5

Digital Imaging and Communications in Medicine (DICOM)

Supplement 223: Repository Query, Inventory IOD, and Related Services

10

15 *Prepared by:*

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

1300 N. 17th Street, Suite 900

Rosslyn, Virginia 22209 USA

20

VERSION: Public Comment 2

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

This is a draft document. Do not circulate, quote, or reproduce it except with the approval of NEMA.

25 Copyright © 2021 NEMA

Table of Contents

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
26 – PC+1	14-Jun-2021	Add definitions to PS3.3; add Part 17 section on access control during production; editorial style fixes

27 – PC+2	24-Jun-2021	Editorial style fixes, conditionality of Removed from Operational Use
28 – PC+3	12-Jul-2021	Incorporate WG-6 feedback; add Study Update DateTime to Study Query/Retrieve Information Model; add Part 17 section on de-identification; clarify meaning of "non-patient"
29 – PC+4	30-Aug-2021	Clarify Study Update DateTime; add Alternate Representation Sequence; finish to do items
30 – PC+5	13-Sep-2021	Add empty value matching and multiple value matching to Query/Retrieve; draft of Inventory Query SOP Class for WG-33 review
31 – PC+6	28-Sep-2021	Revised to Repository Query, full draft including PS3.5 and PS3.7
32	4-Oct-2021	WG-6 status Review
33	26-Oct-2021	Revised Scope of Inventory to correspond to Repository Query Key Attributes
34	3-Nov-2021	For WG33 review prior to PC2; revised PS3.17
35	5-Nov-2021	For WG6 review prior to PC2; revised PS3.17
36	8-Nov-2021	Error code for invalid Prior Record Key
37	10-Nov-2021	Feedback from WG-6; new PS3.17 operational considerations subsections
38	11-Nov-2021	Feedback from WG-6
39	14-Nov-2021	Last WG-33 review before Public Comment 2
40 – PC2	17-Nov-2021	Add File Length in Container, additional Open Issues

310 # Scope and Field of Application

This Supplement introduces a new Repository Query SOP Class to obtain an inventory of a repository system, a composite Inventory IOD that is the equivalent persistent object instantiation of such an inventory, an Inventory Creation SOP Class to initiate asynchronous creation of Inventory SOP Instances, and SOP Classes to transfer/query/retrieve Inventory SOP Instances.

315 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 routinely retained "forever", and data set sizes are increasing with 3D/4D and multimodality
320 studies. Archives in many institutions 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 while 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
325 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
330 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,
335 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.

340 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 to an implementation-specified
345 number appropriate for a human interface, and the Standard is silent on behavior when that number is exceeded. 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,
350 limitations on number of responses and the synchronous protocol require the use of a possibly very large number of partial Query requests, with undefined behavior when query limits are exceeded. This makes producing such an inventory difficult with the current Query Service.

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

355 is needed.

Recognizing the varying technical implementations and business needs of repository manufacturers and of inventory users, this Supplement specifies two mechanisms to achieve a robust inventory production.

This Supplement specifies a new Repository Query SOP Class that includes features supporting a sequential set of Queries intended to produce a complete repository inventory. These features include defined behavior for

360 Queries that reach a system limit for number of responses, and an ability to resume at the next record in a subsequent Query. All key attribute selectors available in the current Query/Retrieve Service are also available in the Repository Query SOP Class, and some new selection mechanisms are defined that are useful for inventory production.

This Supplement also specifies an Information Object Definition capable of encoding an inventory of all studies,

365 series, and instances in a repository. This is functionally equivalent to a Query response that returns an inventory of the entire repository database, or a subset thereof as specified by key attributes. SOP Instances of this Inventory IOD could be produced upon local initiation by the repository administrative controls. The Supplement further defines a mechanism to remotely initiate the production of the inventory through a DICOM network service, and allow production to proceed asynchronously. The inventory SOP Instance(s) would be

370 available for transfer when production is complete through Store/Query/Retrieve mechanisms similar to those specified for other DICOM Non-Patient Object classes (such as Color Palettes).

Only inventory of patient-related studies, series and instances is defined. Inventory of non-patient objects is out of scope for this Supplement.

Since archive systems may optionally support direct filesystem access to DICOM Part 10 compliant files, for all

375 or some of their stored instances, the Repository Query and the Inventory IOD allow a URI 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, which may differ from the values in the stored instances.

While this Supplement defines an IOD and Services to support migration, it is not in itself a complete standard

380 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.

Open Issues for Public Comment 2

385 Supplement 223 was released for Public Comment during March-June 2020. Based on comments received at that time (see Public Comment 1 items 1.A, 1.B, 1.C, 5.A and 5.B below), Sup223 was revised to incorporate a new Repository Query SOP Class and change the approach to encoding the Scope of Inventory to align with that Query. WG-33 requests specific public comment on those changes in the revised draft, and in particular the following issues. Comment on any other aspect of Sup223 is also welcomed.

PC2.1. Repository Query SOP Class	
PC2.1.A	**Should we "raise the bar" for conformance to the new Repository Query?**
	With the definition of a new Query SOP Class, we have the opportunity to make conformance more rigorous and thus perhaps improve effectiveness of the service, but recognizing that any such requirements add burden to implementers. For the Repository Query SOP Class, the draft has minimized additional requirements. But considering the needs of the use cases, should we require support of new Key Attribute matching mechanisms (such as empty value matching), or should we require support for some additional Key Attributes (such as Study Update DateTime)?
PC2.1.B	**What should be the behavior for matching unsupported optional keys?**
	In DICOM Query, when the SCU requests matching on optional Key Attributes that are not supported for matching by the SCP, the baseline response behavior is to treat them as "universal match", i.e., no filtering is performed. Is this the behavior desired for the Repository Query? This may result in a substantial number of extraneous records being returned with performance impacts on both the SCU and SCP. Should requests for matching unsupported Key Attributes result in an error, allowing the SCU to perhaps change the query before proceeding into processing a response of many thousands of records? Should this alternate query response behavior be agreed through a new Extended Negotiation parameter?
	[For example, the SCU may request inventory of Studies updated in the last year by specifying date range matching on Study Update DateTime. If the SCP does not support matching on that attribute, the baseline behavior would be to return inventory of all Studies in the repository. See proposed PS3.17 Section XXXX.7.11]
PC2.1.C	**What should be the query behavior for federated repositories with different sets of supported Attributes and matching mechanisms?**
	Enterprise scale repositories now commonly include federations of multiple semi-independent archives (VNAs). Do we need specified behaviors for how the Repository Query should respond when aggregating responses across multiple federated archives? Or should this be left for the innovation of implementers (at least initially)?
	This is also an issue for the Inventory IOD, and resolution of this comment may require amendment of the IOD specification.
PC2.1.D	**Is there a need for a Repository Query to be able to return all known Attributes?**
	The standard operation of DICOM Query is for the SCU to only return the Attributes that were explicitly requested by the SCU. For the Repository Query, we have added an Attribute "Updated Metadata Sequence" that, if requested by the SCU and supported by the SCP, returns in the Sequence all Attributes with values that differ from those in stored SOP Instances (e.g., updated patient name). Some use cases would benefit from a similar capability that allows the SCP to return all Attributes that it manages in the database, without the SCU needing to enumerate them individually – especially since the SCU does not know what Attributes the SCP manages. Would this be an appropriate addition to the Repository Query?

PC2.2. Non-Patient	
PC2. 2.A	**Is Inventory properly placed in the Non-Patient Object Service Class?** For the purpose of transfer of Inventory SOP Instances, they are included in the category of "Non-Patient Objects" specified in Part 4 Annex GG and Part 18 Section 12, i.e., objects that do not have an individual Patient IE at the root of the IOD. Is this appropriate, as the Inventory does contain protected health information (PHI) about individuals, and the name Non-Patient Objects may imply to some that such objects have no PHI? Note that the name "Non-Patient Objects" or "Non-Patient Instances" occurs only in the descriptive text of the Standard, and does not occur in any network messages. Should we define a new DIMSE and DICOMweb Service Class just for Inventories? Or should we rename the Service Class, and if so, to what? Or is the clarifying text added by Sup223 sufficient?
PC2.3. Direct Access to SOP Instance files	
PC2. 3.A	**Are additional attributes needed for identifying characteristics of multiple versions of stored instances?** There are implementations that store multiple versions of a SOP Instance, e.g., that vary by compression ratio. What Attributes, if any, should be added to the Stored File Access Macro to enable the SCU to identify such variants?
PC2. 3.B	**Should generic octet stream (blob) be defined as a container?** There are implementations that concatenate Part 10 files into a single file or object without an internal structure identifying files (as is the case with TAR and ZIP). This is not in accordance with any defined standard, but is used by at least one major PACS implementation. Should support for this format be added?

390

Issues Closed Following Public Comment 1

During the initial Public Comment, WG-33 requested specific comment on the following issues. Resolution of each of these issues is noted.

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? *The key attribute matching for Scope of Inventory has been changed to align with the general key attribute matching mechanisms for Query C-FIND. It is sufficiently well specified to allow interpretation of arbitrary sets of key attributes.*

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?
	The use-specific mechanisms have been replaced with Query C-FIND equivalent key attribute matching mechanisms. However, two new matching mechanisms, multi-value matching and empty value matching, are being added and made available to all Query SOP Classes through Extended SOP Class Negotiation. This allows the general mechanisms to directly meet all intended use cases, except for list of Patient ID matching (for which there is a workaround using List of UID matching on Study UIDs).
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?
	While no additional use cases were raised that would call for matching on SOP Classes in Study, the change to the Scope of Inventory mechanism now enables matching on the optional attribute SOP Classes in Study (0008,0062) defined for Query (if supported by the SCP for matching).
	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)?
	No problematic use cases were raised. If study/series level optimization through folders is ineffective, a less optimized approach accessing individual SOP Instance files is always available.
2.B	**Should specific non-DICOM file access protocols be prescribed normatively?**
	Non-DICOM file access protocols are left open (PS3.3 Section P.2.2). 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)?
	No needs raised during Public Comment; additional specs can be added in future (or through IHE profiling). The provided URI is as adequate as URIs in any other interoperability context.
	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? *No use cases were raised during Public Comment that would call for an Inventory de-identification profile. A new section is added to the informative annex (PS3.17 Section XXXX.6.8) describing an inventory of a repository of de-identified Studies.*
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? *A new section is added to the informative annex (PS3.17 Section XXXX.6.1.1) to draw attention to this issue for implementers.*
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? *No comments received. Consultation with WG-14 Security identified no well-defined and commonly implemented attributes in PACS. If a need for standard attributes is identified, this may be addressed by a broader change to the Standard. Non-standard attributes may be addressed by private attributes.*
	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? *No comments received during Public Comment.*

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. *No use cases were raised during Public Comment that would call for ability to request a "fresh" inventory.*
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] *In the context of the expected use of Inventory SOP Instances, including access attributes is deemed sufficiently important to override concerns about potential staleness, and is not considered an undue burden on the producer.*
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) *No needs raised during Public Comment.*

395 During Public Comment, the following issues were raised, and their resolution is noted.

5.A	**Is a C-FIND based mechanism preferable to the SOP Instance mechanism?**
	The approach of defining persistent Inventory SOP Instances presents technical and business challenges for some repository applications. Those applications don't generate their own SOP Instances – they only store ones they are given, and implementation of production and management of Inventory objects is a substantial technical burden. And if the primary use case is to facilitate migration (i.e., to another vendor's system), there are also substantial business disincentives to implementation. However, such repository applications are designed to support C-FINDs, and a limited extension of C-FIND to be able to incrementally obtain a complete inventory would be a more tractable solution (both technically and from a business perspective). *WG-33 and WG-6 recognized the merit of the argument, but also the need for persistent Inventory SOP Instances for some use cases and models of use. A Repository Query SOP Class has been added; a new SOP Class that ensures the semantics of incremental return of the inventory are clear. This SOP Class definition is the primary subject of a second solicitation of Public Comment.*

5.B	**Should new attributes be added to existing C-FIND?**
	Study Update DateTime, Removed from Operational Use, and Reason for Removal Code Sequence may be useful for broad operational purposes, and could be added to the basic Patient Root and Study Root Queries.
	Study Update DateTime is proposed to be added to the Query/Retrieve Information Models at the Study level. The attributes Removed from Operational Use and Reason for Removal Code Sequence are proposed for the new Repository Query SOP Class; they may be added at a future time to other Query SOP Classes.

Issues Closed During Development

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.

400

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
1.2.840.10008.5.1.4.1.2.ffff	Repository Query	Query/Retrieve

405

DICOM PS 3.3: Information Object Definitions

Add ZIP, TAR, GZIP, HTTP, NFS, and IHE to Section 2 Normative References

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

410 [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*
415   https://tools.ietf.org/html/rfc1951

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

[RFC7230]   IETF. *Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing.*
        http://tools.ietf.org/html/rfc7230

[RFC7231]   IETF. *Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content*
420   http://tools.ietf.org/html/rfc7231

[RFC7530]   IETF. *Network File System (NFS) Version 4 Protocol* http://tools.ietf.org/html/rfc7530

…

2.6 OTHER REFERENCES

[IHE RAD TF-1] IHE International. *IHE Radiology (RAD) Technical Framework, Volume 1 – Integration Profiles*
425   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/

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

Add Inventory to Section 3 Definitions

3.8 DICOM INFORMATION OBJECT

435 Inventory

    A listing of DICOM Studies, Series, and SOP Instances managed by a repository system.

Scope of Inventory

> The parameters that select the DICOM Studies that are included in an Inventory. Parameters are
> specified as matching rules for Attribute values.

440

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

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 a repository
445 (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 includes
information on the stored SOP Instances, including access mechanisms supported by the repository.

> Notes
>
> 1. This information model is similar to the Study Root Query/Retrieve Information Model (see PS3.4 Section C.6.2.1).
>
> 450 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*).
>   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 if a Study has
>   multiple associated Imaging Service Requests, the request attributes may be encoded at the Series level.

455
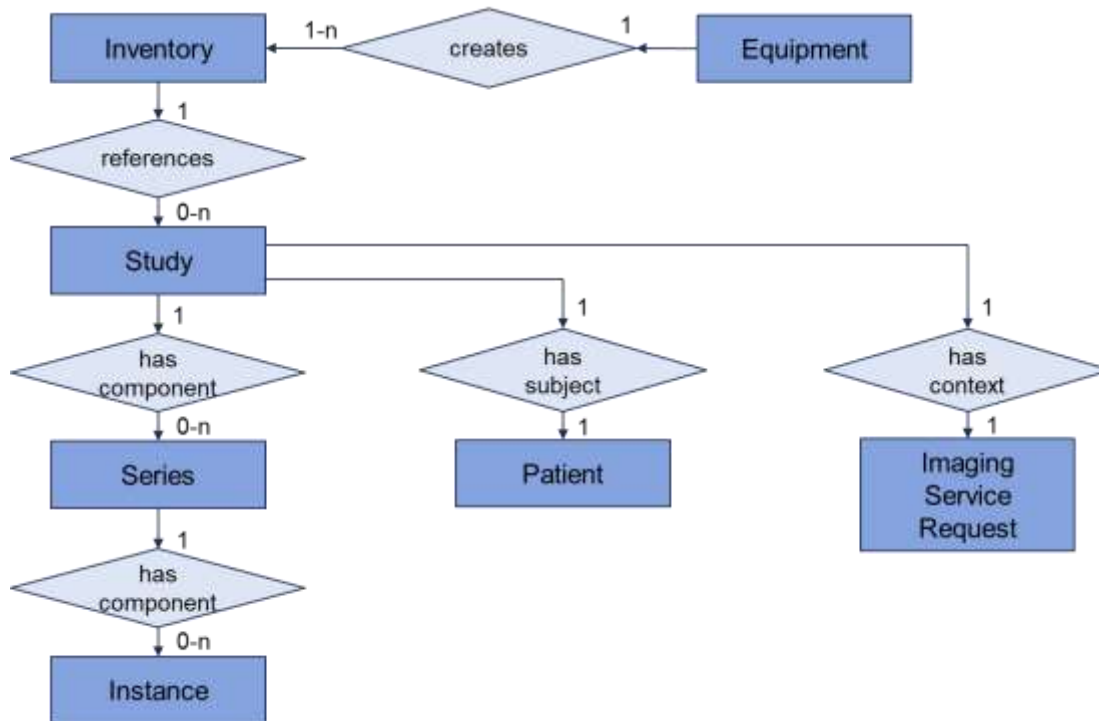
Figure 7.13.x-1. Inventory Information Model E-R Diagram

460 | *Add Inventory IOD to Section A.1.4 summary table*

– 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

…

465 **Table A.1-x. Composite Information Object Modules Overview - TBD**

Module \ IODs	Inventory
Inventory	M
Equipment	M
SOP Common	M

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

A.XX INVENTORY IOD

470 ### A.XX.1 Inventory IOD Description

The Inventory IOD provides an inventory of all Studies, Series, and SOP Instances managed by a repository, 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.

475 For implementation-specific reasons, a complete Inventory may need to be divided across multiple SOP 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:

480 See additional explanatory information in Annex XXXX in PS3.17, including discussion of the use of the Repository Query SOP Class to produce an Inventory SOP Instance.

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

485 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

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
490 Modules in the IOD.

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

495 **B.XX INVENTORY CREATION IOD**

B.XX.1 Inventory Creation IOD Description

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

500 **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

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

C.YY INVENTORY MODULES

505 ## C.YY.1 Inventory Module

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

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 (0008,0033), time point at which the inventory creation began. See Section C.YY.1.1.1
Content Time	(0008,0033)	1	With Content Date (0008,0023), 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*			Selection 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 transaction that created this Inventory SOP Instance. Required if Inventory created by a transaction of the Inventory Creation SOP Class. 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.

