SUPPLEMENT 223: REPOSITORY QUERY, INVENTORY IOD, AND RELATED SERVICES

DICOM WG-33 Data Archive and Management

2022/06/10

History

- WG-33 established by DSC December 2019, SIIM as Secretariat
  - In response to Laitek proposal from April 2019
- Substantive work began March 2020, with broad user and vendor participation
  - Consensus technical approach for Inventory IOD agreed August 2020
- Work Item approved by DSC August 2020
- WG-6 approved Draft Supplement for Public Comment April-May 2021
- Comment from GE requested equivalent query-based mechanism
- DSC direction to include query in Sup 223 August 2021
  - WG-33 and WG-6 consensus for dual approach (IOD and query)
- WG-6 approved revised Draft Supplement for 2nd Public Comment Nov 2021-Jan 2022
- WG-6 approved Letter Ballot April-June 2022

Primary Use Case – PACS migration

- Users typically replace PACS after ~12-15 years, often with change of vendor
  - Became a significant user (PACS administrator) concern in 2010's
- Institutional consolidation often requires merge of disparate PACS
- Replacement/consolidation requires migrating historical data to a new system
  - Digital storage often retained "forever"
  - Volumes increasing with hi-res, 3D/4D, and multimodality studies
  - Many institutions store > 10^9 instances with data volumes > 10^{15}B (petabyte)
- Performance at scale is challenging
  - DICOM Query/Retrieve designed for accessing relatively small number of studies
  - Need mechanisms that account for enterprise-scale data operations

But Sup 223 is not a PACS migration standard

- PACS migration is a multi-faceted endeavor, some aspects of which can be supported by DICOM standards, and others that cannot
- Sup 223 addresses one critical need to support migration – production of a standardized inventory of the DICOM content of a repository
  - Required to prepare for migration, and to validate completeness of migration
  - Many other processes involved in execution of migration
- Sup 223 is a standard for "Inventory IOD and related services"
- Inventory object and services may also support other use cases

Secondary Use Case – Clinical Research and Business Analytics

- Inventory can be important data source (in conjunction with EMR and image instances) for research and analytics
  - Representation of the PACS database in standard non-proprietary format
  - Can be searched or processed without impacting PACS in clinical operation
  - Transferrable to research IT system (data warehouse) – "extract" in paradigm of Extract/Transform/Load (E/T/L)
- Inventory may be produced for a specified selection of studies based on modality and/or study or series description
- De-identification as required in research not in scope of Sup 223
  - De-ID is part of E/T/L "transform", and highly dependent on merge process with other data sources

Secondary Use Case – Safety Backup

- Image access is a patient safety / mission critical resource
- PACS DBMSs typically have fault tolerant designs
  - E.g., redundant online storage and offline backups
  - But data is in proprietary format and dependent on the DBMS for effective use
- DBMS becomes single point of failure
  - May become inoperable, e.g., due to license key expiration or malware attack
- Inventory is a DBMS-independent replica of the critical data content of the PACS database for the managed DICOM SOP Instances
- If the repository instances are in DICOM File Format and referenced in the inventory, there is a complete alternate path to access the images

Secondary Use Case – Quality Control

- Inventory may be produced for studies missing critical attribute values (optional capability of inventory creator)
  - Patient ID, Accession Number, Study Date
  - Allows proactive correction
- Comparison data for wellness check (continuous testing) of PACS operation
  - C-FIND result for sample of studies compared to data in previous inventory (e,g., metadata, number of images)
  - C-MOVE retrieve objects compared to direct filesystem access objects
  - Message digest (MAC) computed on objects retrieved through filesystem compared to MAC in previous Inventory

DICOM
Digital Imaging and Communications in Medicine

Overview of Sup 223

Two approaches for producing inventory

Query-based (synchronous)

- Extension of classic C-FIND to support return of entire database across multiple query transactions
- Minimal additional requirements for existing software facilitates implementation in deployed systems

Persistent Object-based (asynchronous)

- Exchangeable information object that can be processed off-line
- Creation initiated by repository function or by associated DICOM service
- Managed like other DICOM non-patient objects

Both produce Inventory of entire repository (or described subset) in non-proprietary format

- Hierarchical structure based on DICOM information models
- Encoded using DICOM formats and data structures

