ZIP or TAR). The Inventory IOD may identify such container files at the Study or Series levels.

At the destination repository, the process of building the local database for the incoming data may be facilitated by processing the Inventory SOP Instances, rather than parsing the migrating data one object at a time. Image management systems commonly also require order (or imaging service request) information to be received prior
1725 to imaging data for the most efficient integration of new data into the database; the inventory may be processed to provide that data up front before the bulk data transfer is started.

A final step of verification of the migration, ensuring that all data has been transferred, may also use the inventory. In particular, as an initial check the count of the number of series and instances in a study could be compared between the inventories of the source and destination systems.

1730 ## XXXX.4.2 Safety backup

Functions critical to the healthcare mission of an organization, such as access to archived images, should be designed to minimize single points of failure, such that there are multiple paths to accomplish the function under failure or emergency situations. Such reliable access to the images is a key element of patient safety, ensuring timely access to information needed for clinical decisions and treatments.

1735 While the database management systems used by image management systems typically have fault tolerant designs, such as redundant online storage and offline backups, the data is in a proprietary format and dependent on the DBMS software for effective use. The DBMS itself therefore becomes a single point of failure, and can become inoperable, for instance, if a license key expires, or if it is subject to a malware attack.

Note

1740 Malware, and in particular ransomware attacks, may initially seek to disable known DBMS backup mechanisms before attacking the main target, thus preventing alternate recovery mechanisms. DICOM Inventory objects are less likely to be known as a target for malware attacks, or they can be sequestered in an off-line system not accessible to attack.

The Inventory SOP Instances can be used as a DBMS-independent replica of the critical data content of the database for the DICOM SOP Instances it manages. And further, if the repository instances are in DICOM File
1745 Format and referenced in the inventory, there is the possibility of a complete alternate path to access the images in the event of an image management system failure (although certainly not as efficiently as if the system were operational).

There are many ways such a regular safety backup inventory could be organized, using combinations of complete checkpoint inventories, incremental date range update inventories, partition-based inventories, patient-
1750 based inventories, and more. The appropriate approach will vary by the particular needs and workflow of each organization.

XXXX.4.3 Research

While imaging data may be important for research activities, it is rarely used solely by itself. It is generally used in conjunction with other aspects of the patient medical record – diagnoses, treatments, outcomes. Thus,
1755 support for imaging related research needs to support integrated activities with other healthcare informatics systems and data.

Research functions must also not impact ongoing healthcare operations. Data for research is therefore typically extracted from clinical operational systems and transferred to a separate server, often with patient de-identification or anonymization. These systems are sometimes denoted a "data warehouse", an extract of operational data that can be sorted, filtered, and analyzed in any number of ways to support research questions.

The Inventory IOD might thus support research use cases in several ways. In the broadest sense, since it is a representation of the imaging repository database, it can be used for imaging research in conjunction with the image instances and the medical record data. As a DICOM object, it can be transferred to other systems for further processing. Since the data is in a standard format, it can be processed using readily available tools without having to know the proprietary table layouts of the image management system database. And as the inventory has links to the stored SOP Instances, further drill down to the image instances and more detailed metadata is facilitated.

A complete inventory might be used for research purposes, especially if it has already been extracted for other purposes (such as safety backup). Such an inventory may have its data transformed, de-identified, and loaded to a data warehouse. But a more focused inventory might be produced for specific research processes. In particular, if searches of an EMR or data warehouse produces a census of candidate patients, an inventory of studies for just those patients may be created using List of Patient IDs in the Scope of Inventory, and the inventory content could be further constrained, for instance, by Modalities in Study or Study Description. It should be noted, however, that the filtering of studies by the defined Scope of Inventory is not sufficient for most research purposes; it may be sufficient as a first level selection tht requires additional filtering by other processes.

In most research uses, data sets must be de-identified. However, as the inventory must typically be linked (via Patient ID) to other patient medical records, care must be taken in processing of inventories for research to ensure de-identification. The approaches will vary depending on the specific research questions and data used.

The Inventory IOD is defined with the data elements necessary to support the primary use case of migration. However, the image management system may manage additional attributes at the Study, Series, or Instance levels that might be beneficial for research, and that could be included in the Inventory as Standard Extended Conformance.