Name	Tag	Type	Description
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"*			
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)	2	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.7.30 in PS3.17

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

>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason the Study was removed from operational use. See Section C.YY.1.2.2. Required if Removed from Operational Use (00gg,0Fx6) value is Y. 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)	2	Count of stored Series in Study. See Section C.YY.1.2.3
>Number of Study Related Instances	(0020,1208)	2	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

Name	Tag	Type	Description
>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.
>>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 are permitted 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.12.1.1.9-1 "Original Attributes Macro Attributes"*			Recording of prior Attribute values and provenance of metadata changes at the Study level. See Section C.YY.1.2.5

Name	Tag	Type	Description
>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.
>File Set Access Sequence	(00gg,0FyE)	3	Description of non-DICOM protocol access to the set of stored SOP Instances comprising the Study. Only a single Item is permitted in this Sequence.
>>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
>>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

Name	Tag	Type	Description
>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 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 inventoried Series
>>Modality	(0008,0060)	1	Type of device, process or method that originally acquired or produced the data used to create the Instances in this Series.
>>Removed from Operational Use	(00gg,0Fx6)	1C	Flag that this inventoried Series has been removed from operational use. See Section C.YY.1.2.2 Enumerated Values: Y N Required if Series is not to be used for operational purposes, and Removed from Operational Use (00gg,0Fx6) Attribute at Study level does not have the value Y. May be present otherwise.
>>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason the Series was removed from operational use. See Section C.YY.1.2.2. Required if Removed from Operational Use (00gg,0Fx6) value is Y. 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.

Name	Tag	Type	Description
>>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 10-9 "Request Attributes Macro Attributes"*			
>>*Include Table C.12.1.1.9-1 "Original Attributes Macro Attributes"*			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
>>File Set Access Sequence	(00gg,0FyE)	3	Description of non-DICOM protocol access to the set of stored SOP Instances comprising this Series. Only a single Item is permitted in this Sequence.
>>>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.
>>Inventoried Instances Sequence	(0400,06x2)	2C	Description of 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 inventoried SOP Instance
>>>SOP Instance UID	(0008,0018)	1	SOP Instance UID of the inventoried SOP Instance

Name	Tag	Type	Description
>>>Removed from Operational Use	(00gg,0Fx6)	1C	Flag that this inventoried SOP Instance has been removed from operational use. See Section C.YY.1.2.2 Enumerated Values: Y N Required if SOP Instance is not to be used for operational purposes, and Removed from Operational Use (00gg,0Fx6) Attribute at Study or Series level does not have the value Y. May be present otherwise.
>>>Reason for Removal Code Sequence	(00gg,0Fx7)	2C	Reason the SOP Instance was removed from operational use. See Section C.YY.1.2.2. Required if Removed from Operational Use (00gg,0Fx6) value is Y. 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.12.1.1.9-1 "Original Attributes Macro Attributes"*			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
>>>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
>>>Storage Media File-Set ID	(0088,0130)	3	A human readable Identifier that identifies the Storage Media on which the Instance resides
>>>Alternate Representation Sequence	(0008,3001)	3	Reference to an alternate encoding of the inventoried SOP Instance. See Section C.YY.1.4.4. One or more Items are permitted in this Sequence.
>>>>Series Instance UID	(0020,000E)	1C	Unique identifier of the Series of the alternate representation. Required if the alternate representation is in a different Series than the inventoried SOP Instance.
>>>>SOP Class UID	(0008,0016)	1	SOP Class UID of the alternate representation SOP Instance

Name	Tag	Type	Description
>>>>SOP Instance UID	(0008,0018)	1	SOP Instance UID of the alternate representation SOP Instance
>>>>Purpose of Reference Code Sequence	(0040,A170)	1	The nature of the alternate encoding. One Item shall be included in this Sequence.
>>>>>Include *Table 8.8-1 "Code Sequence Macro Attributes"*			Baseline Context Group BCID 7205

510

C.YY.1.1 Inventory Module Attributes

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
515 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).

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.

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

C.YY.1.1.2 Scope of Inventory

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

> 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
530 > 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 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
535 Inventory as of the Content Date (0008,0023) and Content Time (0008,0033) for the repository system identified in the General Equipment Module (see Section C.7.5.1), including all SOP Instances referenced in the Incorporated Inventory Instance Sequence (0400,06x0).

Enumerated Values:

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

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

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

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

    Notes

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

550   2.   The reason for a value other than COMPLETE may be described in Inventory Instance Description (00gg,0Fx2).

    3.   See additional explanatory information and examples of use in Section XXXX.3.3.1 "Scope and Completion
Status" in PS3.17.

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
555 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
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.

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

The inclusion of an Inventory SOP Instance reference in the Incorporated Inventory Instance Sequence
(0400,06x0) 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
565 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 differ from the Content Date (0008,0023) and Content Time
(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
570   Instance, parallel processing across multiple separate subsystems, or reuse of historical Inventories. See Section
XXXX.3.3 in PS3.17.

    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.

575 ### C.YY.1.1.5 Inventoried Studies Sequence

The Inventoried Studies Sequence (0400,06x1) references Studies whose attributes match the Scope of
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.

580   Notes

    1.   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 for other
reasons. The application using an Inventory may need to reconcile such multiple occurrences (see Section
XXXX.7.30 "Study Record Reconciliation" in PS3.17).

585   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.

C.YY.1.2 Study Attributes

C.YY.1.2.1 Study Update DateTime

590 Study Update DateTime is the DateTime of the last update to Study contents (SOP Instances) or to Study metadata within the repository system. Metadata may change either due to coercion of Study related data (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 repository system.

595 As this is an Attribute of the Study as managed within the Repository system, the initial Study Update DateTime is the time that the Study was first received in the repository system, and is updated upon each new SOP Instance 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
600 (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 Repository system, and might not have a value for some or all of the Studies in the Inventory.

2. Study Update DateTime (0020,121x) is not the time of last user access to the Study contents.

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

605 The Removed from Operational Use (00gg,0Fx6) Attribute is defined at the Study, Series, and Instance levels. A value of Y indicates the inventoried Study, Series, or Instance has been removed from operational use related to patient care, although it may be retained in the repository system for other reasons (e.g., for audit of patient radiation exposure). At the Study and Series level, the Attribute indicates whether the entire Study or Series has been removed from operational use. A value of Y at the Study level supersedes any value specified for
610 subsidiary Series and Instances, and a value of Y at the Series level supersedes any value specified for subsidiary Instances.

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

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

2. Studies, Series, or Instances might be marked removed from operational use by actions associated with the processing of specific Key Object Selection Document SOP Instances, e.g., in accordance with [IHE RAD TF-1]
620 *Image Object Change Management Integration Profile* (IOCM). Those Key Object Selection Document SOP Instances, and their Series, may themselves be marked as removed from operational use. The Context Group for Reason for Removal Code Sequence (00gg,0Fx7) includes the Key Object Selection Concept Codes specified in IOCM.

3. The content of the Inventory might not match the information obtained through the Query/Retrieve Service since
625 Studies, Series, or Instances in Inventory SOP Instances marked as removed from operational use might not appear in Query/Retrieve responses (e.g., see [IHE RAD TF-2] *Section 4.66.4.1.3.1 Access to Rejected Instances)*.

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.

630 5. If the Application Entity retains records of deleted Studies, Series, or Instances, even though the actual Instances are physically deleted, it may include those records in the Inventory with an appropriate Reason for Removal Code Sequence (00gg,0Fx7) value. Such instances may have an Instance Availability (0008,0056) value "UNAVAILABLE" (see Section C.YY.1.2.9).

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

If the value of Inventory Level (00gg,0Fx3) is STUDY, an Inventory SOP Instance will only contain references at the Study level and the Inventoried Series Sequence (0400,06x2) will be absent. If the Scope of Inventory indicates Relational matching to the Series or Instance level (see Section C.YY.2.1.1.2), not all Series may be included in the Inventoried Series Sequence (0400,06x2).

In all cases, the Number of Study Related Series (0020,1206) and Number of Study Related Instances (0020,1208) shall give accurate counts of Series and SOP Instances stored in the Repository, respectively. The counts of Series and Instances shall include those that are marked as removed from operational use.

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 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 (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 Section XXXX.7.9.1 "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 (0400,0561) is the top level Data Set of the SOP Instances of the inventoried Study.

The Original Attributes Sequences (0400,0561) in the Inventory SOP Instance might not be identical to the Original Attributes Sequence (0400,0561) in the individual SOP Instances of the Study, and updates recorded in either location are valid (see Section C.YY.1.4.2).

Note

While defined in this specification at the Study, Series, and Instance levels, an Application Entity might manage the Original Attributes Sequence (0400,0561) at only the Instance level. In such case, changes to Study, Patient, Imaging Study Request, and Series Attributes would be recorded at the Instance level for all affected Instances.

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 (00gg,0Fx9) 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).

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 (00gg,0Fx9) Attribute is present with a value.

Notes

1. Stored Instance Base URI (00gg,0Fx9) is optional, and if not present the values of Folder URI (00gg,0FxA) and Stored Instance File URI (00gg,0FxB) at the Study level must be complete URIs. If Stored Instance Base URI (00gg,0Fx9) is present, those other Attributes may still provide complete URIs, rather than relative paths to be merged with the Base URI.

2. For example, the Stored Instance Base URI (00gg,0Fx9) value may be "nfs://pacs.exampleinstitution.org/JZ08555/", which when combined with a Stored Instance File URI (00gg,0FxB) value of "./2.25.916804767294.dcm" gives a URI "nfs://pacs.exampleinstitution.org/JZ08555/2.25.916804767294.dcm" to access a SOP Instance through the Network File System protocol. The initial "./" in the Stored Instance File URI (00gg,0FxB) indicates it is a relative path.

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 specified in Section P.1.2, Stored Instance File URI (00gg,0FxB) provides the URI for accessing that file.

Folder URI (00gg,0FxA) and/or Stored Instance File URI (00gg,0FxB) 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 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.

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 a non-DICOM file 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 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

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.

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

UNAVAILABLE the instance cannot be retrieved

Notes

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 Sequence (00gg,0Fx7) value (11yyy1, DCM, "Data Not Accessible from Storage") (see
Section C.YY.1.2.2).

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.

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 (00gg,0FxA) and/or Stored Instance File URI (00gg,0FxB) 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 URI (00gg,0Fx9) 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 (0018,0015) can be transcoded to their coded equivalent value, such
as by using 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
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

C.YY.1.4.1 File Access Sequence

The File Access Sequence (00gg,0FyF) provides one Item for each URI (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 (00gg,0Fx9) at the Study level.

> Note
>
> A repository system may store a SOP Instance on multiple 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

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.

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 (0400,0561) of the stored
SOP Instance.

Notes:

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

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 (0008,0061) and Number of Study Related Instances
   (0020,1208), and file access Attributes, such as Retrieve AE Title (0008,0054), Stored Instance Base URI
   (00gg,0Fx9), and Stored Instance File URI (00gg,0FxB). 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
   included in the Inventory SOP Instance, especially if they have been updated, such as change of an SR Concept
   Name Code Sequence (0040,A043) from a SNOMED RT to a SNOMED CT coded concept. See Section XXXX.3.5
   "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 coercion, correction or conversion of SOP Instance
   Attributes, the application may add an Item to the Original Attributes Sequence (0400,0561) describing the
   provenance of the change, but preserving any 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 (0400,0561) at the Study level, changes to Series or Instance attributes may
   have their provenance recorded in the Original Attributes Sequence (0400,0561) 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 (0400,0561) in the Inventory at the Study, Series, and Instance levels are added to the Original
   Attributes Sequence (0400,0561) of the stored SOP Instances. See Section XXXX.7.9.1 "Original Attributes
   Sequence" in PS3.17.

C.YY.1.4.3 Expiration DateTime

A repository system 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 time; the SOP Instance may become unavailable at
the access URI at an earlier or later time. This Expiration DateTime (00gg,0Fy9) 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.1.4.4 Alternate Representation Sequence

Alternate Representation Sequence (0008,3001) encodes a reference to an alternate encoding of the
inventoried SOP Instance. This alternate encoding may utilize a different SOP Class or Transfer Syntax, but the
alternate encoding of an image shall be the same image although it may have different image quality
characteristics.

> Notes

> 1. The Alternate Representation Sequence (0008,3001) allows the Inventory record of an original image to reference a lossy compressed version, and vice versa.

> 2. See also PS3.4 Section C.6.1.1.5.1.

> 2. An original SOP Instance that has been operationally replaced by an alternate encoded SOP Instance, and the original then deleted, may still have its Inventory record retained to provide forward referential resolution through the Alternate Representation Sequence (0008,3001) for other SOP Instances that reference the original. This record may have Instance Availability (0008,0056) set to UNAVAILABLE (see Section C.YY.1.2.9), and/or may also be marked as removed from operational use with Reason for Removal Code Sequence (00gg,0Fx7) value (11yyy3, DCM, "Replaced by Alternate Representation") (see Section C.YY.1.2.2).

C.YY.2 Inventory Related Macros

> Note

> 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 Key Attribute values that select the Studies included in the Inventory, i.e., the
Scope of Inventory.

The full semantics of the matching process is specified in Annex C 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 Attribute values for the Studies to be included In the Inventory. See Section C.YY.2.1.1.1. Zero or one Item shall be included in this Sequence. An empty Sequence, or a Sequence with a single empty Item, indicates universal match to all Studies in the Repository.

>Extended Matching Mechanisms	(00gg,0Fy1)	1C	Extended Attribute matching mechanisms applied to the Key Attributes, equivalent to mechanisms accepted during Extended Negotiation of the Query/Retrieve Service Class. See Section C.YY.2.1.1.1.
			
			Enumerated values:
			   RELATIONAL
			   COMBINED_DT
			   FUZZY_SEMANTIC
			   TIMEZONE_ADJUST
			   EMPTY_VALUE
			   MULTIPLE_VALUE
			
			Required if extended matching mechanisms were applied.

>Specific Character Set	(0008,0005)	1C	Expanded or replacement character sets used in the key Attributes. See Section C.12.1.1.2 Required if attributes within the Sequence Item use an expanded or replacement character set.
>Timezone Offset from UTC	(0008,0201)	1C	Offset from UTC to the local timezone for all DA and TM key Attributes. See Section C.12.1.1.8 Required if Extended Matching Mechanisms (00gg,0Fy1) includes value TIMEZONE_ADJUST

>Query/Retrieve Level	(0008.0052)	1C	Matching level for Relational matching mechanism. See Section C.YY.2.1.1.2.
			
			Enumerated Values:
			   STUDY
			   SERIES
			   IMAGE - SOP Instance level
			
			Required if Extended Matching Mechanisms (00gg,0Fy1) includes value RELATIONAL

>Range Matching Sequence	(00gg,0Fy2)	1C	The beginning and end values for range matching of date and time key attributes. See Section C.YY.2.1.1.3 Exactly two Items shall be included in this Sequence. Required if range matching is applied to any Key Attribute
>>*Any range matched DA, DT, or TM Key Attribute*		2	Key matching value for Attribute.
>List of UID Matching Sequence	(00gg,0Fy4)	1C	Key Attribute values for list of UID matching. See Section C.YY.2.1.1.4 Required if List of UID matching is applied to any key attribute One or more Items shall be included in this Sequence.

>>*Any list of UID matched UI Key Attribute*		1	Key matching value for Attribute.
>Empty Value Matching Sequence	(00gg,0Fy5)	1C	Key Attributes for empty value matching. See Section C.YY.2.1.1.5 Required if empty value matching is applied to any Key Attribute Only a single Item is permitted in this Sequence.
>>*Any empty value matched Key Attribute*		2	Zero-length value for Key Attribute.
>General Matching Sequence	(00gg,0Fy6)	1C	Key Attribute values for Single Value, Multiple Value, and Wildcard matching. See Section C.YY.2.1.1.6 Required if Single Value, Multiple Value, or Wildcard matching is applied to any key attribute Only a single Item is permitted in this Sequence.
>>*Any Single Value, Multiple Value, or Wildcard matched Attribute*		1	Key matching value for Attribute.

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 Key Attribute values used to select Studies to be included in an Inventory.

Selection of Studies is done by a matching process between Key Attribute values specified in the Scope of Inventory Sequence (00gg,0Fx0) and the corresponding values in Attributes of Studies in the Repository. The matching process is identical to the Attribute matching specified for the Query/Retrieve Service Class in Section C.2.2.2 "Attribute matching" of PS3.4.

Any extended matching mechanisms beyond the baseline Query/Retrieve behavior used in selection of Studies for the Inventory are identified in the Extended Matching Mechanisms (00gg,0Fy1) Attribute.

> Notes
>
> 1. Attributes of Studies that may be matched include those additional Query/Retrieve Attributes that may be managed by the Repository system for managing stored SOP Instances, but are not specified in Composite IODs. See Sections C.3.4 and C.6.x4.1 in PS3.4. Of particular interest for some inventory use cases are Study Update DateTime (0020,121x), Modalities in Study (0008,0061), and Removed from Operational Use (00gg,0Fx6).
>
> 2. In Query/Retrieve, extended mechanisms are agreed through SOP Class Extended Negotiation. In the Inventory SOP Instances, the use of such mechanisms is simply encoded in the Extended Matching Mechanisms (00gg,0Fy1) Attribute by the Instance creator.

Key Attributes for matching in the Query/Retrieve Service Class are allowed to have values that are exceptions to the VR for those attributes in other contexts (see Section 6.1. in PS3.5), or to have multiplicities that are exceptions to the VM of those attributes as specified in PS3.6. The Scope of Inventory Sequence (00gg,0Fx0) specifies Sequence Attributes that provide equivalent encoding of Key Attributes while maintaining conformance to VR and VM specifications.

Universal matching Key Attributes, i.e., those with zero-length value or with the value "*", do not select Studies, and they shall not be included in the Scope of Inventory. Zero-length values are permitted in Key Attributes of the Range Matching Sequence (00gg,0Fy2), where they represent an unbounded beginning or end of a range, and in Key Attributes of the Empty Value Matching Sequence (00gg,0Fy5).

C.YY.2.1.1.2 Relational Matching

850 If Extended Matching Mechanisms (00gg,0Fy1) does not include the value RELATIONAL, only Key Attributes defined at the Study level, as specified in Section C.6.2.1.2 "Study Level" in PS3.4, shall be included in Scope of Inventory.