DICOM
Digital Imaging and Communications in Medicine

Repository Query Service

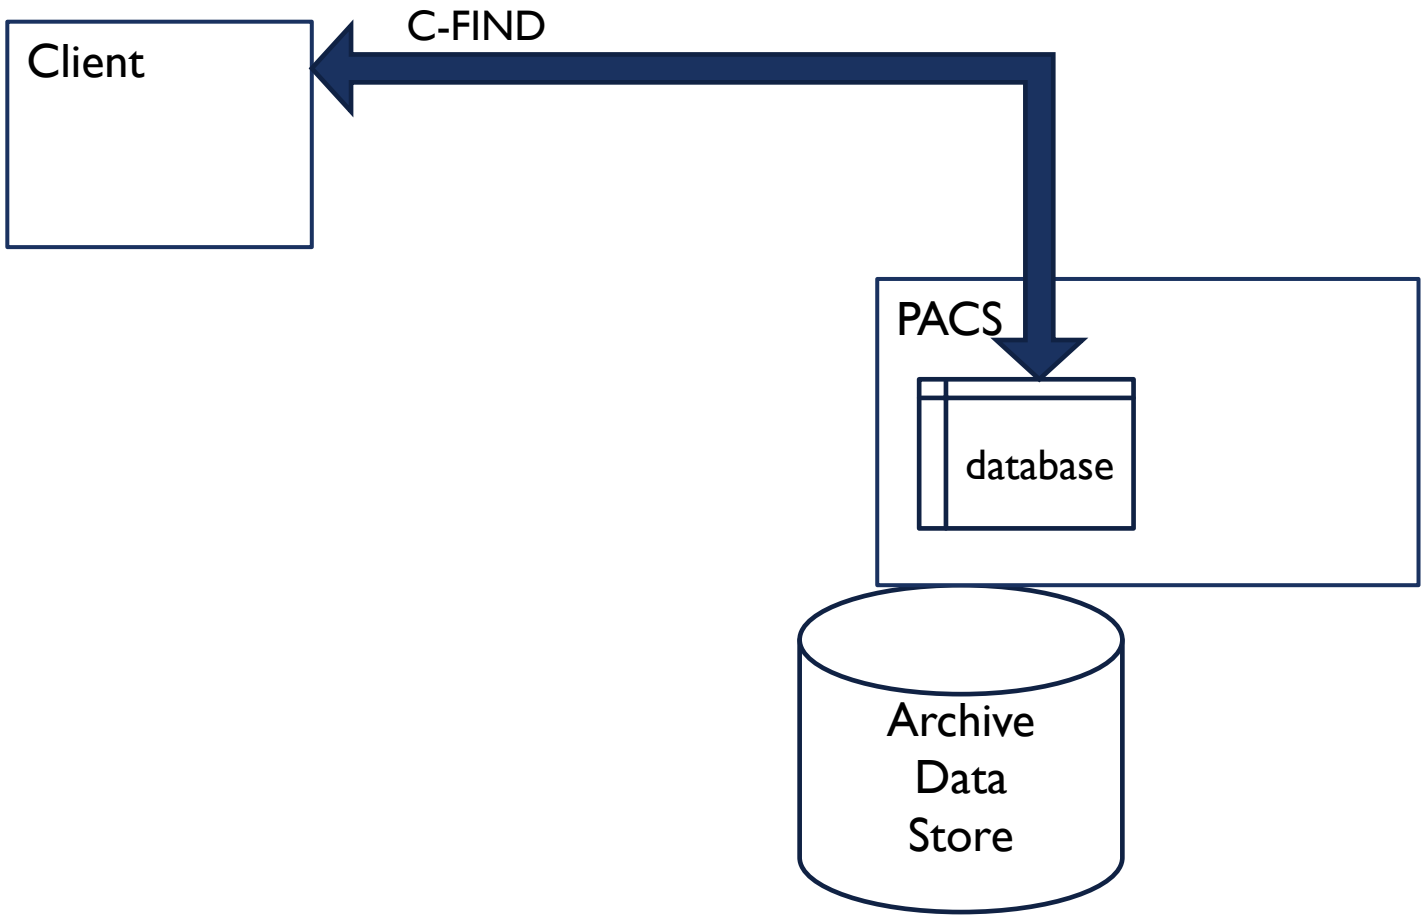

Repository Query obtains a complete inventory in a sequence of query transactions

Query features

- Repository Query SOP Class is limited extension of Study Root Query SOP Class with semantics of partial results/continuation
  - Three new required Attributes
  - Does not require Extended Negotiation, nor change to any other C-FIND semantics
- Additional Key Attributes supporting inventory use case are optional
  - Additional Attributes optional for both SCU and SCP, explicitly requested in Query
  - Query Response like all other matching Key Attributes
- Two new optional Key Attribute matching mechanisms
  - Agreed through Extended Negotiation, like other enhanced matching (fuzzy semantic, combined date time, etc.)