XXXX.4.4 Quality Assurance

Certain study attributes provide linkage to other aspects of the patient medical record. In particular, Patient ID links the study to the medical record, Study Date allows correlation to other patient medical events, and Accession Number links the study to the relevant imaging order and study workflow. However, DICOM specifies these three critical attributes as Type 2 in composite SOP Instances, and they might therefore be empty in studies in the repository.

As a general quality assurance principle, but especially during migration, it is important that these attributes have correct values, but there is no way to construct a query for empty attributes using DICOM Standard queries (the Study Root Information Model FIND SOP Class). However, the Inventory Creation SOP Class and its Scope of Inventory allows study selection based on the absence of these attributes, and the Inventory Level allows limiting the inventory to Study level records. This allows creation of an inventory with only the Study level records with empty attributes. As a quality assurance process, such inventories may be produced on a regular basis and identified studies corrected as needed.

XXXX.4.5 Wellness Check

With all healthcare critical IT systems, and especially with enterprise scale systems, periodic checks for abnormal functioning are warranted. This includes not only monitoring and evaluation of error logs, but also active probing for fault conditions. In the context of an enterprise image data repository, this could include comparison of real-time storage system query responses with expected results, e.g., as recorded in a prior inventory. It might similarly include retrieving a sample set of studies using DICOM protocols and comparing the results with the same studies retrieved using a non-DICOM protocol recorded in the inventory. This use case aligns with current trends in continuous testing of cloud- and premises-deployed applications.

1805 **XXXX.5 SECURITY CONSIDERATIONS**

XXXX.5.1 Access Control and Secure Transport

DICOM is not prescriptive with respect to user identification, authorization, access control, or secure transport. However, DICOM does provide enabling capabilities for security features (see Section D.3.3.7 in PS3.7), and specifies available profiles for some aspects of secure access and transport (see Annex B "Secure Transport
1810 Connection Profiles" in PS3.15). As DICOM deals with exchange of legally protected health information, every real-world deployment must address these security features through institutional policies, procedures, and technical mechanisms. The specifics will vary with the organization and the capabilities of the technical infrastructure, including DICOM applications.

Inventories may potentially include data on all patients within a healthcare organization. Unauthorized access to
1815 inventory objects may thus potentially be a data breach affecting all patients. The breadth of the inventory makes it of particular concern for access control and transport security, and may require special attention in the institutional security policies, procedures, and technical mechanisms.

The Standard describes the use of DICOM and non-DICOM protocols to access stored SOP Instances, both Inventory objects and DICOM data in the repository (see Annex P "Stored File Access Through Non-DICOM
1820 Protocols" in PS3.3). All such protocols support technical means for access control and transport security, which must be used in accordance with institutional security policies and procedures. Although the Inventory identifies the available access mechanisms, there are no data elements for storing access credentials, as placing them in the Inventory would present significant security vulnerabilities. Processes for a reading application to obtain access credentials must be handled by non-DICOM mechanisms.

1825 Access control mechanisms must also address audit logs for recording access to protected health information. Both the technical means of recording user identity and the organizational policies and procedures to effectively use those technical means need to be considered.

XXXX.5.2 File Format

The DICOM File Format has security considerations that will apply whenever that format is used, e.g., for the
1830 Inventory SOP Instances or the referenced DICOM SOP Instances in the repository. See Section 7.5. "Security Considerations for DICOM File Format" in PS3.10.

The ZIP and TAR container file formats, which are defined formats for DICOM data in the repository, are known to have vulnerabilities and to be the target of malware attacks. Implementations that create or read container files should utilize appropriate defenses and safeguards such as:

1835 - Virus scanners for container content
- Sandbox execution and processing
- Full format and content validation
- Overrun detection

Applications that store container files for later use by other systems should consider the environments of those
1840 systems. This means the scanning and validation should detect attacks against at least Windows, MacOS, and Linux operating systems and applications.

Container files should not contain any directly or indirectly executable content (see Section P.1.2 in PS3.3). Container content validation should include a test for any form of executable content and consider the detection of executable content to be a risk of malicious content. The presence of malicious content may indicate a
1845 security breach of the source system or other upstream system.

XXXX.5.3 Network Protocols

Aside from the access control and transport security concerns of DICOM and non-DICOM network protocols, each protocol may have additional vulnerabilities, and considerations and warnings related to the

implementation and use of the protocol. The specific details of any such considerations are outside the scope of
1850 DICOM.