If Extended Matching Mechanisms (00gg,0Fy1) includes the value RELATIONAL, Studies are selected for the Inventory in accordance with the specifications for Relational-queries matching in Section C.4.1.3.2.2 in PS3.4
855 using the Query/Retrieve Level (0008.0052) and Key Attributes specified. The Inventory Level (00gg,0Fx3) may differ from this Query/Retrieve Level and determines the level of information that is included in the Inventory, In the case that the Inventory Level is at or lower than the Query/Retrieve Level, then only the records for entities matching all Key Attributes will be included in the Inventory.

Note

860 For example, specification of Relational Matching, Query/Retrieve Level (0008.0052) SERIES, and a Key Attribute Modality (0008,0060) with value CT will include in the Inventory all Studies that have a Series with Modality CT. If the Inventory Level (00gg,0Fx3) is INSTANCE, only those Series and Instances with modality CT will be included in the Inventory; Series and Instances with Modality PR or SR, for instance, in the same Study will not be included.

However, baseline (non-relational) matching of Modalities in Study (0008,0061) with value CT will also include in the
865 Inventory all Studies that have a Series with Modality CT. But since that Key Attribute performs selection at the Study level, an Inventory will include all Series of any modality for that Study.

C.YY.2.1.1.3 Range Matching Sequence

Range Matching Sequence (00gg,0Fy2) provides values for Key Attributes that use range matching of dates and times using the mechanism specified in Section C.2.2.2.5 in PS3.4.

870 The same set of Key Attributes shall be present in both Items of this Sequence. For each Key Attribute, at least one Item shall have a non-zero length value.

Values in the first Item represent the beginning of the range, and values in the second Item represent the end of the range. Either the first or the second value may be empty, in which case the range has an unbounded beginning or end, respectively. The range is inclusive; matched Attributes whose value is at or after the first
875 value and before or after the second value match the range.

Notes

In the equivalent Query/Retrieve Service Class range matching, the beginning and end values are encoded in a single Data Element, and delimited by "-".

If Extended Matching Mechanisms (00gg,0Fy1) includes the value COMBINED_DT, related date and time
880 attributes, such as Study Date (0008,0020) and Study Time (0008,0030), are considered together as a datetime value for matching.

If Extended Matching Mechanisms (00gg,0Fy1) includes the value TIMEZONE_ADJUST, dates and time are adjusted for the local timezone specified in Timezone Offset from UTC (0008,0201), as specified in Section C.2.2.2.1.3 in PS3.4.

885 ### C.YY.2.1.1.4 UID List Matching Sequence

List of UID Matching Sequence (00gg,0Fy4) provides a list of UIDs, one per Item, to be matched using the list of UIDs matching mechanism specified in Section C.2.2.2.2 in PS3.4.

Note

In the equivalent Query/Retrieve Service Class List of UIDs matching, the VM=1 restriction is overridden for Key
890 Attributes.

C.YY.2.1.1.5 Empty Value Matching Sequence

Empty Value Matching Sequence (00gg,0FyD) specifies Key Attributes that match when the corresponding Attribute in the Study has no value, as specified in Section C.2.2.2.x7 in PS3.4. Key Attributes identified in this Sequence shall have zero-length values.

895 Empty Value Matching Sequence (00gg,0FyD) shall only be present if Extended Matching Mechanisms (00gg,0Fy1) includes the value EMPTY_VALUE.

> Note
>
> In the equivalent Query/Retrieve Service Class empty value matching, the value consisting of exactly two QUOTATION MARK characters signifies an empty value, and its use is permitted in Key Attributes with a VR of DA, DT, or TM.

900 ### C.YY.2.1.1.6 General Matching Sequence

General Matching Sequence (00gg,0Fy6) specifies Key Attributes that are matched using Single Value and Wildcard matching, as specified in Section C.2.2.2.1 and Section C.2.2.2.4 in PS3.4.

If Extended Matching Mechanisms (00gg,0Fy1) includes the value MULTIPLE_VALUE, General Matching Sequence (00gg,0Fy6) may also specify Key Attributes that are matched using Multiple Value matching, as 905 specified in Section C.2.2.2.x8 in PS3.4.

> Note
>
> If Extended Matching Mechanisms (00gg,0Fy1) does not include the value MULTIPLE_VALUE, Key Attributes with multiple values have undefined meaning.

If Extended Matching Mechanisms (00gg,0Fy1) includes the value TIMEZONE_ADJUST, dates and time are 910 adjusted for the local timezone specified in Timezone Offset from UTC (0008,0201), as specified in Section C.2.2.2.1.3 in PS3.4.

If Extended Matching Mechanisms (00gg,0Fy1) includes the value FUZZY_SEMANTIC, matching of Attributes with VR PN may be done using implementation-specific fuzzy semantic matching, as specified in Section C.2.2.2.1.1 in PS3.4.

915 ### 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 (00gg,0FxB) 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 (00gg,0FxB) 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
File Length in Container	(00gg,0FxF)	3	Length of the SOP Instance file within a container 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)	1C	The algorithm used for generating a Message Authentication Code. See Table C.12.1.1.3.1.2-1 for Defined Terms. Required if MAC (0400,0404) is present.
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

920 **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.

Stored Instance File URI may be a relative path reference beginning with the single-dot-segment "./" (see
925 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, File Offset in Container, and File Length in Container

If the stored SOP Instance is included in a container file as specified in Section P.1.2, Filename in Container (00gg,0FxD) is the filename within that container file of the file containing the SOP Instance. File Offset in
930 Container (00gg,0FxE) is an optional byte offset for the start of the file for the SOP Instance, and File Length in Container (00gg,0FxF) is the file length. 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
> 935 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
940 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 differs 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.

945

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

950   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). See also Section XXXX.7.3 "Using Referenced Inventories" in PS3.17.

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 (see Section XX in PS3.4)
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

955 ## 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*		Key Attribute values 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 (00gg,0FyB)
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 through Inventory MOVE or Inventory GET SOP Classes (see Section XX in PS3.4)
Retrieve URL	(0008,1190)	Base URI of the origin server to retrieve this Inventory SOP Instance through the Non-Patient Instance Service (see Section 12 in PS3.18)
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 time; the SOP Instances may become unavailable at an earlier or later time. This Expiration DateTime (00gg,0Fy9) applies to all available access methods (both DICOM and non-DICOM protocols).

Note

See Section XXXX.7.5 "Inventory Lifecycle Management" in PS3.17 for further discussion.

Add new Inventory Directory Record Type to Section F.3.2.2

Table F.3-3. Directory Information Module Attributes

980

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).

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

Table F.4-1. Relationship Between Directory Records

985

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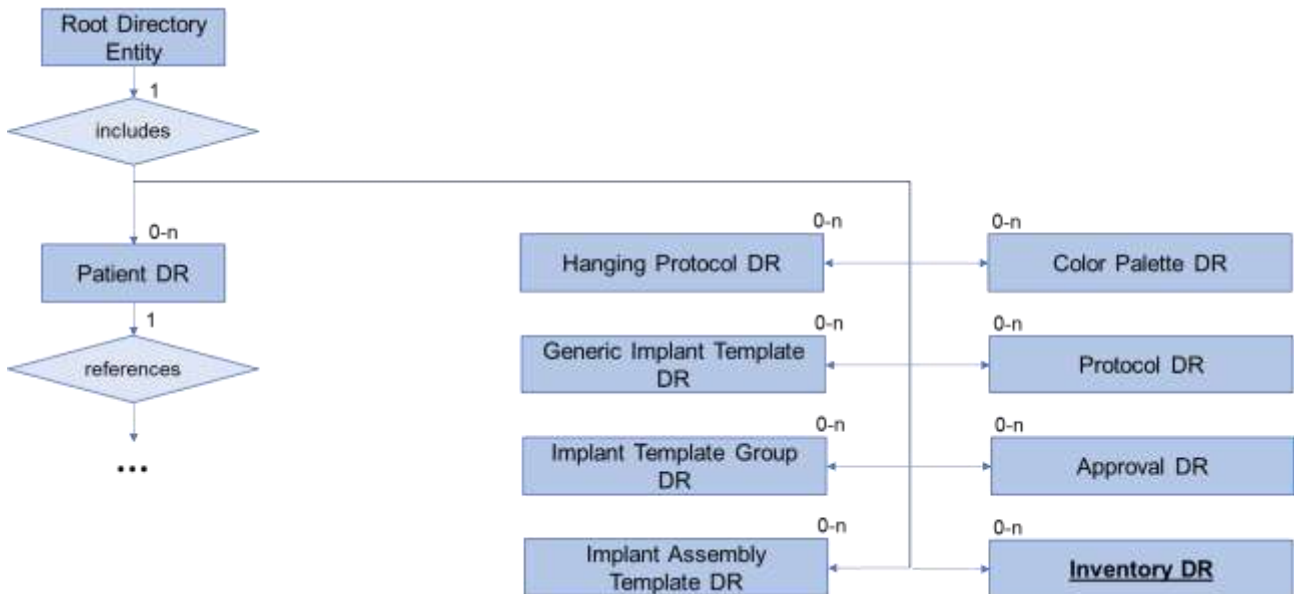
Figure F.4-1. Basic Directory IOD Information Model

990 *Add new Inventory Directory Record Definition to Section F.5*

F.5.x Inventory Directory Record Definition

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
995 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.

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	
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	

1000

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

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
1005 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 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).

1010 Note

See also Section XXXX.6 "Security Considerations" in PS3.17.

P.1 FILES AND SETS OF FILES

P.1.1 DICOM File Format

Each stored SOP Instance that is a target of an Inventory URI link shall be encoded in accordance with the
1015 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.

P.1.2 Container File Formats

1020 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 container file, and should be mitigated against when reading a container file. See Section XXXX.6.2 in PS3.17.

1025 **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] DEFLATE compressed files, and does not permit encryption.

1030 **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

"ustar" is the formal name, derived from "Unix Standard Tape ARchive", although TAR is the common term and is used in
1035 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].

Note

1040   The GZIP format supports compression of the contained file using [RFC1951] DEFLATE compression.

P.1.2.4 TARGZIP

Multiple SOP Instance files may be included in a TAR container file in accordance with [POSIX], which is then contained within a GZIP container file in accordance with [RFC1952].

Note

1045   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

Depending on the storage mechanism, these structures may alternately be denoted "directories". The term "folder" is
1050   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.

P.2 ACCESS PROTOCOLS

1055 ## 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 boundary, and if a slash ("/") character occurs at that boundary, it shall be placed in the base URI data element
1060   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 Conformance Statement for the implementation shall specify the protocol(s) in its description of conformance to
1065   the Inventory SOP Class. Some 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/

2.   Several protocols are layered on HTTPS. While the specific protocol should be specified in the Conformance
1070   Statement, the only feature required is the ability to read an object, which is generally implemented simply as an HTTP GET in such protocols, and may be implemented with only a URI.

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

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	[RFC7230], [RFC7231] Includes various cloud storage implementations

1075

DICOM PS3.4: Service Class Specifications

Add empty value matching and multiple value matching to Query/Retrieve

(See also change to PS3.5 to allow empty value matching characters for VRs DA, TM, and DT)

C.2.2.2 Attribute Matching

1080 The following types of matching may be performed on Key Attributes in the Query/Retrieve Service Class:

- Single Value Matching
- List of UID Matching
- Universal Matching
- Wild Card Matching
1085 - Range Matching
- Sequence Matching
- **Empty Value Matching**
- **Multiple Value Matching**

Matching requires special characters (i.e., "*","?","-", "=", ~~and~~ "\"**, and QUOTATION MARK**), which need not be
1090 part of the character repertoire for the VR of the Key Attributes.

…

C.2.2.2.1 Single Value Matching

If the value specified for a Key Attribute in a request is non-zero length and if it is not of VR SQ and:

- of VR of AE, CS, LO, LT, PN, SH, ST, UC, UR or UT and contains **a single value with** no wild card
1095 characters, **and if extended negotiation of Empty Value Matching is successful it does not have the value of exactly two QUOTATION MARK characters,** or
- of VR of DA, TM or DT and contains a single value with no "-" **and no QUOTATION MARK characters**, or
- of any other VR

1100 then Single Value Matching shall be performed.

…

C.2.2.2.x7 Empty Value Matching

If extended negotiation of Empty Value Matching is successful, and for attributes with VR of AE, CS, DA, DT, LO, LT, PN, SH, ST, TM, UC, UR, or UT, a Key Attribute value of exactly two QUOTATION MARK
1105 **characters (22H in the Default Character Set) specified in a request indicates matching to an empty value. Such a match is successful to entities with a zero-length value for the corresponding Attribute, or for which the Attribute is not present.**

Notes

1. **Other quotation mark characters defined in an expanded or replacement character set, such as double**
1110 **angle quotation marks, left and right curved quotation marks, etc., are not valid for specifying empty value matching.**

2. **VRs AE, LO, LT, PN, SH, ST, UC, and UT allow the QUOTATION MARK character. Matching of a string of two QUOTATION MARK characters in those attributes can be done with wild card matching using the ASTERISK character (e.g., ""*)**

1115 **C.2.2.2.x8 Multiple Value Matching**

If extended negotiation of Multiple Value Matching is successful, and for attributes with VR of AE, CS, LO, LT, PN, SH, ST, UC, UR, or UT and a defined VM greater than 1, a Key Attribute with multiple values specified in a request indicates matching to a multi-valued entity attribute. Such a match is successful only if each and every value in the Key Attribute matches a value in the corresponding attribute in the
1120 **entity.**

Key Attribute Values shall not include wild cards.

The order of Attribute Values in the request need not match the order in the corresponding entity Attribute.

> **Notes**

1125 1. **This is a logical 'AND' function for all values in the Key Attribute. E.g., a value of CT\PT for Modalities in Study (0008,0061) in the Request Identifier will match only Studies that include both a CT Series and a PT Series.**

2. **There may be additional values in the entity Attribute that do not match values in the Key Attribute.**

3. **Multiple Value Matching does not apply to Specific Character Set (0008,0005), which is not a Key Attribute**
1130 **for matching, but rather specifies the encoding of characters in the Identifier.**

4. **In the absence of successful Extended Negotiation for Multiple Value Matching, the semantics of multiple values in a Key Attribute is undefined.**

C.2.2.3 Matching Multiple Values

1135 When **using Single Value Matching, Wild Card Matching, or Range Matching for** matching an Attribute that has a value multiplicity of greater than one, if any of the values match, then all values shall be returned.

When using Multiple Value Matching, if matching is successful, then all values from the entity Attribute shall be returned.

1140 *Add empty value matching and multiple value matching to Query/Retrieve Extended Negotiation*

C.5.1.1 SOP Class Extended Negotiation

The SOP Class Extended Negotiation allows, at Association establishment, peer DICOM AEs to exchange application Association information defined by specific SOP Classes. This is achieved by defining the Service-class-application-information field. The Service-class-application-information field is used to define support for
1145 relational-queries, combined date time matching, fuzzy semantic matching of person names, timezone query adjustment, ~~and~~ Enhanced Multi-Frame Image Conversion**, Empty Value Matching, and Multiple Value Matching**.

This negotiation is optional. If absent, the default conditions shall be:

- no relational-query support
1150 - separate (independent) Range Matching of date and time Attributes
- literal matching of person names with case sensitivity unspecified
- timezone query adjustment unspecified
- no Enhanced Multi-Frame Image Conversion support
- **no empty value matching**
1155 - **no multiple value matching**

The Association-requester, for each SOP Class, may use one SOP Class Extended Negotiation Sub-Item. The SOP Class is identified by the corresponding Abstract Syntax Name (as defined by PS3.7) followed by the Service-class-application-information field. This field defines one or more sub-fields:

- relational-query support by the Association-requester
1160
- combined date and time Range Matching by the Association-requester
- literal or fuzzy semantic matching of person names by the Association-requester
- timezone query adjustment by the Association-requester
- Enhanced Multi-Frame Image Conversion support by the Association-requester
- **empty value matching support by the Association-requester**
1165
- **multiple value matching support by the Association-requester**

…

C.5.1.1.1 SOP Class Extended Negotiation Sub-Item Structure (A-ASSOCIATE-RQ)

The SOP Class Extended Negotiation Sub-Item consists of a sequence of mandatory fields as defined by PS3.7. Table C.5-1 defines the Service-class-application-information field for DICOM Query/Retrieve SOP
1170 Classes and Specialized DICOM Query/Retrieve SOP Classes that include the C-FIND operation. This field may be either one or more bytes in length (i.e., item bytes 2~~, 3, 4, and 5~~ **through 7** are optional).

Table C.5-1. SOP Class Extended Negotiation Sub-Item
(Service-Class-Application-Information Field) - A-ASSOCIATE-RQ

Item Bytes	Field Name	Description of Field
1	Relational-queries	This byte field defines relational-query support by the Association-requester. It shall be encoded as an unsigned binary integer and shall use one of the following values 0 - relational queries not supported 1 - relational queries supported
…		
6	**Empty Value Matching**	**This byte field defines whether or not Empty Value Matching may be requested by the Association-requester. It shall be encoded as an unsigned binary integer and shall use one of the following values** **0 - empty value matching not requested** **1 - empty value matching requested**
7	**Multiple Value Matching**	**This byte field defines whether or not Empty Value Matching and Multiple Value Matching may be requested by the Association-requester. It shall be encoded as an unsigned binary integer and shall use one of the following values** **0 - multiple value matching not requested** **1 - multiple value matching requested**

1175 ### C.5.1.1.2 SOP Class Extended Negotiation Sub-Item Structure (A-ASSOCIATE-AC)

The SOP Class Extended Negotiation Sub-Item is made of a sequence of mandatory fields as defined by PS3.7. Table C.5-2 defines the Service-class-application-information field for DICOM Query/Retrieve SOP Classes and Specialized DICOM Query/Retrieve SOP Classes that include the C-FIND operation. This field may be either one or more bytes in length (i.e., item bytes 2~~, 3, 4, and 5~~ **through 7** are optional).

1180 **Table C.5-2. SOP Class Extended Negotiation Sub-Item (Service-Class-Application-Information Field) - A-ASSOCIATE-AC**

Item Bytes	Field Name	Description of Field
1	Relational-queries	This byte field defines relational-query support for the Association-acceptor. It shall be encoded as an unsigned binary integer and shall use one of the following values 0 - relational queries not supported 1 - relational queries supported
…		
6	**Empty Value Matching**	**This byte field defines Empty Value Matching support by the Association-acceptor. It shall be encoded as an unsigned binary integer and shall use one of the following values** **0 - empty value matching not supported** **1 - empty value matching supported**
7	**Multiple Value Matching**	**This byte field defines Multiple Value Matching support by the Association-acceptor. It shall be encoded as an unsigned binary integer and shall use one of the following values** **0 - multiple matching not supported** **1 - multiple matching supported**

Add empty and multiple value matching to Query/Retrieve Conformance

C.6.1.2.1.1 C-FIND SCU Conformance

1185 …

An implementation that conforms to one of the SOP Classes of the Patient Root SOP Class Group as an SCU shall state in its Conformance Statement whether or not it supports extended negotiation of combined date-time matching**,** ~~and/or~~ fuzzy semantic matching of person names**, empty value matching, and/or multiple value matching**.

1190 …

C.6.1.2.2.1 C-FIND SCP Conformance

…

An implementation that conforms to one of the SOP Classes of the Patient Root SOP Class Group as an SCP shall state in its Conformance Statement whether or not it supports extended negotiation of combined date-time
1195 matching**,** ~~and/or~~ fuzzy semantic matching of person names**, empty value matching, and/or multiple value matching**. If fuzzy semantic matching of person names is supported, then the mechanism for fuzzy semantic matching shall be specified.

…

C.6.2.2.1.1 C-FIND SCU Conformance

1200 …

An implementation that conforms to one of the SOP Classes of the Study Root SOP Class Group as an SCU shall state in its Conformance Statement whether or not it supports extended negotiation of combined date-time matching**,** ~~and/or~~ fuzzy semantic matching of person names**, empty value matching, and/or multiple value matching**.

1205 …

C.6.2.2.2.1 C-FIND SCP Conformance

…

An implementation that conforms to one of the SOP Classes of the Study Root SOP Class Group as an SCP shall state in its Conformance Statement whether or not it supports extended negotiation of combined date-time 1210 matching, ~~and/or~~ fuzzy semantic matching of person names**, empty value matching, and/or multiple value matching**. If fuzzy semantic matching of person names is supported, then the mechanism for fuzzy semantic matching shall be specified.

…

1215 | *Add Study Update DateTime to common Query/Retrieve Information Model*

C.3.4 Additional Query/Retrieve Attributes

Some optional Attributes that may be used in Query/Retrieve Information Models that are not Attributes of an Information Object Definition and, therefore, are not defined in PS3.3. These Attributes are defined in Table C.3-1.

1220 **Table C.3-1. Additional Query/Retrieve Attributes**

Attribute Name	Tag	Attribute Description
…		
Study Update DateTime	**(0020,121x)**	**DateTime of last update to Study metadata (e.g., Patient, Study or Procedure, Imaging Service Request, or Series Attributes) or to the set of SOP Instances in the Study** **Note** **May specify DateTime of update to the storage location of Study SOP Instances if the SCP supports Attributes for stored SOP Instance access through a non-DICOM protocol (see Section C.6.x4.1.3).**

…

C.6.1.1.3 Study Level

Table C.6-2 defines the keys at the Study Information level of the Patient Root Query/Retrieve Information
1225 Model.

Table C.6-2. Study Level Keys for the Patient Root Query/Retrieve Information Model

Attribute Name	Tag	Type
…		
Number of Study Related Instances	(0020,1208)	O
Study Update DateTime	**(0020,121x)**	**O**
All other Attributes at Study Level		O

…

1230 ### C.6.2.1.2 Study Level

Table C.6-5 defines the keys at the Study Information level of the Study Root Query/Retrieve Information Model.

Table C.6-5. Study Level Keys for the Study Root Query/Retrieve Information Model

Attribute Name	Tag	Type
…		
Number of Study Related Instances	(0020,1208)	O
Study Update DateTime	**(0020,121x)**	**O**
All other Attributes at Study Level		O

Add Repository Query SOP Class to Annex C.6

C.6.x4 Repository Query SOP Class

The Repository Query SOP Class uses the C-FIND Service and the Study Root Query/Retrieve Information Model. The SOP Class specifies additional semantics and behaviors for the SCU and SCP beyond those defined for the Study Root Query/Retrieve Information Model – FIND SOP Class.

1240 In particular, the Repository Query SOP Class supports incremental query/response for a large number of entities using the following features:

- Both the SCU and SCP may set limits on the number of entity records (Pending status responses) returned in a C-FIND transaction.

- If the number of responses would exceed the SCU or SCP limit, the C-FIND query processing is
1245 terminated with a defined Warning status response indicating an incomplete set of responses.

- If requested by the SCU, each response shall return a unique Record Key by which the SCP has ordered the C-FIND responses.

- The SCU may include the last Record Key it receives in one C-FIND transaction as the Prior Record Key in the C-FIND Request for a subsequent transaction. The SCP processes that subsequent C-FIND
1250 beginning with the record following that Prior Record Key.

The Repository Query SOP Class also supports return of URI links for access to stored SOP Instances using a non-DICOM file access protocol. Since stored SOP Instances accessed through a non-DICOM protocol might not include all current metadata Attributes (such as updated patient names or IDs), the Repository Query SOP Class also supports return of current metadata Attributes.

1255 Note:

See additional explanatory information in Annex XXXX in PS3.17, including discussion of the use of the Repository Query SOP Class to produce an Inventory SOP Instance.

C.6.x4.1 Additional Query Information Model Attributes

The Repository Query SOP Class uses the Study Root Query/Information model specified in Section C.6.2.1,
1260 but specifies additional Key Attributes. These Key Attributes, like those defined in Section C.3.4, are not specified in the Composite IODs of PS3.3, but represent information that may be used by a repository system for managing stored SOP Instances. Table C.6.x4.1-1 defines the additional Key Attributes.

Table C.6.x4.1-1 Additional Keys for Repository Query

Attribute Name	Tag	Type	Attribute Description
Record Key	(00gg,0Fz0)	R	Implementation-specific unique identifier of the entity record. This Attribute is supported only with Universal Matching. See Section C.6.x4.1.1
Removed from Operational Use	(00gg,0Fx6)	O	Specifies whether entity at the Query/Retrieve Level specified in the Identifier of the C-FIND and its set of Composite Object Instances have been removed from operational use related to patient care. See Section C.6.x4.1.2
Reason for Removal Code Sequence	(00gg,0Fx7)	O	Reason the entity at the Query/Retrieve Level specified in the Identifier of the C-FIND was removed from operational use.

Zero or one Item may be included in this Sequence. |

>*Include Table C.6-2b "Basic Code Value Keys Macro with Optional Keys"*			Defined Context Group is CID 7xxx
File Set Access Sequence	(00gg,0FyE)	O	Description of non-DICOM protocol access to sets of stored SOP Instances. This Attribute is defined only at the Study and Series levels. This Attribute is supported only with Universal Matching, i.e., if present in the Request Identifier it is zero-length or has only a zero-length Item. See Section C.6.x4.1.3 Zero or one Item may be included in this Sequence in a Response Identifier.
>Stored Instance Base URI	(00gg,0Fx9)	-	Base URI for accessing SOP Instances through a non-DICOM protocol.
>Folder URI	(00gg,0FxA)	-	Access URI for a folder containing all SOP Instances for this Study or Series
>Stored Instance File URI	(00gg,0FxB)	-	Access URI for a container file containing all the SOP Instances for this Study or Series

>Container File Type	(00gg,0FxC)	-	Type of container file.
			
			Defined Terms:
			    ZIP
			    TAR
			    GZIP
			    TARGZIP
			
			Required in Response Identifier if Stored Instance File URI (00gg,0FxB) is present.

File Access Sequence	(00gg,0FyF)	O	Description of non-DICOM protocol access to a stored SOP Instance. This Attribute is defined only at the Composite Instance level. This Attribute is supported only with Universal Matching, i.e., if present in the Request Identifier it is zero-length or has only a zero-length Item. See Section C.6.x4.1.3 Zero or more Items may be included in this Sequence in a Response Identifier.
>Stored Instance File URI	(00gg,0FxB)	-	Access URI for file containing the SOP Instance. Required in Response Identifier.

>Container File Type	(00gg,0FxC)	-	Type of container file.
			
			Defined Terms:
			    ZIP
			    TAR
			    GZIP
			    TARGZIP
			
			Required in Response Identifier if Stored Instance File URI (00gg,0FxB) references a container file.

>Filename in Container	(00gg,0FxD)	-	Filename within a container file of the file containing the SOP Instance. Required in Response Identifier if Stored Instance File URI (00gg,0FxB) references a container file.

>File Offset in Container	(00gg,0FxE)	-	Byte offset (zero-based) within a container file for the start of the SOP Instance file
>File Length in Container	(00gg,0FxF)		Length of the SOP Instance file within a container file
>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	-	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)	-	The algorithm used for generating a Message Authentication Code. Required in Response Identifier if MAC (0400,0404) is present.
>MAC	(0400,0404)	-	Message Authentication Code computed across stored instance file for verification of file integrity.
Updated Metadata Sequence	(00gg,0Fz2)	C	Attributes at the Query/Retrieve Level and higher for which values are different from the values contained in stored SOP Instance files. This Attribute shall be supported by the SCP if File Access Sequence (00gg,0FyF) is supported. This Attribute is supported only with Universal Matching, i.e., if present in the Request Identifier it is zero-length or has only a zero-length Item. See Section C.6.x4.1.4 Zero or one Item may be included in this Sequence in a Response Identifier.
>Any Attributes at Query/Retrieve Level or higher		-	

C.6.x4.1.1 Record Key

Record Key (00gg,0Fz0) is defined at the Study, Series, and Instance query levels. It is an implementation-specific unique identifier of the entity record in the SCP. The SCP of the Repository Query SOP Class shall return non-zero length values of Record Key (00gg,0Fz0). The content of Record Key (00gg,0Fz0) is opaque to applications other than the SCP.

Record Key (00gg,0Fz0) values shall have the property of ordering. The SCP shall construct the Attribute Value such that each value is ordered with respect to all other such values, and such that C-FIND Response Identifiers are returned in the ordering of Record Key (00gg,0Fz0) values. The SCP shall be able to determine from a given value the next entity record to be returned that matches the given Query Request Identifier, without repeating any records.

Record Key (00gg,0Fz0) values shall have the property of persistence sufficient that a value can be used as the Prior Record Key (00gg,0Fz1) value in a subsequent Query Request (see Section C.6.x4.5.1). The SCP may establish implementation specific conditions after which a Record Key (00gg,0Fz0) value is not valid to allow continuation of a sequence of Query operations. The SCP shall be able to determine from a given Prior Record Key (00gg,0Fz1) value whether that value is still valid for determining the next record to be returned.

1280   Notes

1. UIDs may not have the properties of ordering for record return and/or determinability of validity for use by an SCP as the Record Key.

2. The structure and content of Record Key (00gg,0Fz0) values is totally SCP implementation-specific, and opaque to the SCU. Values may be permanent, or may be constructed dynamically during query processing.

1285   3. The intention of the ordering and persistence requirements for the Record Key is to allow an SCU to obtain the complete inventory matching the Key Attributes in a sequence of Queries. See Section XXXX.2.2 "Record Key and Continuation" in PS3.17.

C.6.x4.1.2 Removed from Operational Use

The Removed from Operational Use (00gg,0Fx6) Attribute is defined at the Study, Series, and Instance query
1290   levels. The Enumerated Values are:

Y   the instances are not to be used for operational purposes

N   the instances may be used for operational purposes

A value of Y indicates the Study, Series, or Instance has been removed from operational use related to patient care, although it may be retained in the repository system for other reasons (e.g., for audit of patient radiation
1295   exposure). At the Study and Series level, the Attribute indicates whether the entire Study or Series has been removed from operational use. A value of Y at the Study level supersedes any value specified for subsidiary Series and Instances, and a value of Y at the Series level supersedes any value specified for subsidiary Instances.

While defined at the Study, Series, and Instance levels, an SCP might not support this Attribute at some, or any,
1300   of those levels. E.g., an SCP may only manage this Attribute at the Instance level, and is not required to infer a value for the Series or Study level.

Notes

1. The meaning of "operational use related to patient care" is implementation or site specific, but generally includes diagnostic, clinical, and therapeutic uses, as well as administrative uses necessary for providing care (e.g.,
1305   insurance authorization).

2. Studies, Series, or Instances might be marked removed from operational use by actions associated with the processing of specific Key Object Selection Document SOP Instances, e.g., in accordance with [IHE RAD TF-1] *Image Object Change Management Integration Profile* (IOCM). Those Key Object Selection Document SOP Instances, and their Series, may themselves be marked as removed from operational use. The Context Group for
1310   Reason for Removal Code Sequence (00gg,0Fx7) includes the Key Object Selection Concept Codes specified in IOCM.

3. An SCP might exclude entities marked as removed from operational use from the C-FIND Responses of other Query/Retrieve SOP Classes (e.g., see [IHE RAD TF-2] *Section 4.66.4.1.3.1 Access to Rejected Instances).*

4. Removed from Operational Use (00gg,0Fx6) is independent of Instance Availability (0008,0056). A composite
1315   instance may have been removed from operational use but is still accessible at the rapidity specified by Instance Availability (0008,0056). Conversely, an instance may not have been removed from operational use but is UNAVAILABLE for retrieval.

5. If the SCP retains records of deleted Studies, Series, or Instances, even though the actual Instances are physically deleted, it may include those entities in the C-FIND Response with an appropriate Reason for Removal Code
1320   Sequence (00gg,0Fx7) value. Such instances may have an Instance Availability (0008,0056) value "UNAVAILABLE" (see Section C.4.1.1.3.2).

C.6.x4.1.3 File Set Access Sequence and File Access Sequence

The SCP may support optional Attributes providing a URI link to SOP Instances stored in the DICOM File Format (see Section 7 "DICOM File Format" in PS3.10) and accessible through a non-DICOM file access
1325 protocol (see Annex P "Stored File Access Through Non-DICOM Protocols" in PS3.3).

> Note
>
> See Section XXXX.3.4 "Access Mechanisms for Repository Data" in PS3.17

For a query at the Study or Series level, return of Stored Instance Base URI (00gg,0Fx9) within the File Set Access Sequence (00gg,0FyE) establishes an [RFC3986] base URI that shall be combined with relative path
1330 reference URIs for non-DICOM protocol access to SOP Instances of the Study or Series. If all of the stored SOP Instance files of that Study or Series entity are catalogued in a single folder, Folder URI (00gg,0FxA) provides the URI for protocol operations on that folder. If all of the stored SOP Instance files are in a single container format file, Stored Instance File URI (00gg,0FxB) provides the URI for accessing that file. Folder URI (00gg,0FxA) and/or Stored Instance File URI (00gg,0FxB) may be a relative path reference URI beginning with
1335 the single-dot-segment "./", and the URI is relative to the name space of the Stored Instance Base URI (00gg,0Fx9).

> Note
>
> Stored Instance Base URI (00gg,0Fx9) is optional, and if not present the values of Folder URI (00gg,0FxA) and Stored Instance File URI (00gg,0FxB) must be complete URIs. If Stored Instance Base URI (00gg,0Fx9) is present, those other
1340 Attributes may still provide complete URIs, rather than relative paths to be merged with the Base URI.

For a query at the Composite Instance level, Items of the File Access Sequence (00gg,0FyF) in the Response Identifier each provide an [RFC3986] URI to access the stored SOP Instance. Stored Instance File URI (00gg,0FxB) may be a URI relative path reference beginning with the single-dot-segment "./", and the URI is relative to the name space of the Stored Instance Base URI (00gg,0Fx9) specified within the File Set Access
1345 Sequence (00gg,0FyE) at the Series level, if present, or otherwise to the Stored Instance Base URI (00gg,0Fx9) specified within the File Set Access Sequence (00gg,0FyE) at the Study level.

> Notes
>
> 1. The SCP may store a SOP Instance on multiple 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
1350 non-DICOM protocol at multiple URIs.
>
> 2. If the Stored Instance File URI (00gg,0FxB) for a SOP Instance is a relative path reference URI, the SCU will need to have obtained the Stored Instance Base URI (00gg,0Fx9) from the Study and Series level queries.
>
> 3. See Section XXXX.7.2 "Using Non-DICOM Protocols" in PS3.17.

C.6.x4.1.4 Updated Metadata Attributes

1355 Although a stored SOP Instance referenced by Stored Instance File URI (00gg,0FxB) must be conformant to its IOD (per the requirements of the DICOM File Format), some Attributes might not have current values (e.g., Patient Name may have been corrected or changed after the Instance was stored).

Updated Metadata Sequence (00gg,0Fz2) contains all Attributes at the Query level and at higher Query levels whose values are different from the values contained in the stored SOP Instance file. It may also contain
1360 Attributes that do not have differing values. It shall not include Attributes specified in Section C.3.4 or in Table C.6.x4.1-1.

Note

If the SCP does not track which specific Attributes have changed values, it may simply return all current values that it knows.

1365 As Updated Metadata Sequence (00gg,0Fz2) provides current metadata for SOP Instances accessed via the Stored Instance File URI (00gg,0FxB), the SCP shall support this attribute if it also supports File Access Sequence (00gg,0FyF).

At any Query level Updated Metadata Sequence (00gg,0Fz2) may include the Original Attributes Sequence (0400,0561) describing the provenance of changes to Attributes at that level or at higher Query levels.

1370 If present in a Request Identifier, the absence of Updated Metadata Sequence (00gg,0Fz2) in the Response Identifier indicates the SCP does not support the Attribute (see Section C.2.2.1.3). A zero-length value or a single empty Item in a Response Identifier indicates the SCP supports the Attribute, but there are no differing Attribute values.