An implementation that supports direct file access using non-DICOM protocols should incorporate mechanisms
mitigating the particular risks from those protocols. This includes supply chain protection for software
components, update and patching mechanisms, site-specific configuration differing from the default, and other
administrative issues.

1855 XXXX.5.4 Application Validation

Introduction of software applications into a healthcare organization IT network has the potential to open security
vulnerabilities, and must be managed in accordance with institutional policy preventing unapproved applications
being installed and obtaining access to patient data. Applications that deal with the inventory and with its linked
data (i.e., the entire DICOM repository) should be thoroughly validated with regard to appropriateness of data
1860 use, including ensuring patient data privacy, as well as conformance to the DICOM Standard.

As the Inventory provides links to stored SOP Instances that may not have been updated to current metadata
(e.g., Patient Name may have been corrected or changed after the Instance was stored), an application
accessing those files through a non-DICOM protocol needs to obtain the current metadata values from the
Inventory SOP Instance. Applications for which current metadata is required should be specifically validated to
1865 ensure current metadata is applied.

XXXX.5.5 Inventory Resource Use

Inventory production may consume significant system resources, so policies and system implementations must
assure that such activities do not adversely affect the clinical operations of the organization (denial of service).
This may involve special authorization for initiating broad inventories, and appropriate setting of software task
1870 priorities for the inventory application.

XXXX.5.6 Encryption of Data at Rest

An organization may have policies requiring encryption of data at rest (i.e., as stored in the files of the storage
system). Encryption both limits access to applications that have (securely) obtained the decryption keys, and
also ensures file integrity. DICOM specifies methods for secure (encrypted) files (see Annex D "Media Storage
1875 Security Profiles" in PS3.15 and Section 7.4 "Secure DICOM File Format" in PS3.10), and other file-based
encryption mechanisms might be employed by a storage system. However, issues such as key management
and distribution are implementation- and site-specific.

Of particular interest to Inventories, the URI link to a stored SOP Instance may point to a Secure DICOM File or
a file encrypted by another mechanism. There are no specifications regarding key management to access that
1880 file, but storing the key in the Inventory would present significant vulnerabilities, and would be an inappropriate
mechanism unless the inventory itself were encrypted. Processes for a reading application to access such
secured files must be handled by non-DICOM mechanisms.

XXXX.5.7 Message Digest

The integrity of a stored SOP Instance file (unencrypted) may be verified by a Message Authentication Code
1885 (MAC, also known as a message digest, hash, or cryptographic checksum) computed across the file. This value
may be recomputed whenever a file is accessed, and that value compared to a previously computed MAC to
assure that no changes have occurred to the file.

The Inventory IOD supports recording a MAC computed by the storage system (the writing application) for files
in the repository that will be accessible through a non-DICOM protocol. The file reading application can
1890 independently perform the MAC computation to assure integrity of the file as read or transferred.

XXXX.6 OPERATIONAL CONSIDERATIONS

XXXX.6.1 Using Referenced Inventories

Section XXXX.2.3 describes the tree of Inventory SOP Instances whose contents are included by reference in the complete inventory described by the root SOP Instance. As there is no requirement for implementation of
1895 the Inventory Query/Retrieve Service (see Section XXXX.3.3), a potential challenge for the user of the inventory is locating and retrieving all referenced inventories in the tree. The Inventory IOD therefore includes access information with each link to a referenced Inventory SOP Instance. This access may be through the DIMSE Inventory Query/Retrieve Service, the DICOM web-based Non-Patient Instance Service, or a non-DICOM file access protocol.

1900 The specification of the Incorporated Inventory Instance Sequence (0400,06x0), which provides the links to subsidiary SOP Instances, recursively includes itself (see Section C.YY.2.3 "Inventory Reference Macro" in PS3.3). This allows encoding of a tree structure. (A similar specification approach is used for the tree of observation records in DICOM Structured Reports.) Each Inventory SOP Instance includes the entire set of links for the tree of which it is the root.