C.6.x4.2 Repository Query Request Identifier Additional Attributes

1375 In addition to the attributes specified in Section C.4.1.1.3.1, a Repository Query Request Identifier may conditionally contain the Attributes shown in Table C.6.x4.2-1 to request certain C-FIND response content.

Table C.6.x4.2-1. Repository Query Request Identifier Additional Attributes

Attribute Name	Tag	Type	Attribute Description
Prior Record Key	(00gg,0Fz1)	C	Unique ordered identifier of an entity returned in Record Key (00gg,0Fz0) in a prior request. Entities returned for the current request will begin with the first matching record with a record key subsequent to this record key. See Section C.6.x4.1.1 Required if Repository Query response is to continue from an incomplete prior Repository Query response
Maximum Number of Records	(00gg,0Fz2)	C	Maximum number of matched entities to be returned for this request. Required if SCU has a resource limitation for handling responses.

C.6.x4.3 Additional Status Codes

1380 In addition to the statuses specified in Section C.4.1.1.4 and general status code values and fields related to status code values defined for C-FIND DIMSE Service in PS3.7, the Repository Query SOP Class specifies the additional status code values of Table C.6.x4.3-1.

Table C.6.x4.3-1. Repository Query Additional Status Values

Service Status	Further Meaning	Status Codes	Related Fields
Failure	Invalid Prior Record Key	A710	None
Warning	Matching reached response limit, subsequent request may return additional matches	B001	None

1385 **C.6.x4.4 C-FIND SCU Behavior**

In addition to the behaviors specified in Section C.4.1.2, the SCU shall interpret a response with a status equal to Warning to convey the end of Pending responses, and an indication that the SCP is not further processing the request in this C-FIND operation.

C.6.x4.5 C-FIND SCP Behavior

1390 The Repository Query SOP Class adds the following behaviors to both Baseline and Extended Behavior of the SCP as described in Section C.4.1.3.

C.6.x4.5.1 Record Key

The SCP of the Repository Query SOP Class shall support return of a non-zero length value of Record Key (00gg,0Fz0).

1395 The SCP shall order C-FIND Responses by value of Record Key (00gg,0Fz0) if that Attribute is included in the Request Identifier.

Note

If the Record Key (00gg,0Fz0) is not requested, there is no requirement on ordering, and the SCU will not be able to issue a subsequent query request with a Prior Record Key.

1400 **C.6.x4.5.2 Incomplete Response and Warning Status**

The SCP may limit the number of Pending responses (i.e., returned matching records) in a Query transaction to an implementation-defined resource limit. Further, the SCP shall send no more than the number specified by the SCU in the Request Identifier Attribute Maximum Number of Records (00gg,0Fz2). If either limit is reached, and there are additional records that satisfy the Key Attribute matching, the SCP shall send a final Status value B001
1405 "Warning".

Note

The SCP resource limit may be a number of records, or a time limit on database operations, or another limit that allows timely completion of a Repository Query transaction.

C.6.x4.5.3 Subsequent Response

1410 If Prior Record Key (00gg,0Fz1) is present in the Repository Query Request Identifier, the SCP shall return entity records, continuing with the next record key, that satisfy the Key Attribute matching. The SCP shall return a Failure status code A710 if the value of Prior Record Key (00gg,0Fz1) is invalid for identifying an entity determining a point of continuation.

For a sequential set of Query transactions with the same Key Attribute Values, and for which each C-FIND
1415 Request after the first includes a Prior Record Key (00gg,0Fz1) value equal to the last Record Key (00gg,0Fz0) in the C-FIND Responses to the prior Query, the SCP shall return the same set of C-FIND Responses as if the entire Query would have been completed in a single transaction. Whether new entities and updates received by the SCP after the time of initial Query might be included in the C-FIND Responses is implementation specific.

Notes

1420　1. For example, an entity added after the time of initial Query might be assigned a Record Key that precedes a Record Key already sent in a C-FIND Response, or the content of a Study may be updated after a Query is in process. The implementation determines whether and how such changes are reported in C-FIND Responses. However, any entity should only appear once in such a sequential set of Query transactions.

2. The SCP should not assume that the last Record Key it sent will be the used as the Prior Record Key in a
1425　　subsequent transaction. Error conditions may cause the SCU to need to resume from an earlier Record Key.

3. There is no requirement for the Key Attributes in a Request Identifier to be identical to those in a prior Repository Query transaction, even though the presence of Prior Record Key (00gg,0Fz1) implies a continuation of a prior transaction. Each Repository Query transaction is independent. The SCP does not need to retain any state information between Repository Query transactions, although it may include state
1430 information in the Record Key (00gg,0Fz0).

C.6.x4.5.4 Entities Removed from Operational Use

If the SCP supports the Key Attribute Removed from Operational Use (00gg,0Fx6) (see Section C.6.x4.1.2), query responses shall include all entities regardless of the value of that Attribute.

Note

1435 As the Study Root Query/Retrieve Information Model – FIND SOP Class is typically used for operational purposes, an SCP implementation of that SOP Class by policy might not return entities with Removed from Operational Use (00gg,0Fx6) value "Y".

C.6.x4.6 Conformance Requirements

An implementation may conform to the Repository Query SOP Class as an SCU, SCP or both. The
1440 Conformance Statement shall be in the format defined in PS3.2.

C.6.x4.6.1 C-FIND SCU Conformance

An implementation that conforms to the Repository Query SOP Class shall support queries against the Study Root Query/Retrieve Information Model described in Section C.6.2.1 and Section C.6.x4.1 using the C-FIND SCU behavior described in Section C.4.1.2 and Section C.6.x4.4.

1445 An implementation that conforms to the Repository Query SOP Class as an SCU shall be capable of generating queries using Hierarchical Search. It shall not generate queries using Relational-queries unless the Relational-queries option has been successfully negotiated. An implementation that conforms to Repository Query SOP Class as an SCU shall state in its Conformance Statement whether it may generate Relational-queries.

An implementation that conforms to the Repository Query SOP Class as an SCU shall state in its Conformance
1450 Statement:

- whether or not it supports extended negotiation of combined date-time matching, fuzzy semantic matching of person names, empty value matching, and/or multiple value matching.

- how it makes use of Specific Character Set (0008,0005) and Timezone Offset From UTC (0008,0201) when encoding queries and interpreting responses.

1455 - any limitations it places on the number of C-FIND responses through the Maximum Number of Records (00gg,0Fz2) Attribute.

C.6.x4.6.2 C-FIND SCP Conformance

An implementation that conforms to the Repository Query SOP Class shall support queries against the Study Root Query/Retrieve Information Model described in Section C.6.2.1 and Section C.6.x4.1 using the C-FIND
1460 SCP behavior described in Section C.4.1.3 and Section C.6.x4.5.

An implementation that conforms to the Repository Query SOP Class as an SCP shall state in its Conformance Statement:

- which Optional Keys it supports for existence and which for matching.

- whether it supports Relational Search. If it supports Relational Search, then it shall also state its support
1465 for extended negotiation of relational-queries as an SCP.

- whether or not it supports extended negotiation of combined date-time matching, fuzzy semantic matching of person names, empty value matching, and/or multiple value matching. If fuzzy semantic

matching of person names is supported, then the mechanism for fuzzy semantic matching shall be specified.

1470
- whether it supports case-insensitive matching for PN VR Attributes, and list Attributes for which this applies.

- how it makes use of Specific Character Set (0008,0005) and Timezone Offset From UTC (0008,0201) when interpreting queries, performing matching and encoding responses.

- any limitations it places on the number of C-FIND responses.

1475
- the specific conditions after which a Record Key (00gg,0Fz0) value is not valid to allow resumption of a C-FIND operation.

- any non-DICOM mechanisms used to access SOP Instances that may be specified by a Stored Instance File URI (00gg,0FxB), including protocols (such as NFS, SMB, or HTTP), and use of folders and/or container files (ZIP, TAR, TARGZIP) for the SOP Instances. If such mechanisms are used, the
1480   use of Updated Metadata Sequence (00gg,0Fz2) shall be described.

C.6.x4.7 SOP Classes

The UID for the Repository Query SOP Class is listed in Table C.6.x4.7-1.

Table C.6.x4.7-1. SOP Classes for Study Root Query/Retrieve

SOP Class Name	SOP Class UID
Repository Query	1.2.840.10008.5.1.4.xxuid.1

1485

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

---

GG.1.1 Scope

The Non-Patient Object Storage Service Class defines an application-level class-of-service that allows one DICOM AE to send a SOP Instance of a non-patient-related information object to another DICOM AE.

1490 **Information Object Definitions for Instances that are transferred under the Non-Patient Object Storage Service Class shall adhere to a Composite Instance IOD Information Model specified in PS3.3, but which does not have at its root the Patient Information Entity representing an individual real world Patient.**

> **Note**
>
> 1495 **"Non-patient-related" does not imply that there is no patient-related protected health information (PHI) in the Instances. E.g, the Inventory IOD does include Attributes of PHI, but it does not have a Patient IE at the root of its information model.**

…

GG.3 SOP CLASSES

The application-level services addressed by the Non-Patient Object Storage Service Class definition are
1500 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**

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

---

1505 ## 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:

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

1510 - 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.

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

1515 | *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
1520 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
1525 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
1530 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
1535 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)

1540 **XX.2.1 E-R Model**

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

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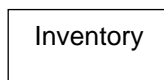
Figure XX.2-1. Inventory Q/R Information Model E-R Diagram

1545 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.

XX.2.2 Inventory Q/R Information Model Attributes

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

1550 **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 an Inventory Creation SOP Class transaction, may be present otherwise.
Content Date	(0008,0023)	R	1	
Content Time	(0008,0033)	O	1	
Scope of Inventory Sequence	(00gg,0Fx0)	O	2	May only be requested with Universal matching
Inventory Purpose	(00gg,0Fx1)	O	1	
Inventory Instance Description	(00gg,0Fx2)	O	2	
Inventory Level	(00gg,0Fx3)	O	1	
Inventory Completion Status	(0400,06x5)	O	1	
Number of Study Records	(0400,06x6)	O	1	

Note

1. See also the additional returned attributes in the Response Identifier specified in Section XX.4.2.2.1.

1555 2. Further definition of these attributes is specified in Section C.YY.1 in PS3.3

XX.3 DIMSE-C SERVICE GROUPS

XX.3.1 C-FIND Operation

See the C-FIND Operation definition for the Basic Worklist Management Service Class (Section K.4.1), and
1560 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.

XX.3.2 C-MOVE Operation

See the C-MOVE Operation definition for the Query/Retrieve Service Class (Section C.4.2). No Extended
1565 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.

Note

1570 More than one Inventory SOP Instance may be retrieved, using List of UID matching.

XX.3.3 C-GET Operation

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.

1575 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.

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

1580

XX.4 SOP CLASS DEFINITIONS

XX.4.1 SOP Classes

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:

1585 **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

XX.4.2 Conformance Requirements

An implementation may conform to one of the Inventory Query/Retrieve Service SOP Classes as an SCU, SCP
1590 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

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
1595 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 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).

1600 ##### 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 Service Class (see Section C.4.3.2.1 and Section XX.3.3).

XX.4.2.2 SCP Conformance

1605 ##### 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 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
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)
File Access Sequence	(00gg,0FyF)	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).
>Stored Instance File URI	(00gg,0FxB)	1	
>Container File Type	(00gg,0FxC)	1C	Required if Stored Instance File URI (00gg,0FxB) is present and references a container format file.
>Filename in Container	(00gg,0FxD)	1C	Required if Stored Instance File URI (00gg,0FxB) is present and references a container format file.
>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	1	
>MAC Algorithm	(0400,0015)	3	
>MAC	(0400,0404)	3	

Note

Further definition of these attributes is specified in Section C.YY.3 in PS3.3. In that Section, the Attributes apply to
stored patient-related Composite SOP Instances, e.g., images, while here they apply to Inventory SOP Instances.

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
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
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
1635 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.
1640 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

1645 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
1650 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

1655

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.

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.

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.6 "Security Considerations" in PS3.17.

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

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-requestor) 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 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 Instance specified in Table ZZ.1-3 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 SOP Instance is conceptually created during start-up of each Storage Management Service Class SCP Application.

Table ZZ.1-3. Storage Management SOP Instance UID

SOP Instance UID	UID Name
1.2.840.10008.1.20.1.x	Storage Management SOP Instance (well-known)

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.

1700 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 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.

1705 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
1710

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 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).

1715 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

1720 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).
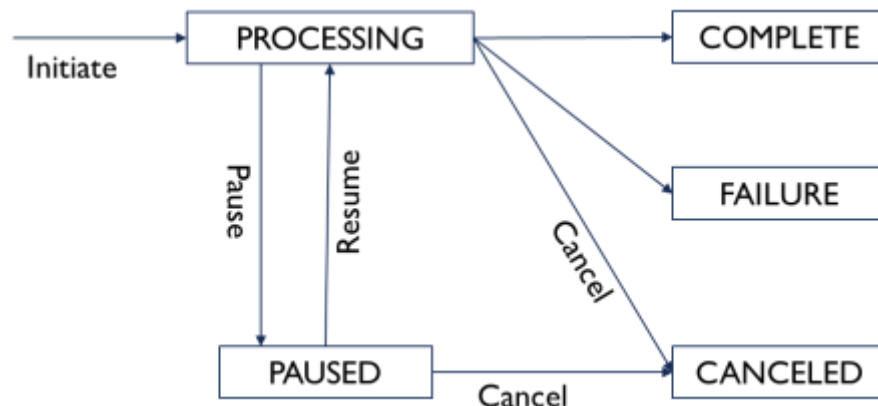
Figure ZZ.2-1 Inventory Production Statuses and Status Transitions

1725 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.

1730 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
		>Extended Matching Mechanisms	(00gg,0Fy1)	1C/1C
		>Specific Character Set	(0008,0005)	1C/1C
		>Timezone Offset from UTC	(0008,0201)	1C/1C
		>Range Matching Sequence	(00gg,0Fy2)	1C/1C
		>>*Any range matched DA, DT, or TM Key Attribute*		2/2
		>List of UID Matching Sequence	(00gg,0Fy4)	1C/1C
		>>*Any list of UID matched UI Key Attribute*		1/1
		>Empty Value Matching Sequence	(00gg,0Fy5)	1C/1C
		>>*Any empty value matched Key Attribute*		2/2
		>General Matching Sequence	(00gg,0Fy6)	1C/1C
		>>*Any Single Value, Multiple Value, or Wildcard matched Attribute*		1/1
		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. Studies that match all specified keys shall be included in the
produced Inventory.

The Scope of Inventory Sequence may be zero-length to indicate universal match to all Studies in the repository (i.e., a complete Inventory).

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

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 (0008,1195) will be included by the SCP in N-EVENT-REPORT primitives (see Section ZZ.2.3.1). Use of the Transaction UID (0008,1195) Attribute allows the SCU to match requests and results that may 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 Key Attributes for Studies to be included in the Inventory. The SCU shall specify in Extended Matching Mechanisms (00gg,0Fy1) any requested matching mechanisms or matching semantics beyond the baseline specified for Query in Section C.2.2.2. The SCU may request the SCP to regularly report on the progress of producing the Inventory associated with the Transaction UID (0008,1195) 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 (0008,1195). The SCU may request the SCP to change the interval for progress reports as specified by Requested Status Interval (00gg,0Fy7).

> 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 (0008,1195).

The SCU shall use the N-ACTION primitive with Action Type "Pause" (Action Type ID = 14) to request the SCP to pause production of the Inventory associated with the Transaction UID (0008,1195). 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-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.

For Action Type "Initiate" (Action Type ID = 11), the SCP shall return an error status 0212H (Mistyped argument) if the SCP does not support all of the character sets identified in Specific Character Set (0008,0005) or all of the matching mechanisms identified in Extended Matching Mechanisms (00gg,0Fy1). The SCP shall return a warning status B010H if one or more of the Key Attributes specified in the Scope of Inventory Sequence (00gg,0Fx0) are not supported for matching (see Table ZZ.2.2.3-1).

1790

Table ZZ.2.2.3-1. N-ACTION Response Status Values

Service Status	Further Meaning	Status Codes	Related Fields
Failure	Processing Failure	0110	(0000,0902)
	Refused: Not authorized	0124	(0000,0902)
	Mistyped argument - one of the parameters supplied has not been agreed for use	0212	(0000,0902)
	Resource Limitation	0213	(0000,0902)
	Other failure status as specified in Section 10.1.4.1.10 in PS3.7		
Warning	One or more of Key Attributes are not supported for matching	B010	(0000,1005)
Success	Success	0000	

Note

See Section E.1 "Registry of DICOM Command Elements" in PS3.7 for definition of the Related Fields.

If the SCP conveys a success status for Action Type "Initiate" (Action Type ID = 11), it shall initiate production of
1795 an Inventory and periodic status reporting in accordance with the parameters of the Scope of Inventory
Sequence (00gg,0Fx0) and Inventory Level (00gg,0Fx3) provided in the N-ACTION request. Any value provided
by the SCU for Inventory Purpose (00gg,0Fx1) shall be 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
1800 Inventory SOP Instance of Study records already obtained, and set Inventory Completion Status (0400,06x5) to
"FAILURE".

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

1805 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).

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 (0008,1195). If the Cancel request includes Retain Instances
(00gg,0Fy8) value "N", the SCP may delete any Inventory SOP Instances associated with the Transaction UID
1810 (0008,1195). 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 "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.

1815 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 (0008,1195). The SCP may complete any processing
necessary to reach 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 (0008,1195).

1820 Note

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

ZZ.2.3 Notifications

ZZ.2.3.1 Event Information

1825 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
		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

1830

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)
1835 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

1840 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
1845 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
1850 "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,
1855 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 DateTime 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
1860 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
1865 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
1870 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

1875 • 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)

1880 • 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).

1885 ## 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.

1890 ### 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
1895   periodic reporting of status if that attribute is not provided by the SCU

- character repertoires supported in Specific Character Set (0008,0005) in the Scope of Inventory Sequence (00gg,0Fx0)

- support for Extended Matching Mechanisms (00gg,0Fy1) in the Scope of Inventory Sequence (00gg,0Fx0)

- Key Attributes supported for matching in the Scope of Inventory Sequence (00gg,0Fx0)

1900   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

1905 • 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
1910   Instances Notification, including the specific non-DICOM protocols (if any) referenced in the Stored Instance File URI

DICOM PS 3.5: Data Structures and Encoding

Add empty value matching characters for VRs DA, TM, and DT to Section 6.2

1915

Table 6.2-1. DICOM Value Representations

VR Name	Definition	Character Repertoire	Length of Value
…			
DA Date	A string of characters of the format YYYYMMDD; where YYYY shall contain year, MM shall contain the month, and DD shall contain the day, interpreted as a date of the Gregorian calendar system. … **In the context of a Query with empty value matching (see PS3.4), specifications for the component format do not apply**	"0"-"9" of Default Character Repertoire In the context of a Query with range matching (see PS3.4), the character "-" is allowed, and a trailing SPACE character is allowed for padding. **In the context of a Query with empty value matching (see PS3.4), the QUOTATION MARK character is allowed.**	8 bytes fixed In the context of a Query with range matching (see PS3.4), the length is 18 bytes maximum. **In the context of a Query with empty value matching (see PS3.4), the length is 2 bytes fixed.**
…			
DT Date Time	A concatenated date-time character string in the format: YYYYMMDDHHMMSS.FFFFFF&ZZXX … **In the context of a Query with empty value matching (see PS3.4), specifications for the component format do not apply.**	"0"-"9", "+", "-", "." and the SPACE character of Default Character Repertoire **In the context of a Query with empty value matching (see PS3.4), the QUOTATION MARK character is allowed.**	26 bytes maximum In the context of a Query with range matching (see PS3.4), the length is 54 bytes maximum.
…			
TM Time	A string of characters of the format HHMMSS.FFFFFF; where HH contains hours (range "00" - "23"), MM contains minutes (range "00" - "59"), SS contains seconds (range "00" - "60"), and FFFFFF contains a fractional part of a second as small as 1 millionth of a second (range "000000" - "999999"). A 24-hour clock is used. Midnight shall be represented by only "0000" since "2400" would violate the hour range. The string may be padded with trailing spaces. Leading and embedded spaces are not allowed. … **In the context of a Query with empty value matching (see PS3.4), specifications for the component format do not apply.**	"0"-"9", "." and the SPACE character of Default Character Repertoire In the context of a Query with range matching (see PS3.4), the character "-" is allowed. **In the context of a Query with empty value matching (see PS3.4), the QUOTATION MARK character is allowed.**	14 bytes maximum In the context of a Query with range matching (see PS3.4), the length is 28 bytes maximum.

DICOM PS 3.6: Data Dictionary

1920 *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,0FxF)	File Length in Container	FileLengthInContainer	UL	1	
(00gg,0Fy0)	Stored Instance Transfer Syntax UID	StoredInstanceTransferSyntaxUID	UI	1	
(00gg,0Fy1)	Extended Matching Mechanisms	ExtendedMatchingMechanisms	CS	1-n	
(00gg,0Fy2)	Range Matching Sequence	RangeMatchingSequence	SQ	1	
(00gg,0Fy4)	List of UID Matching Sequence	UIDListMatchingSequence	SQ	1	
(00gg,0Fy5)	Empty Value Matching Sequence	EmptyValueMatchingSequence	SQ	1	
(00gg,0Fy6)	General Matching Sequence	GeneralMatchingSequence	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,0FyE)	File Set Access Sequence	FileSetAccessSequence	SQ	1	
(00gg,0FyF)	File Access Sequence	FileAccessSequence	SQ	1	
(00gg,0Fz0)	Record Key	RecordKey	OB	1	
(00gg,0Fz1)	Prior Record Key	PriorRecordKey	OB	1	
(00gg,0Fz2)	Updated Metadata Sequence	UpdatedMetadataSequence	SQ	1	
(0020,121x)	Study Update DateTime	StudyUpdateDateTime	DT	1	

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

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

1925

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
1.2.840.10008.5.1.4.1.2.ffff	Repository Query	SOP Class	PS3.4
1.2.840.10008.1.20.1.x	Storage Management SOP Instance	Well-known Instance	PS3.4

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	

1930

DICOM PS 3.7: Message Exchange

Add Warning status to C-FIND

1935

9.1.2.1.6 Status

Indicates the status of the response. It may have any of the following values (see also Annex C):

a. Success (0000H) - This indicates that processing of the matches is complete. It shall not contain a matching Identifier.

1940 …

f. Failed (Status value is Service Class specific) - Indicates that the C-FIND operation failed at the performing DIMSE Service User.

g. Warning (Status value is Service Class specific) - Indicates that the performing DIMSE Service User has terminated the C-FIND operation, and the set of returned matching Identifiers may not be
1945 **complete.**

…

9.1.2.2 C-FIND Service Procedures

…

1950 The following C-FIND service procedures apply to the performing DIMSE Service User:

…

f. When the C-FIND operation completes (either in success or in failure), the performing DIMSE Service User issues a C-FIND response with the status set to either Refused, Failed, **Warning,** or Success to the DIMSE Service Provider.

1955 …

<mark>DICOM PS 3.16: Content Mapping Resource</mark>

Add to Section 2 Normative References

1960 ## 2.1 GENERAL

[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

---

Add Context Group to Annex B DCMR Context Groups

1965

CID 7XXX REASON FOR REMOVAL FROM OPERATIONAL USE

Type: Extensible

Version: 2021mmdd

UID: 1.2.840.10008.6.1.nnn

1970 **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 Document
DCM	11yyy3	Replaced by Alternate Representation

Add terminology concepts to Annex D DCMR Context Groups

1975 **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 Document	A Key Object Selection Document 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*)	
11yyy3	Replaced by Alternate Representation	SOP Instance has been replaced by a SOP Instance with an alternate encoding. The alternate encoding may utilize a different SOP Class or have different image quality characteristics (e.g., lossy compression).	

DICOM PS 3.17: Explanatory Information

Add to Section 2 Normative References

2.2 OTHER REFERENCES

[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*
http://tools.ietf.org/html/rfc7233 .

Add explanatory Annex

Annex XXXX Inventories (Informative)

XXXX.1 THE DICOM DATA MANAGEMENT ENVIRONMENT

DICOM data in a healthcare organization is typically 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 routinely retained "forever", and data set sizes are increasing with 3D/4D and multimodality studies. Repositories in many institutions store over a billion instances across 10's of millions of studies, with data volumes over one petabyte. Enterprise-scale management tools and data are required, including interoperability features that operate at large scales.

XXXX.1.1 Inventories

An important feature supporting repository management is the ability to obtain an inventory of the repository contents in a standard format. DICOM provides two complementary methods - an interactive query-based mechanism and a persistent inventory information object. Both approaches address the issues associated with a large inventory with over a billion records.

The two methods, query-based and persistent object, each satisfy distinct approaches to implementation and use of inventories. Generally, repository systems that already implement query for patient-oriented operations may find implementation of a query-based inventory to be expeditious, but there may be repository systems that may want to implement production of an inventory object. Many user applications need to have a persistent object that can be processed offline in a bulk operation, such as E/T/L (extract, transfer, and load) to a data warehouse, but some inventory using applications may desire to use an interactive query model. There may also be applications that can mediate between queries and persistent objects (see Section XXXX.7.1)

> Note
>
> The following sections describing the Repository Query and the Inventory Information Object, respectively, are written to be read independently of one another; there is therefore significant overlap between the sections.

XXXX.2 REPOSITORY QUERY

XXXX.2.1 Overview

The Repository Query SOP Class is an extension of the Study Root Query/Retrieve – FIND SOP Class with features that support very large response sets. For queries that might return millions of records, it allows both the SCU and SCP to set constraints on the number of records to be returned in a single C-FIND transaction. It specifies well-defined behavior for a "partially completed" status to be returned if not all entities selected by the key matching attributes in the request were returned, and allows the SCU to specify a "continuation point" in a subsequent query to return responses from that point onward. This provides a mechanism to "window" through the entries in a deterministic way without overburdening either the SCU or the SCP.

XXXX.2.2 Record Key and Continuation

Deterministic behavior is achieved by the SCP imposing a sorting order on the returned records that is based on a unique value for each entry, the Record Key. In the query, the SCU can request return of the Record Key in each response. When a "partially completed" status is returned, or if there is a communications failure during the transaction, the value in the last received response can be used in the next query to request the SCP to continue returning responses for matching entities with next higher unique value as per the sorting order.

The structure and content of Record Key values is totally SCP implementation-specific, and opaque to the SCU. Values may be permanent, or may be constructed dynamically during query processing. For SCPs that use a relational database, the database primary record key might be used as the unique Record Key value, although an implementation might choose to use some other element or some type of session-oriented key. The intention is that the SCP manage Record Keys such that the SCU will be able to use them to obtain a complete inventory in a sequence of Queries in a reasonable time period, recognizing that for large inventories that time period may be substantial. If there are limitations on the lifetime of the Record Key, they should be documented in the SCP Conformance Statement.

XXXX.2.3 Key Matching Attributes

The Repository Query SOP Class uses Key Attribute matching exactly as defined for Study Root C-FIND. However, several additional Key Attributes are defined in support of repository management.

XXXX.2.3.1 Objects Removed from Operational Use

A repository system may manage Studies, Series, and Instances that are marked in the database as removed from operational use. The associated SOP Instances may have been physically deleted, or they may be left in storage, commonly denoted as 'soft deleted', 'deprecated' or 'hidden' (see Section C.6.x4.1.2 "Removed from Operational Use" in PS3.4). The SCU may desire to receive records of these entities in the inventory, especially to determine which entities were removed since the last inventory (see example in Section XXXX.7.10.1).

The Repository Query SOP Class defines Removed from Operational Use (00gg,0Fx6) and Reason for Removal Code Sequence (00gg,0Fx7) as Key Attributes. Studies, Series, or Instances might be marked removed from operational use by local user actions, or by actions associated with the processing of specific Key Object Selection Document SOP Instances, e.g., in accordance with [IHE RAD TF-1] *Image Object Change Management Integration Profile (IOCM).*

XXXX.2.3.2 Access to stored objects

A repository system may support multiple access mechanisms for each stored instance – DIMSE C-MOVE 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). The DIMSE Retrieve AE Title (0008,0054) is returned in the query response without specific request of the SCU (see Section C.4.1.1.3.2 in PS3.4). The

Repository Query SOP Class defines File Set Access Sequence (00gg,0FyE) at the Study and Series levels and File Access Sequence (00gg,0FyF) at the Instance level as Key Attributes. The SCU can request these Attributes to obtain a non-DICOM protocol URI link to the stored instance (see Section C.6.x4.1.3 "File Set Access Sequence and File Access Sequence" in PS3.4).

2060　See Section XXXX.7.2 for discussion of using non-DICOM protocols.

XXXX.2.3.3 Updated metadata

Many repository applications do not update stored SOP Instances with changes to metadata that occur over time (e.g., patient name and ID). Therefore, applications that use non-DICOM direct file access need to obtain the current metadata, which can be retrieved in a query using the Updated Metadata Sequence (00gg,0Fz2)
2065　Key Attribute, defined by the Repository Query SOP Class (see Section C.6.x4.1.4 "Updated Metadata Attributes" in PS3.4). It is the responsibility of the using application to use the metadata in these returned Attributes (see Section XXXX.7.9 "Metadata Updates").

XXXX.2.3.4 Study Update DateTime

Query of the Study Update DateTime (0020,121x) Key Attribute using date and time range matching allows an
2070　SCU to identify changes that have occurred since a prior inventory was obtained. This enables incremental processing of updates to synchronize the SCU's data to that held by the SCP, and is crucial to the migration use case. It may also support important quality assurance and other processes. In order to maintain backward compatibility with existing repository databases, this Attribute is identified as optional, but its importance to a number of use cases and its potential for significant performance improvement makes implementation highly
2075　desirable.

XXXX.3 THE INVENTORY INFORMATION OBJECT

XXXX.3.1 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
2080　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 "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.3-1. The Inventory is created by an identified piece
2085　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 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.

2090　Note

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
2095　associated Imaging Service Requests, the request attributes may be encoded at the Series level in the Request Attributes Sequence (0040,0275). The Inventory IOD includes the Request Attributes Sequence to support this use.
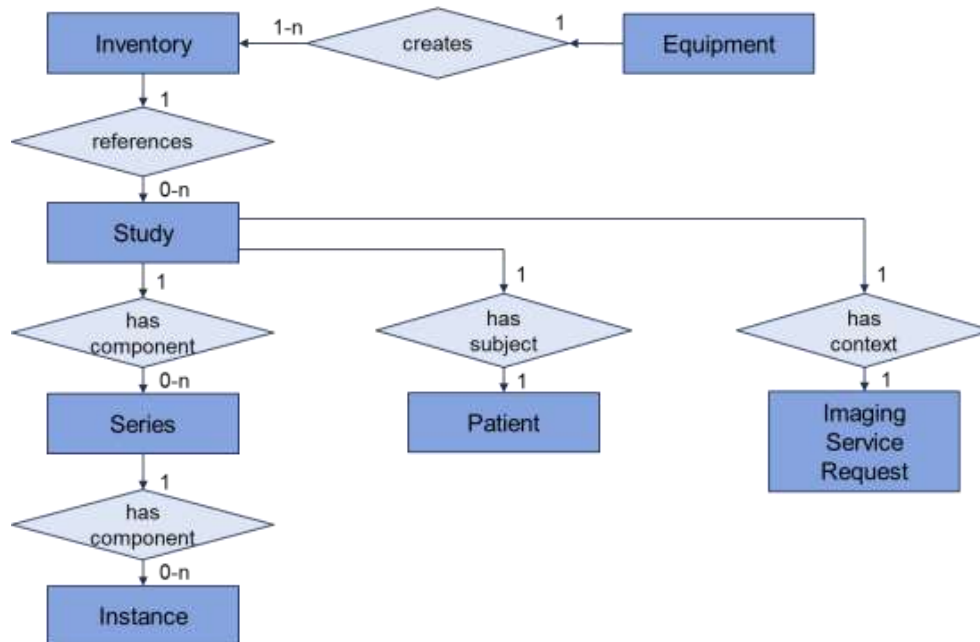
Figure XXXX.3-1. Inventory Information Model E-R Diagram
(reproduced from Figure 7.13.x-1 in PS3.3)

2100   These IOD Information Entities include all the required Attributes specified for Query processing by the repository system (see Section C.6 in PS3.4). An Inventory is thus a standard DICOM representation of the key content of a repository 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.

2105   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 many 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.

XXXX.3.2 Scope of Inventory

2110   The Inventory IOD supports Inventories of subsets of the repository based on a set of Key Attributes that specify the Scope of Inventory. The values of those Key Attributes are used to match the corresponding attributes of Studies in the repository, similar to the Key Attribute matching used in Query services, to select the Studies that are included in the Inventory (see Section C.YY.2.1 "Scope of Inventory Macro" in PS3.3). Any Key Attributes allowed for Query services can be specified in the Scope of Inventory (see Section XXXX.2.3 Key Matching Attributes).

2115   The scope is also implicitly limited to records available at 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, at the discretion of the implementation.

XXXX.3.3 Inventory Instance Tree

2120   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

2125   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
2130   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
2135 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.3-2. 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.
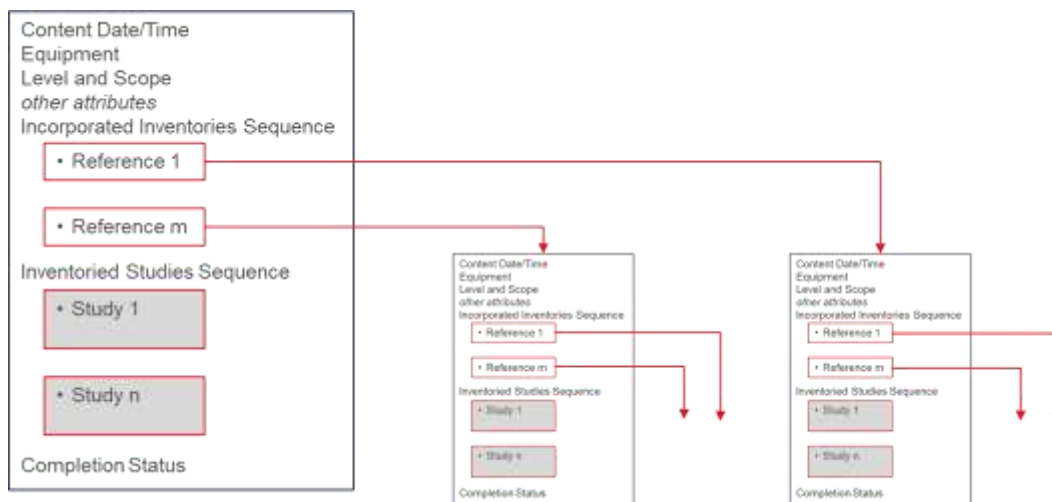
2140   **Figure XXXX.3-2. Inventory IOD Schematic Structure**

XXXX.3.3.1 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 the value 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,
2145 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
2150 accurately describes the content of the tree relative to the Scope of Inventory at the Content Date and Time for
the repository system identified in the General Equipment Module.