1905 Therefore, when an application creates an inventory and includes another inventory by reference, it simply needs to add the access information to the referenced SOP Instance into the Incorporated Inventory Instance Sequence together with a copy of the referenced object's Incorporated Inventory Instance Sequence (see Figure XXXX.6-1). Note that including the entire tree of object references ensures that the tree is acyclic.
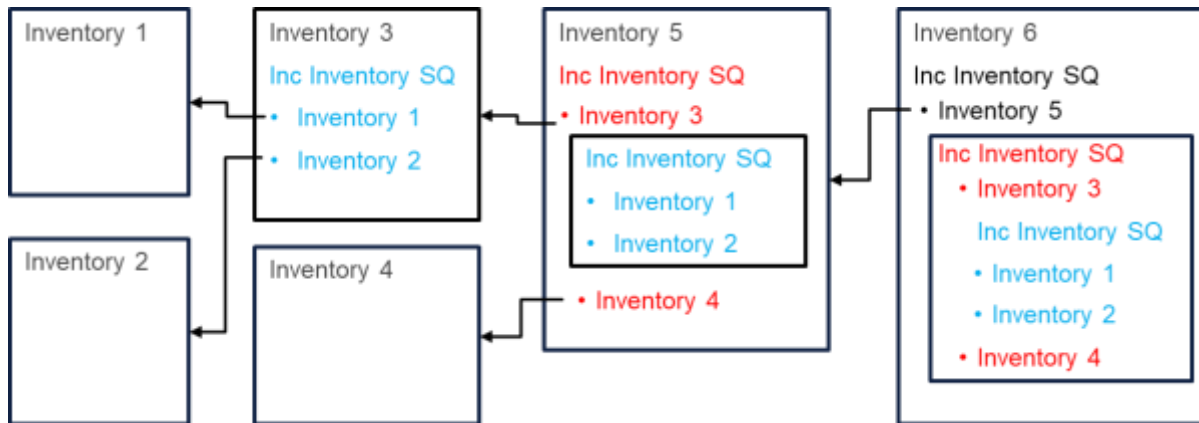
1910 **Figure XXXX.6-1 Inclusion of Inventory references**

The Inventory IOD also requires the SOP Instance to provide a count of the Number of Study Records (0400,06x6), which includes inventories included by reference. Since each Inventory SOP Instance computes the value for its tree, this simply means that when an instance includes others by reference the value is the sum of the Number of Study Records for each of its immediately referenced instances plus the number of Items in its
1915 own Inventoried Studies Sequence (0400,06x1).

XXXX.6.2 Incremental Inventories

Like all DICOM composite objects, Inventory SOP Instances are static, so an inventory of a repository in dynamic operation can never be complete. The challenge is to obtain a close enough approximation of completeness for the purposes of continuing work for the intended task, and then to obtain the incremental
1920 update since the prior inventory if needed.

The IOD design allows full inventory to be created with an incremental update inventory object including a baseline inventory by reference (see Section XXXX.2.3.2.2), thus minimizing the processing resource cost of a full inventory. However, there is no assumption that a creating application will utilize such an approach. If it does

not, the user application needs to request just an incremental update inventory to avoid the cost of creating a full
1925 inventory.

The attribute Study Update DateTime is intended to support obtaining an inventory of studies that have changed
since the time of the prior inventory. However, many repository implementations do not manage this attribute for
some (or any) of the stored studies. Although the desired functionality would be achieved by requesting a Scope
of Inventory with time range for Study Update DateTime beginning at the Content Date / Time of the prior
1930 inventory, the SCP might use Study Date / Time for the matched attribute as a fall-back when there is no
available Study Update DateTime. There are a number of reasons this is a poor approximation. First, there is an
inherent delay between the Study Date / Time (typically captured on the modality as the start of data acquisition)
and the time at which that study arrives at the repository; this delay may vary depending on the workflow of the
department (e.g., cardiology studies might be sent to the enterprise repository only after reading in department,
1935 with a 1-2 day delay). Second, studies may be updated with additional analytic or annotation series
(segmentations, presentation states, reports) well after the Study Date / Time. Third, studies received from
external organizations may have Study Date / Time significantly in the past, especially for imported prior exams.
Fourth, patient metadata may be updated much later than the Study Date / Time based on events totally outside
the imaging department. Fifth, there may be studies added to the repository that do not have a Study Date /
1940 Time (which are Type 2 attributes of the General Study Module).

The user of the inventory may need to mitigate this discrepancy by a variety of means in order to obtain an
inventory of the incremental changes to the repository. For instance, this may require adjusting the requested
time range to account for typical workflow delays and reconciling differences from the prior inventory. Some of
these methods may require data and processes outside the scope of DICOM, such as using external sources
1945 (e.g., audit logs) to identify imported studies and requesting inventory on those (using List of Study IDs in the
Scope of Inventory), or using external sources (e.g., HL7 ADT message logs) to identify changed metadata and
requesting inventory on those (using List of Patient IDs).

Comparison of the Number of Study Related Series and Number of Study Related Instances between a prior
and a current inventory may help identify studies that have changed content.

1950 ## XXXX.6.3 Inventory Lifecycle Management

Production and storage of inventory objects may use significant system resources, so effective system
management requires appropriate policies and controls on those services and objects to minimize necessary
resources. In addition to typical authorizations or permissions allowing specific users to create inventories, such
management policies may constrain when or how often inventories may be created, what scopes of inventories
1955 are permitted to which users, and when inventory objects should be deleted.

For instance, an organization that uses Inventory SOP Instances for safety backup (see Section XXXX.4.2),
may have policies to create a complete inventory each month, to maintain the two most recent inventories, and
automatically delete prior ones. Such a policy would allow the assignment of a value for Expiration Datetime
(00gg,0Fy9) for the Inventory SOP Instances.

1960 An organization may set a shorter retention period for Inventory SOP Instances associated with a canceled
inventory creation request.

A system that supports the Inventory Creation SOP Class (see Section XXXX.3.2) might reject requests that
duplicate the scope of an inventory recently created and that is available through the Inventory Query/Retrieve
Service. A system that produces a regular full inventory, e.g., monthly, might allow Inventory Creation requests
1965 only with a Study Update Datetime range after the last full inventory Content Date / Time.

XXXX.6.4 Interactive Access to Inventory Content

Inventory SOP Instances may be very large, and may reside on a server separate from the application that
needs to use them. The objects may be transferred to the using application using the DIMSE Non-Patient Object
Storage Service (see Annex GG in PS3.4), but that service transfers whole SOP Instances, and the using
1970 application may not require or want to store a whole inventory object.

If the origin and destination support the DICOM web-based Non-Patient Instance Service and Resources (see Section 12 in PS3.18), and if the origin supports HTTP Range request headers, the destination application can interactively retrieve specific byte ranges of the SOP Instance using the mechanism of [RFC7233].

1975 If the Inventory SOP Instances are made available through a non-DICOM protocol, that protocol may support interactive remote application reading of the file. Support for such protocols is typically integrated into the filesystem I/O capabilities of the using application's operating system.

XXXX.6.5 Multiple Application Entity Titles

A repository may have multiple Application Entities, with distinct DICOM protocol addresses (AE Titles). One common use is a PACS that has multiple separate archive subsystems, each of which supports DICOM protocol
1980 services (for example, as shown in Figure XXXX.2-4).

Another use for multiple AE Titles may be to provide separate views of the repository, and hence separate inventory content, for restricted subsets of the stored data. For example, the repository may include data that has patient consent for research use and data without such consent. This distinction does not have an associated key attribute for the Scope of Inventory. The system may therefore present one AE Title for
1985 operations on the entire repository, and a different AE Title for operations only on the research qualified data. This approach could be used for any other subsets of the repository that the system manages, but for which there is no standard key attribute for the Scope of Inventory.

A similar use of multiple AE Titles may provide separate views of the repository to different sets of users. An example of this is described in the next section for views of the Patient ID.

1990 ## XXXX.6.6 Multiple Patient IDs

The basic DICOM Patient Information Entity as used in the Inventory IOD supports a primary Patient ID, with an optional issuer or assigning authority, plus additional IDs and issuers in the Other Patient IDs Sequence (0010,1002). The DICOM attributes describing the assigning authority have mappings to corresponding HL7v2 CX Data Type components (see Section 10.15 in PS3.3).

1995 As PACS migration or consolidation often involves Patient IDs from multiple assigning authorities, organizations should establish well-defined assigning authority identifiers. The implementer of the inventory production application and the user organization should consider whether to include values for Issuer of Patient ID (0010,0021) in production of an inventory. Such values may especially facilitate consolidation of multiple repositories.