## XXXX.3.3.2 Examples

XXXX.3.3.2.1 Serial Production

As an example of how this tree structure might be used, consider an application producing a large Inventory. It
2155 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.3-3). The last object becomes the root of the tree of the complete inventory.
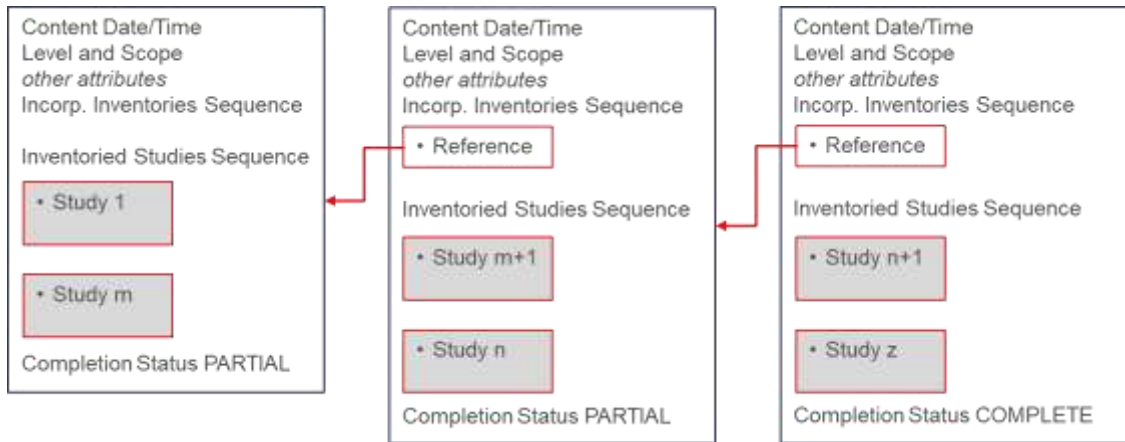
2160

Figure XXXX.3-3. 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.)

2165 **XXXX.3.3.2.2 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 2170 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).
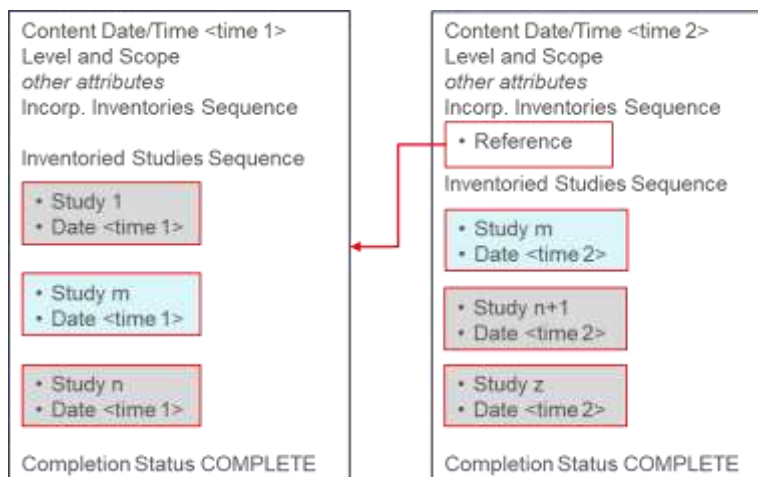
Figure XXXX.3-4. Baseline with incremental update

2175 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.3.3.2.3 Parallel Production

2180 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.3-4). Note that the Scope of Inventory will be the same for all objects, but the Equipment identifiers will differ.
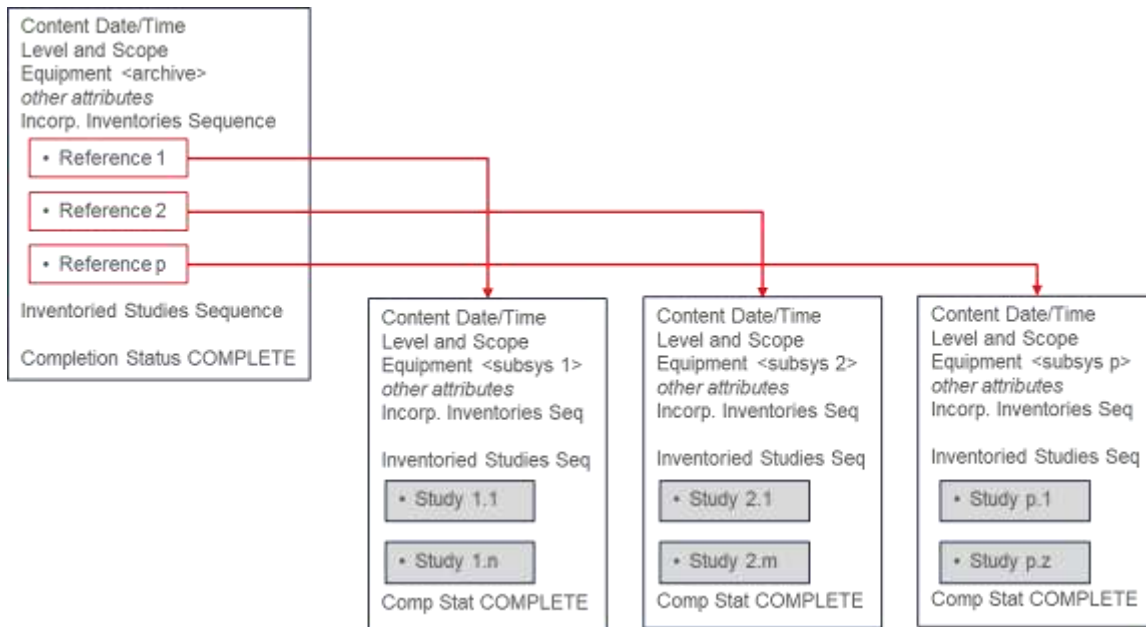
2185 **Figure XXXX.3-5. Federated or parallel production example**

XXXX.3.3.2.4 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.3-5), where each node is the root of a
2190 subtree or a terminal leaf.
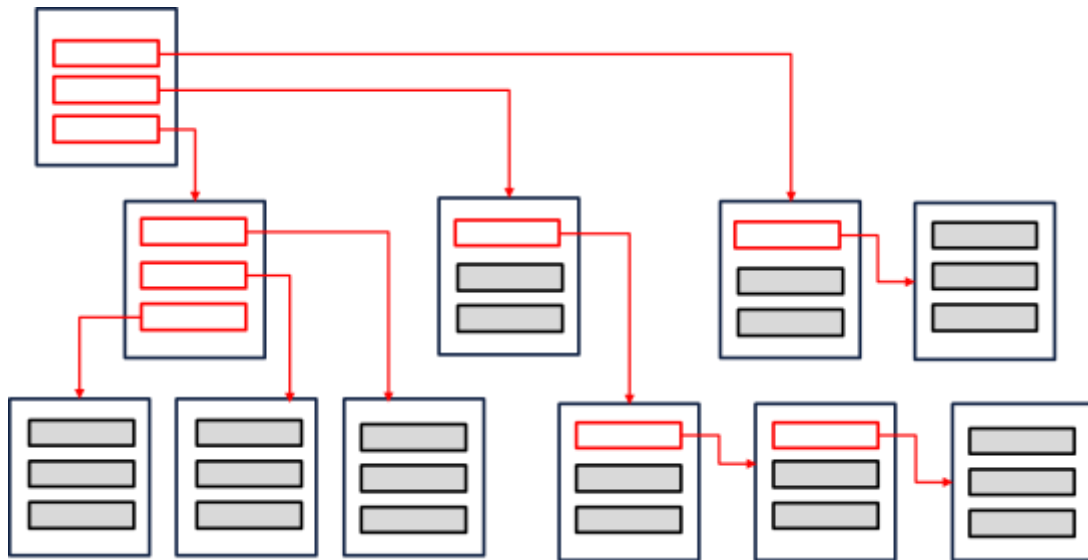
Figure XXXX.3-6. Arbitrary tree structure example

XXXX.3.3.2.5 Empty inventory

A repository system may be tasked with producing an Inventory, but for which there are no stored studies that 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 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
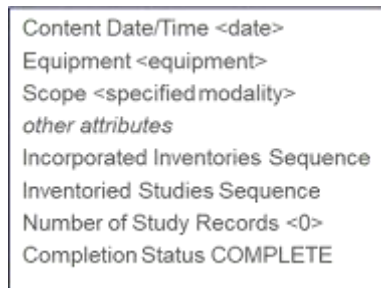
Figure XXXX.3-7. Empty inventory example

XXXX.3.4 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 (AE Title) or web (origin server address), or both, must be provided for each stored SOP Instance; the non-DICOM protocols are optional. See Section XXXX.7.2 for discussion of using non-DICOM protocols

Many repository applications do not keep the stored SOP 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.7.9).

XXXX.3.5 Additional Data Elements

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

2215 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 database record keys in Private Data Elements for implementation-specific processing.

> Note

2220 > 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 Instances in accordance with the [IHE RAD TF-1] *Evidence Documents Profile.*

XXXX.3.6 Producer vs. Consumer Implementation

2225 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 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 2230 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 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 2235 XXXX.5.1), 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 (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 2240 to introduce a single way of managing such data. Rather, the repository 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 Document SOP Instance, e.g., in accordance with the [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.

2245 ## XXXX.4 RELATED SERVICES FOR INVENTORY SOP INSTANCES

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.6).

XXXX.4.1 Inventory Storage and Query/Retrieve

The Inventory IOD is defined in the category of non-patient-related DICOM composite objects. As such, its basic 2250 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).

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 2255 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 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.

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 a whole to the destination system (see Section XXXX.7.6).

XXXX.4.2 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 Inventory.

The Inventory Creation SOP Class is in many ways similar to the Repository Query SOP Class (see Section XXXX.2). In both cases, the SCU requests a list of Studies, managed by the SCP, that match Key Attribute values. However, the Repository Query operates synchronously (i.e., the query and response occur on the same Association). The SCP is expected to respond within a typical transactional timeout period, and the SCU must interactively process responses and sequentially initiate queries to continue after partial completion responses (or errors).

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.

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 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.) 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.

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 Creation IOD use many of the same Attributes, including the Scope of Inventory Sequence.

XXXX.4.3 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
2305 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.

2310 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
2315 implement a DICOM method for accessing the produced Inventory instances.

Figure XXXX.4-1 illustrates the relationships of the Inventory-related services to the information definitions.
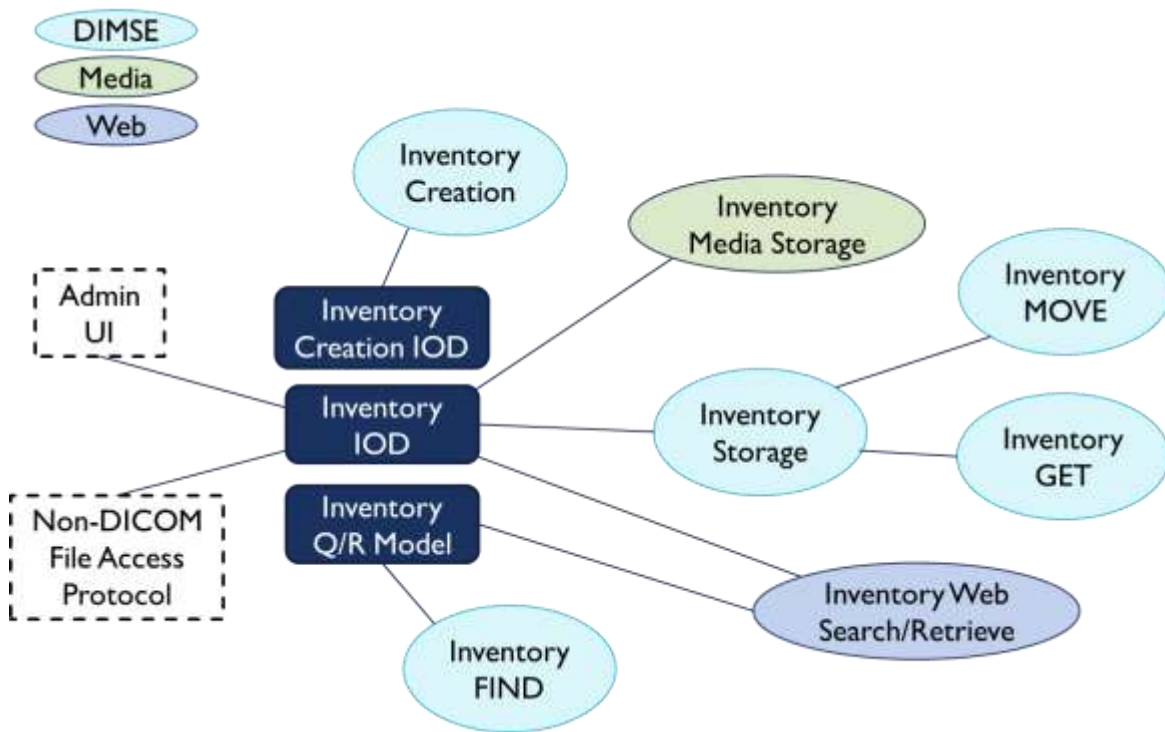
Figure XXXX.4-1 Inventory-related information definitions and services

As noted above, an application implementing the Inventory Creation SOP Class must support matching on all of
2320 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.5 USE CASES

XXXX.5.1 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 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 Repository Query and the Inventory IOD and its related Services. Additional steps may include data reconciliation between the source repository and the databases of the radiology 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 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 in their DICOM protocol implementation for the additional massive input/output requirements of migration.

The Inventory, whether obtained through the Repository Query responses or through Inventory SOP Instances, 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 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 includes optional access Attributes identifying available non-DICOM file access protocols 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 into a single container file (ZIP or TAR). The access Attributes may identify such container files.

At the destination repository, the process of building the local database for the incoming data may be facilitated by processing the Inventory, 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 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.

XXXX.5.2 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.

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

2370 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. Further, if the repository instances are in DICOM File
2375 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,
2380 patient-based Inventories, and more. The appropriate approach will vary by the particular needs and workflow of each organization.

XXXX.5.3 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,
2385 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
2390 operational data that can be sorted, filtered, and analyzed in any number of ways to support research questions.

The Inventory 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
2395 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
2400 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
2405 research purposes, but it may be sufficient as a first level selection that simplifies 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,
2410 and the overall medical record architecture and systems of the organization.

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.5.4 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. The Repository Query SOP Class and the Inventory IOD's Scope of Inventory allow Study selection using the extended (optional) capability "empty value matching". If such matching is implemented in the repository system, it allows creation of an Inventory of Studies with empty Attributes. As a quality assurance process, such inventories may be produced on a regular basis, identified studies corrected as needed, and root causes for missing values identified and corrected as a process improvement task.

XXXX.5.5 Wellness Check / Continuous Testing

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 repository 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.

XXXX.6 SECURITY CONSIDERATIONS

XXXX.6.1 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 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 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 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.

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.6.1.1 Access Control in Production of Inventory

A repository might limit disclosure or retrieval of SOP instances, studies, or patients following a variety of authorization policies and data protection rules, often based on the user's identity and/or attributes of the

instance data. Such limits might be implemented at different layers of the repository software architecture (file system, database management system, application, etc.).

The identity of the initiator of Inventory production might be known to the production application, e.g., if initiated from a local user interface, or if conveyed in the secure transport layer of the Inventory Creation service. That
2465 user identity might affect the content of an Inventory by triggering various data protection rules.

The Inventory IOD has no means of identifying whether such protection rules have been invoked, and thus whether the inventory may be incomplete with respect to the restricted data. In some cases, the fact that protection rules have been invoked, or even the existence of such rules, is not disclosed to using applications.

The implementers and users of an Inventory production application should be cognizant of the potential effect of
2470 user identity and permissions on the content of the produced Inventory SOP Instances. Implementers may disclose in the product DICOM Conformance Statement section on Application Level Security any access control features that might impact Inventory production. Users should verify that access controls are not inappropriately impacting Inventory production.

XXXX.6.2 File Format

2475 The DICOM File Format has security considerations that will apply whenever that format is used, e.g., for the 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
2480 files should utilize appropriate defenses and safeguards such as:

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

2485 Applications that store container files for later use by other systems should consider the environments of those 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
2490 security breach of the source system or other upstream system.

XXXX.6.3 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
2495 implementation and use of the protocol. The specific details of any such considerations are outside the scope of 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
2500 administrative issues.

XXXX.6.4 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

2505 data (i.e., the entire DICOM repository) should be thoroughly validated with regard to appropriateness of data 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
2510 Inventory SOP Instance. Applications for which current metadata is required should be specifically validated to ensure current metadata is applied.

XXXX.6.5 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).
2515 This may involve special authorization for initiating broad inventories, and appropriate setting of software task priorities for the Inventory application.

XXXX.6.6 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
2520 also ensures file integrity. DICOM specifies methods for secure (encrypted) files (see Annex D "Media Storage 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 repository 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
2525 a file encrypted by another mechanism. There are no specifications regarding key management to access that 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.6.7 Message Digest

2530 The integrity of a stored SOP Instance file (unencrypted) may be verified by a Message Authentication Code (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.

Inventories support recording a MAC computed by the storage system (the writing application) for files in the
2535 repository that will be accessible through a non-DICOM protocol. The file reading application can independently perform the MAC computation to assure integrity of the file as read or transferred.

XXXX.6.8 De-identification

While some research use cases may involve de-identification of protected health information (PHI), where that de-identification occurs in the data processing pipeline may vary with the specific research objective and the
2540 capabilities of the systems involved. DICOM specifies a profile with many options for de-identification of SOP Instances, the Basic Application Level Confidentiality Profile (see Annex E in PS3.15). That specification is designed for patient-related SOP Instances with PHI attributes in the top-level data set, and there would be substantial technical challenges to applying that profile to an Inventory SOP Instance.

However, an Inventory may be produced for a repository of de-identified Studies. That is, the SOP Instances in
2545 the repository are first de-identified in accordance with a confidentiality profile and options appropriate to the specific research use case, and then an Inventory is produced for the repository, or for a subset thereof in accordance with the Scope of Inventory. There are no specific PHI de-identification requirements on the Inventory itself.

XXXX.7 OPERATIONAL CONSIDERATIONS

This section describes topics relevant to implementation and use of Inventories.

XXXX.7.1 Transforming Repository Query Responses into Inventory SOP Instances

For implementation reasons, there will be situations where the repository implements the Repository Query SOP Class, but the using application wants to work asynchronously from an Inventory SOP Instance. Because the Repository Query SOP Class and the Inventory IOD are aligned technically, it is feasible for an intermediary application to transform Repository Query responses into Inventory SOP Instances.