2000 Notes

1. See the [IHE RAD TF-1] *Scheduled Workflow.b Integration Profile* and its Enterprise Identity Option.

2. Similar considerations may be applied to the Issuer of Accession Number Sequence (0008,0051).

Some repository management systems, particularly those that support independent but related organizations, handle multiple Patient ID schemes. In such an environment, Query/Retrieve from applications in one
2005 organization may be returned with the Patient IDs for that organization, while the same queries from a different organization will have the Patient IDs for that second organization; the same approach may be used for production of inventories that have different views of the data for different users. To distinguish queries from the different organizations, the repository management system may use Application Entity Titles in two different ways. First, it may associate the SCU's Calling AE Title with an organization context; this requires the SCP to
2010 know all SCU AE Titles. A second approach has the SCP implement multiple Called AE Titles, each assigned to a different organization; each SCU is then configured to call the SCP AE Title appropriate to its organization.

If the SCP for the Inventory Creation Service (see Section XXXX.3.2) provides separate data views for different organizations, the name for the view may be encoded in the Station Name (0008,1010) or the Inventory Instance Description (00gg,0Fx2) attribute.

2015 XXXX.6.7 Metadata Updates

To maintain synchronization between the image repository and other electronic medical record systems, the PACS may support updates to patient, order, and procedure data that correspond to data managed by those other EMR systems. The PACS may also update series and instance level information as part of quality control processes. Examples of metadata updates include correction of patient name, change of patient ID, update of

2020 procedure descriptions or codes to a standard format, or correction of body part laterality. Such updates are managed by processes outside the scope of the DICOM Standard.

> Note
>
> See, for example, the [IHE RAD TF-1] *Patient Information Reconciliation Integration Profile*.

A common PACS design implementation stores received SOP Instances to disk in the DICOM File Format, but

2025 any metadata updates are retained in its database and not propagated to the stored instances. Applications that use non-DICOM protocols to access the files of stored SOP Instances must therefore also have access to current metadata.

The Inventory SOP Instance provides the current metadata for the stored instances, and the values in the attributes of the Inventory are considered authoritative. Therefore, the producer of the Inventory should ensure

2030 that it is created with current values, and the Item Inventory DateTime (00gg,0Fx5) records the time at which those values were extracted from the PACS database and were correct.

XXXX.6.7.1 Original Attributes Sequence

An optional additional capability is for the Inventory to record the provenance of metadata updates in the Original Attributes Sequence (0400,0561). While the current correct values are in the attributes of the

2035 Inventoried Studies Sequence, the Original Attributes Sequence records the prior (replaced) values, the datetime of the change, and the identity of the modifying system.

In Composite IODs, attributes of the Study, Series, and Instance levels are all encoded in the top level data set. The Original Attributes Sequence is defined in the SOP Common Module, and it aggregates all changes at any level of the information model. However, in the Inventory IOD the Original Attributes Sequence is defined

2040 separately at the Study, Series, and Instance levels, so that it can record updates at the higher levels without needing to replicate into the records for each referenced instance.

As an example, Table XXXX.6-1 shows what would be a portion of an Inventory SOP Instance for a study where the patient name was updated based on a master patient index, and one series was updated with a Body Part Examined that had been missing in the data received from the modality.

2045 **Table XXXX.6-1. Example Updated Study Record with Original Attributes Sequences**

Attribute	Tag	VR	Value	Comment
…				
Inventoried Studies Sequence	(0400,06x1)	SQ		
>Study Date	(0008,0020)	DA	20190506	
>…				
>Patient's Name	(0010,0010)	PN	Smith^Jane	current name
>…				
>Item Inventory DateTime	(00gg,0Fx5)	DT	20221103000450	
>Original Attributes Sequence	(0400,0561)	SQ		
>>Source of Previous Values	(0400,0564)	LO		unknown
>>Attribute Modification DateTime	(0400,0562)	DT	20190508110956	

Attribute	Tag	VR	Value	Comment
>>Modifying System	(0400,0563)	LO	GinHealthSystem PACS	
>>Reason for the Attribute Modification	(0400,0565)	CS	COERCE	
>>Modified Attributes Sequence	(0400,0550)	SQ		
>>>Patient's Name	(0010,0010)	PN	Doe^Jane	prior name
…				
>Inventoried Series Sequence	(0400,06x2)	SQ		
>>Series Date	(0008,0021)	DA	20190506	
…				
>>Body Part Examined	(0018,0015)	CS	LIVER	current value
…				
>>Original Attributes Sequence	(0400,0561)	SQ		
>>>Source of Previous Values	(0400,0564)	LO		unknown
>>>Attribute Modification DateTime	(0400,0562)	DT	20190508152157	
>>>Modifying System	(0400,0563)	LO	GinHealthSystem PACS	
>>>Reason for the Attribute Modification	(0400,0565)	CS	ADD	
>>>Modified Attributes Sequence	(0400,0550)	SQ		
>>>Body Part Examined	(0018,0015)	CS		prior value missing
…				

When updating the stored SOP Instance with the metadata values from the Inventory, the items of the Original Attributes Sequences at the Study, Series, and Instance levels from the Inventory are added to the items (if any) already in the Original Attributes Sequence of the stored SOP Instance. While there may be duplication, 2050 duplicate Items are not an issue for the audit purposes of the Original Attributes Sequence.

XXXX.6.8 Study Record Reconciliation

Within the tree of linked Inventory SOP Instances, a given Study may be referenced multiple times among the Inventoried Studies Sequence Items. The Items may have different content, but each Item is a complete record of the contents of the Study as known by the creator of that Item.

2055 Differences in content may occur due to changes to the metadata or content (SOP Instances) of the Study during the production of the inventory, or due to different Series of a Study being stored on different media or storage subsystems, or for other reasons. The application using an Inventory may need to reconcile such multiple occurrences.

DICOM is not prescriptive regarding methods of reconciliation, but the Inventory IOD does provide attributes that 2060 can assist in the process, in particular the various timestamps associated with the study content and the process of inventory creation, as shown in Table XXXX.6-2. These timestamp attributes might be used to establish a timeline of changes to study content and metadata, and of record extraction for inclusion in the inventory. For example, a study record may differ from a record with an earlier Item Inventory DateTime only with the presence of an additional Series whose Series Date is after the prior Item Inventory DateTime. The later record might 2065 reasonably be considered to be a more current replacement. However, two Study records might have entirely different sets of Series, and in that case simply choosing one record based on timestamp is probably not correct; the Study records would have to be further evaluated for the underlying reason for the difference, and the records potentially merged in some way.

Table XXXX.6-2. Timestamp Attributes Assisting in Reconciliation

Attribute	Tag
Content Date	(0008,0023)
Content Time	(0008,0033)
Inventoried Studies Sequence	*(0400,06x1)*
>Item Inventory DateTime	(00gg,0Fx5)
>Study Update DateTime	(0020,121x)
>*Original Attributes Sequence*	*(0400,0561)*
>>Attribute Modification DateTime	(0400,0562)
>*Inventoried Series Sequence*	*(0400,06x2)*
>>Series Date	(0008,0021)
>>Series Time	(0008,0031)
>>*Original Attributes Sequence*	*(0400,0561)*
>>>Attribute Modification DateTime	(0400,0562)

2070

In general, a major factor in reconciling diverse records is a full understanding of how the storage system manages the repository, and which timestamps and change auditing data it actually records. The reconciliation process will typically need to account for such system design features, which are not conveyed in Inventory SOP Instance attributes or in DICOM Conformance Statements.

2075 Note that a task for study record merge is reconciliation of access paths to stored SOP Instances of the study. This may present challenges if the study records link to different access methods, target folders, or container files. In the case of conflicting information, it may be necessary to disregard Study or Series level access specifications, and use only the access links to each SOP Instance of the study as recorded in the Instance level record.

2080 XXXX.6.8.1 Example – Deleted Study

An example will show the dependency on system design for study record reconciliation. Consider two inventories, a baseline made at time A and an increment made at a later time B, and during the intervening time a study is deleted (perhaps because it was assigned to the wrong patient). The migration source storage system might have taken one of several approaches, with the associated result in the time B inventory (this is not an 2085 exhaustive list):

1) It marks the study as deprecated, but otherwise retains the data – the time B incremental inventory includes the entire study record with the Removed from Operational Use attribute value Y.
2) It deletes the references to the Series and SOP Instances of the study in the database, retains the study level database record, but does not support a deprecation flag – the time B incremental inventory
2090 includes a study item, but no series items.
3) It deletes all study information – the time B incremental inventory simply does not record the study.

In case 1), the consumer application knows exactly what has happened, and can make a determination whether to move the deprecated study data to the migration target repository. That determination would be based, among other factors, on the data retention policies of the organization, and on the technical approach the target 2095 system takes to identifying and managing deleted studies.