The basic operation is for the using app to first inform the intermediary app about its desired Scope of Inventory, which can be done through the Inventory Creation SOP Class or some non-DICOM method (such as manual configuration). The intermediary app performs Repository Queries at the Study level using the desired Key Attributes, including universal matching for all Attributes that will be included in the Inventory (minimally, the Type 1 and Type 2 Attributes specified in the Inventory IOD). For each response, the app performs a query for the Series, and for each Series a query for instances, down to the level for which the Inventory is being produced. The responses from the three hierarchical levels are encoded in the Sequence Items of an Inventory SOP Instance. A separate Inventory SOP Instance might be created for each Study level Repository Query transaction, with the Instances chained as described in Section XXXX.3.3.2.1 Serial Production.

The intermediary app must account for Attributes not supported for matching by the SCP (see Section XXXX.7.11). Such Attributes need to be excluded from the Scope of Inventory in the SOP Instances as they were not used for selection of Studies.

The process for transforming a Relational query to an Inventory SOP Instance is more complex, and requires knowledge of the hierarchical level of each Attribute. The initial query is not at the Study level, but rather a relational query at the Series or Instance level. Depending on the Key Attributes requested to be returned for the level(s) higher than the Query level, the app may need to perform follow-on queries at those higher levels. In particular, the Attributes defined at multiple levels (such as Removed from Operation Use (00gg,0Fx6) and File Set Access Sequence (00gg,0FyE)) need to be requested at each specific level.

While such transforms are relatively straightforward, there are some differences between Repository Query and the Inventory IOD that need to be addressed. First, the specification of Key Attributes differs, as the Query uses constructs unavailable to SOP Instances, and the IOD uses Sequence Attributes for different types of matching. Second, the Query uses the Updated Metadata Sequence (00gg,0Fz2) to be able to obtain changed attributes without needing to enumerate them, while in the Inventory IOD all attributes are encoded in the top-level data set for the appropriate entity Sequence Item. Of course, the intermediary application needs to handle error conditions that may occur, which should be expected in a process that may extend over several days, with applications that may also be involved in other tasks that may interrupt the inventory production.

XXXX.7.2 Using Non-DICOM Protocols

The direct file access URI links of the Repository Query SOP Class and the Inventory IOD do not limit the protocol used, which is specified by the scheme of the URI (e.g., "https:", "nfs:", "smb:", etc.). Applications that intend to use direct file access may need to be adapted to use the protocol specified by the repository. Not all capabilities of the access mechanism may be evident from the URI scheme, e.g., HTTP is used with several different cloud-based storage protocols that differ in the ways they use HTTP headers. The specifics of the protocol are conveyed in the Conformance Statement, rather than in DICOM Attributes.

The target resource of a non-DICOM protocol must be a SOP Instance stored in the DICOM File Format as specified in Section 7 in PS3.10 (commonly denoted a Part 10 file). At the SOP Instance level, the link may be a complete URI conveyed in the Stored Instance File URI (00gg,0FxB). However, many (or all) of the links for SOP Instance files of a Study or Series are likely to have the same URI base. The Inventory can factor out that commonality by specifying a Stored Instance Base URI (00gg,0Fx9) at the Study and/or Series level and using relative reference URIs (starting with ./) for the individual Instances.

The Part 10 file can also be contained within a defined container file structure – ZIP, TAR, GZIP, or TARGZIP. In that case, the Stored Instance File URI (00gg,0FxB) links to the container file, and the specific Part 10 file is identified by Filename in Container (00gg,0FxD).

2600 If all the files for a Study or a Series are in a single container file, the Inventory Study or Series record can specify a URI link to that file. Similarly, if all the files for a Study or a Series are in a single folder (which is an operating system "container" mechanism), the Inventory can specify a URI link to that folder.

These permutations are shown in Table XXXX.7-1, and an example is shown in Figure XXXX.7-1.

Table XXXX.7-1 Use of URI-related attributes

Use	URI Attributes Used	Notes
SOP Instance in a Part 10 file; File URI is complete	File Access Sequence (00gg,0FyF) >Stored Instance File URI (00gg,0FxB)	
SOP Instance in a Part 10 file; File URI is relative reference	File Set Access Sequence (00gg,0FyE) >Stored Instance Base URI (00gg,0Fx9) File Access Sequence (00gg,0FyF) >Stored Instance File URI (00gg,0FxB)	Base URI specified in Series level File Set Access or, if not there, in Study level File Set Access
SOP Instance Part 10 file in a container file; File URI is complete	File Access Sequence (00gg,0FyF) >Stored Instance File URI (00gg,0FxB) >Container File Type (00gg,0FxC) >Filename in Container (00gg,0FxD)	
SOP Instance Part 10 file in a container file; File URI is relative reference	File Set Access Sequence (00gg,0FyE) >Stored Instance Base URI (00gg,0Fx9) File Access Sequence (00gg,0FyF) >Stored Instance File URI (00gg,0FxB) >Container File Type (00gg,0FxC) >Filename in Container (00gg,0FxD)	Base URI specified in Series level File Set Access or, if not there, in Study level File Set Access
All files in Study or Series in a container file; File URI is complete	File Set Access Sequence (00gg,0FyE) >Stored Instance File URI (00gg,0FxB) >Container File Type (00gg,0FxC)	
All files in Study or Series in a container file; File URI is relative reference	File Set Access Sequence (00gg,0FyE) >Stored Instance Base URI (00gg,0Fx9) >Stored Instance File URI (00gg,0FxB) >Container File Type (00gg,0FxC)	For Series, if Base URI is not specified at Series level, it must be specified at Study level
All files in Study or Series in a folder; Folder URI is complete	File Set Access Sequence (00gg,0FyE) >Folder URI (00gg,0FxA)	
All files in Study or Series in a folder; Folder URI is relative reference	File Set Access Sequence (00gg,0FyE) >Stored Instance Base URI (00gg,0Fx9) >Folder URI (00gg,0FxA)	For Series, if Base URI is not specified at Series level, it must be specified at Study level

2605   **Table XXXX.7-2 Example Use of URI-related Attributes**

Attribute			Tag	VR	Value	Notes
…Inventory attributes						
Inventoried Studies Sequence			(0400,06x1)	SQ		
Item Study	>…*Study attributes*					
	>File Set Access Sequence		(00gg,0FyE)	SQ		Note 1
	Item	>>Stored Instance Base URI	(00gg,0Fx9)	UR	nfs://vna.exampleinstitution.org/ JZ08555/2.25.916804767294.zip	
	>Inventoried Series Sequence		(0400,06x2)	SQ		
	Item Series 1	>>… *Series attributes*				
		>>File Set Access Sequence	(00gg,0FyE)	SQ		Note 2
		Item	>>>Stored Instance File URI	(00gg,0FxB)	UR	./	
			>>>Container File Type	(00gg,0FxC)	CS	ZIP	
		>>Inventoried Instances Sequence	(0400,06x3)	SQ		
		Item Inst 1.1	>>> …*Instance attributes*				
			>>>File Access Sequence	(00gg,0FyF)	SQ		Note 3
			Item	>>>>Stored Instance File URI	(00gg,0FxB)	UR	./	
				>>>>Container File Type	(00gg,0FxC)	CS	ZIP	
				>>>>Filename in Container	(00gg,0FxD)	UR	2.25.192734871000545.dcm	
				>>>>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	UI	1.2.840.10008.1.2.4.70	
		Item Inst 1.2	>>> …*Instance attributes*				
			>>>File Access Sequence	(00gg,0FyF)	SQ		Note 3
			Item	>>>>Stored Instance File URI	(00gg,0FxB)	UR	./	
				>>>>Container File Type	(00gg,0FxC)	CS	ZIP	
				>>>>Filename in Container	(00gg,0FxD)	UR	2.25.192734871076985.dcm	
				>>>>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	UI	1.2.840.10008.1.2.4.70	
	Item Series 2	>>… *Series attributes*				
		>>Inventoried Instances Sequence	(0400,06x3)	SQ		
		Item Inst 2.1	>>> …*Instance attributes*				
			>>>File Access Sequence	(00gg,0FyF)	SQ		Note 4
			Item 1	>>>>Stored Instance File URI	(00gg,0FxB)	UR	nfs://vna.exampleinstitution.org/ JZ08555/2.25.460800590520.dcm	
				>>>>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	UI	1.2.840.10008.1.2.4.70	
			Item 2	>>>>Stored Instance File URI	(00gg,0FxB)	UR	smb://pacs.exampleinstitution.org/ cachesrv/2.25.460800590520.dcm	
				>>>>Stored Instance Transfer Syntax UID	(00gg,0Fy0)	UI	1.2.840.10008.1.2.1	

Notes to Table XXXX.7-2

1. Most of the Study content is in a ZIP file, but not all, so Stored Instance File URI cannot be used at Study level. However, the Study level sets the Base URI, which is a complete URI referencing the ZIP file.

2610   2. All of the content of Series 1 is in the ZIP file, so the Series level Item can include the Stored Instance File URI. Since the Base URI is not present at the Series level, it defaults to the Study level Base URI, which is the complete URI reference to the ZIP file. The Stored Instance File URI can then be a relative reference that adds nothing, and so it links to the ZIP specified at the Study level.

3. Since the Base URI is not present at the Series level, the Instances of Series 1 also default to the Study level Base
2615   URI. Their filenames within the ZIP are identified.

4. The SOP Instance of Series 2 is not in the ZIP, so its Stored Instance File URI is a complete URI reference to the Part 10 file. There are two copies known to the inventory, and they are each referenced.

Although the Inventory identifies the available access mechanisms for repository stored instances, the security features associated with those access mechanisms and with container file structures are outside the scope of
2620   DICOM, and will need to be implemented in applications that use the Inventory (see Section XXXX.6).

XXXX.7.3 Using Referenced Inventories

Section XXXX.3.3 describes the tree of Inventory SOP Instances whose contents are included by reference in the complete Inventory described by the root SOP Instance. A user of the Inventory may retrieve referenced Inventories in the tree through the Inventory Query/Retrieve Service (see Section XXXX.3.3), or the DICOM
2625   web-based Non-Patient Instance Service, if either of those is implemented in the system. The Inventory IOD may also include alternative access information for a non-DICOM file access protocol with each link to a referenced Inventory SOP Instance.

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
2630   PS3.3). This is used to encode a tree structure containing the entire set of links for the tree of which it is the root.

Therefore, when an application creates an Inventory and includes another Inventory by reference, it adds 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.7-2). Note
2635   that including the entire tree of object references ensures that the tree is acyclic.
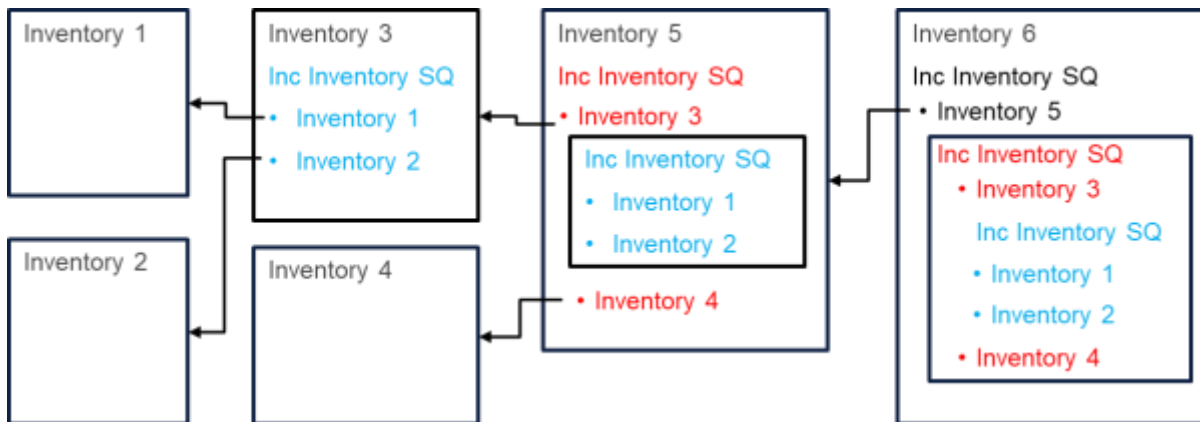

Figure XXXX.7-2 Inclusion of Inventory references

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

2640 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 own Inventoried Studies Sequence (0400,06x1).

XXXX.7.4 Incremental Inventories

Like all DICOM composite objects, Inventory SOP Instances are static, so an Inventory of a repository in
2645 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 update since the prior inventory if needed.

The IOD design allows a full Inventory to be created with an incremental update Inventory object including a baseline Inventory by reference (see Section XXXX.3.3.2.2), thus minimizing the processing resource cost of a
2650 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 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
2655 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 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
2660 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, 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
2665 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 / 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
2670 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 (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).

2675 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.

XXXX.7.5 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
2680 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 Inventory 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.5.2), may have policies to create a complete Inventory each month, to maintain the two most recent Inventories, and
2685 automatically delete prior ones. Such a policy would allow the assignment of a value for Expiration DateTime (00gg,0Fy9) for the Inventory SOP Instances.

An organization might 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 only with a Study Update DateTime range after the last full Inventory Content Date / Time.

XXXX.7.6 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 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 server supports HTTP Range request headers, the destination application can interactively retrieve specific byte ranges of the SOP Instance using the mechanism of [RFC7233].

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.7.7 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 services (for example, as shown in Figure XXXX.3-4).

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 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.

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

XXXX.7.8 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).

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, even if those are not managed in the repository database. Such values may especially facilitate consolidation of multiple repositories.

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 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 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 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.

XXXX.7.9 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 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 implementation stores received SOP Instances to disk in the DICOM File Format, but 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 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.7.9.1 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 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 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.7-2 shows what would be a portion of an Inventory SOP Instance for a study where the Patient's 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.

Table XXXX.7-2. 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	
>>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)
2780 already in the Original Attributes Sequence of the stored SOP Instance. While there may be duplication, duplicate Items are not an issue for the audit purposes of the Original Attributes Sequence.

XXXX.7.10 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
2785 of the contents of the Study as known by the creator of that Item.

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.

2790 DICOM is not prescriptive regarding methods of reconciliation, but the Inventory IOD does provide Attributes that 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.7-3. 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
2795 with the presence of an additional Series whose Series Date is after the prior Item Inventory DateTime. The later record might 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.

2800 **Table XXXX.7-3. 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)

In general, a major factor in reconciling diverse records is a full understanding of how the repository system manages the storage of Studies, 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
2805 Inventory SOP Instance attributes or in DICOM Conformance Statements.

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
2810 level record.

XXXX.7.10.1 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
2815 system might have taken one of several approaches, with the associated result in the time B inventory (this is not an exhaustive list):

1) It marks the Study as deprecated, but otherwise retains the data – the time B incremental inventory includes the entire set of Study, Series, and Instance records, each with the Removed from Operational Use attribute value Y.

2) It marks the Study as deprecated, and deletes all the Series and Instance data – the time B incremental inventory includes only the Study record with the Removed from Operational Use attribute value Y.

3) 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 includes a Study item, but no Series items.

4) It deletes all Study information, with only a record in an audit trail – the time B incremental inventory simply does not record the Study.

In cases 1) and 2), 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 system takes to identifying and managing deleted Studies.

In case 3), 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 4), without an explicit Study record 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, and this situation may trigger an audit investigation.

XXXX.7.11 Key Attributes Unsupported for Matching

In DICOM Query, when the SCU requests matching on optional Key Attributes that are not supported for matching by the SCP, the baseline response behavior is for the SCP to treat them as "universal match", i.e., no filtering is performed by the SCP. In the Repository Query or Inventory Creation SOP Classes, such behavior may result in a substantial number of records not desired by the SCU being returned. For example, the SCU may request inventory of Studies updated in the last year by specifying a date range match on Study Update DateTime. If the SCP does not support matching on that attribute, the baseline behavior would be to return inventory of all Studies in the repository. This could have significant performance impacts on both the SCU and SCP.

The Inventory Creation SOP Class specifies a Warning response to an Initiate N-ACTION request, B010 – "One or more of Key Attributes are not supported for matching", with the list of unsupported Attributes provided in the N-ACTION response field Attribute Identifier List (0000,1005). The SCU can evaluate the Warning and, if desired, send a Cancel N-ACTION.

The Repository Query SOP Class does not provide such a warning. However, the SCP's Conformance Statement is required to identify Attributes supported for matching, although if that list is site-configurable the Conformance Statement may not provide the requisite information. The SCU could, however, request a relatively small Maximum Number of Records (00gg,0Fz2) in the initial Query, evaluate the Query responses, and check that responses do not exceed the requested match values before continuing with a subsequent Query.

DICOM PS 3.18: Web Services

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

2865 **12.1 OVERVIEW**

The Non-Patient Instance (NPI) Storage Service enables a user agent to retrieve, store, and search an origin server for instances that are not related to a**n individual** patient.

> **Notes:**
>
> 2870 1. **Non-Patient Instances adhere to a Composite Instance IOD Information Model that does not have at its root the Patient Information Entity representing an individual Patient.**
>
> 2. **"Non-patient" does not imply that there is no patient-related protected health information (PHI) in the Instances. E.g., the Inventory IOD does include Attributes of PHI, but it does not have a Patient IE at the root of its information model.**

An NPI Storage Service manages a collection of resources belonging to the categories specified in Table 2875 12.1.1-1.

All NPI Storage Service origin servers shall support the Retrieve Capabilities, Retrieve, and Search transactions. Support for the Store transaction is optional. All NPI Storage Service user agents support one or more of the Retrieve Capabilities, Retrieve, Store, or Search transactions.

12.1.1 Resource Descriptions

2880 An NPI Service manages resources from the same NPI Category. Target URIs have the following templates:

/{npi-name}
/{npi-name}/{uid}

Where

2885 npi-name   = "color-palettes"
        / "defined-procedure-protocols"
        / "hanging-protocols"
        / "implant-templates"
        / "inventories"
2890 uid   ; is the Unique Identifier of an NPI Instance

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

Resource Category	URI Template and Description	Corresponding IOD	Storage Class	Information Model
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**

2895 The NPI SOP Classes are listed in Table GG.3-1 "Standard SOP Classes" in PS3.4.

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

---

2900 ## 12.6.1.2 Query Parameters

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
2905 in the corresponding sections in Table 12.6.1-2.

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**