In case 2), it might not be clear just from the content of the inventories what is the appropriate status of the study. This is further complicated if the SOP Instance files listed in the time A baseline inventory are still accessible from storage, perhaps indicating that the study was not supposed to be empty. If the consumer

application knows that this is the expected behavior of the source system for study deletion, it might proceed with migration in accordance with organizational policy. However, the application may need to consult external information, such as audit trails or human authorization, before proceeding.

In case 3), without an explicit study record somehow indicating deletion, the incremental inventory record for a deleted study is identical to a record for an unchanged study (i.e., no record in the inventory). The migration application would have no reason to suspect that the study was deleted until it tries to migrate the SOP Instances, and cannot find them. Studies that have gone missing are a patient safety issue, as opposed to studies that are known to have been deleted for a valid reason.

DICOM PS 3.18: Web Services

Add Inventories to Section 12.1.1 Non-Patient Instance Service and Resources / Resource Descriptions

12.1.1 Resource Descriptions

An NPI Service manages resources from the same NPI Category. Target URIs have the following templates:

```
/{npi-name}
```
2115 `/{npi-name}/{uid}`

Where

```
npi-name   = "color-palettes"
```
2120 `       / "defined-procedure-protocols"`
`       / "hanging-protocols"`
`       / "implant-templates"`
`       / "inventories"`
uid   ; is the Unique Identifier of an NPI Instance
2125

Table 12.1.1-1 contains the templates for the NPI Resource Categories.

Table 12.1.1-1. Resource Categories, URI Templates and Descriptions

Resource Category	URI Template and Description	Corresponding IOD	Storage Class	Information Model
Color Palette	/color-palettes{/uid}	Section A.58 "Color Palette IOD" in PS3.3	Section GG "Non-Patient Object Storage Service Class" in PS3.4	Section X.1.3 "Query/Retrieve Information Model" in PS3.4
Defined Procedure Protocol	/defined-procedure-protocols{/uid}	Section A.82 "Procedure Protocol Information Object Definitions" in PS3.3	Section GG "Non-Patient Object Storage Service Class" in PS3.4	Section HH.1.3 "Query/Retrieve Information Model" in PS3.4
Hanging Protocol	/hanging-protocols{/uid}	Section A.44 "Hanging Protocol IOD" in PS3.3	Section GG "Non-Patient Object Storage Service Class" in PS3.4	Section U.1.3 "Query/Retrieve Information Model" in PS3.4
Implant Template	/implant-templates{/uid}	Section A.61 "Generic Implant Template IOD" in PS3.3	Section GG "Non-Patient Object Storage Service Class" in PS3.4	Section BB.1.3 "Query/Retrieve Information Model" in PS3.4
Inventory	**/inventories{/uid}**	**Section A.XX "Inventory IOD" in PS3.3**	**Section GG "Non-Patient Object Storage Service Class" in PS3.4**	**Section XX.2 "Inventory Q/R Information Model" in PS3.4**

The NPI SOP Classes are listed in Table GG.3-1 "Standard SOP Classes" in PS3.4.

2130

Add Inventories to Section 12.6.1.2 Non-Patient Instance Service and Resources / Query Parameters

12.6.1.2 Query Parameters

2135 The user agent shall supply, and the origin server shall support, the Common Query Parameters in Section 12.1.2.

The origin server shall support Query Parameters as required in Table 8.3.4-1.

The user agent shall supply in the request Query Parameters as required in Table 8.3.4-1.

For each Resource Category the origin server supports, it shall support the behaviors and matching key Attributes specified in the corresponding sections in Table 12.6.1-2.

2140 **Table 12.6.1-2. NPI Resource Search Attributes**

Resource Category	Behaviors and Matching Key Attributes
Color Palette	Section X.6.1.2 "Color Palette Attributes" in PS3.4.
Defined Procedure Protocol	Section HH.6.1.2 "Defined Procedure Protocol Attributes" in PS3.4.
Hanging Protocol	Section U.6.1.2 "Hanging Protocol Attributes" in PS3.4.
Implant Template	Section BB.6.1.2 "Implant Template Attributes" in PS3.4.
Inventory	**Section XX.2.2 "Inventory Q/R Information Model Attributes" in PS3.4**