Partial result and continuation

- SCP and SCU may limit number of query response records returned
- Each returned record includes a Record Key
- SCU may send Prior Record Key in next Query
- SCP resumes responses with next Record Key

Partial result

- SCP and SCU may limit number of response records returned
  - SCP internal limit (number, duration of processing, time of day, etc.)
  - SCU requested maximum number
  - Purpose 1: allow inventory to be retrieved across multiple Query transactions, allowing management of resources on both SCP and SCU sides
  - Purpose 2: addressing possible network communication failures during long duration synchronous association
- When limit is reached, SCP sends final message with appropriate status code indicating incomplete response to query
  - New status code B001 Warning (C-FIND Warning added to PS3.7)
  - Similar to C-MOVE B000 Warning when some Instances could not be sent

Record Key

- Each returned record includes a Record Key
  - Must be requested by SCU (Type R Key Attribute, universal match ONLY)
- Record Key is opaque (to SCU) binary string of implementation-specific length (VR=OB)
- Response records ordered by Record Key
- SCU may send Prior Record Key in next Query
  - Last received Record Key upon receipt of B001 Warning status, or upon abnormal end of network association
- SCP resumes responses with next Record Key

Record Key characteristics

- Responses must be ordered by Record Key to allow continuation in subsequent Query
- Handling of inserted new record keys during sequential set of Queries is implementation-specific
- Record Key may include other state information for continuation
  - SCP does not need to retain such state information, as it will be returned in next Query
  - Implementation specific key VR=OB allows carrying of as much state information as desired
- UID is usually not sufficient as Record Key

New Repository Query SOP Class defines behavior

- Use of Warning Status and Record Key defined in new SOP Class
  - Partial results imperfectly defined for existing C-FIND SOP Classes, and we're not fixing that
  - No resumption behavior in existing C-FIND SOP Classes
- New capability agreed by basic negotiation of SOP Class

Additional Query features

- Matching mechanisms and Extended Negotiation
- Generally available Key Attribute – Study Update DateTime
- Repository Query specific Key Attributes
- Repository Query control attributes

Multiple Value and Empty Value Matching

- For attributes with VM > 1, Key Attribute with multiple values matches target attribute if all Key values are present in target
  - Example purpose: find multi-modality studies by matching against Modalities in Study
- Key Attribute value "" matches target Attribute with no value (or absent)
  - Example purpose: find studies missing date, accession number, patient ID, etc. values for quality assurance
- New matching capabilities optional, but generally available for all C-FIND
  - Agreed through new Extended Negotiation items (similar to combined date time or fuzzy semantic matching extensions to base matching capabilities)
  - Modification to DA, DT, and TM VRs in PS3.5 to allow " character in Query (similar to allowing – for range matching of DA)

Study Update DateTime

- New attribute for Study Update DateTime
  - 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
  - Defined at Study Level, but not in composite SOP Instances (similar to Number of Study Related Instances (0020,1208))
  - Example purpose: find Studies that were added or updated since the previous inventory
- New optional Key Attribute generally available for all C-FIND
  - Strongly recommended for implementation in all PACS/VNAs

Repository Query specific Key Attributes

- Further information in description of Inventory IOD
- Removed from Operational Use, Reason for Removal Code Sequence
  - Supports deprecated objects
- File Set Access Sequence, File Access Sequence
  - Non-DICOM protocol file access URIs at Study/Series level and Instance level, respectively
- Metadata Sequence, Updated Metadata Sequence
  - Mechanism to request/return all metadata attributes managed in PACS DB, or all updated attributes not propagated into stored instance
  - Support for one or both attributes required if Non-DICOM protocol file access supported
  - Universal match on SQ allows SCP to populate Item as appropriate without SCU needing to enumerate Attributes

Repository Query control attributes

- Prior Record Key, Maximum Number of Records (see above on Partial result and continuation)
  - Control attributes similar to Specific Character Set, Timezone Offset From UTC
- Considered alternative of Extended Negotiation
  - Undesirable due to more complex implementation
  - Needs to be fundamental to SOP Class (not a negotiable option)

Inventory IOD

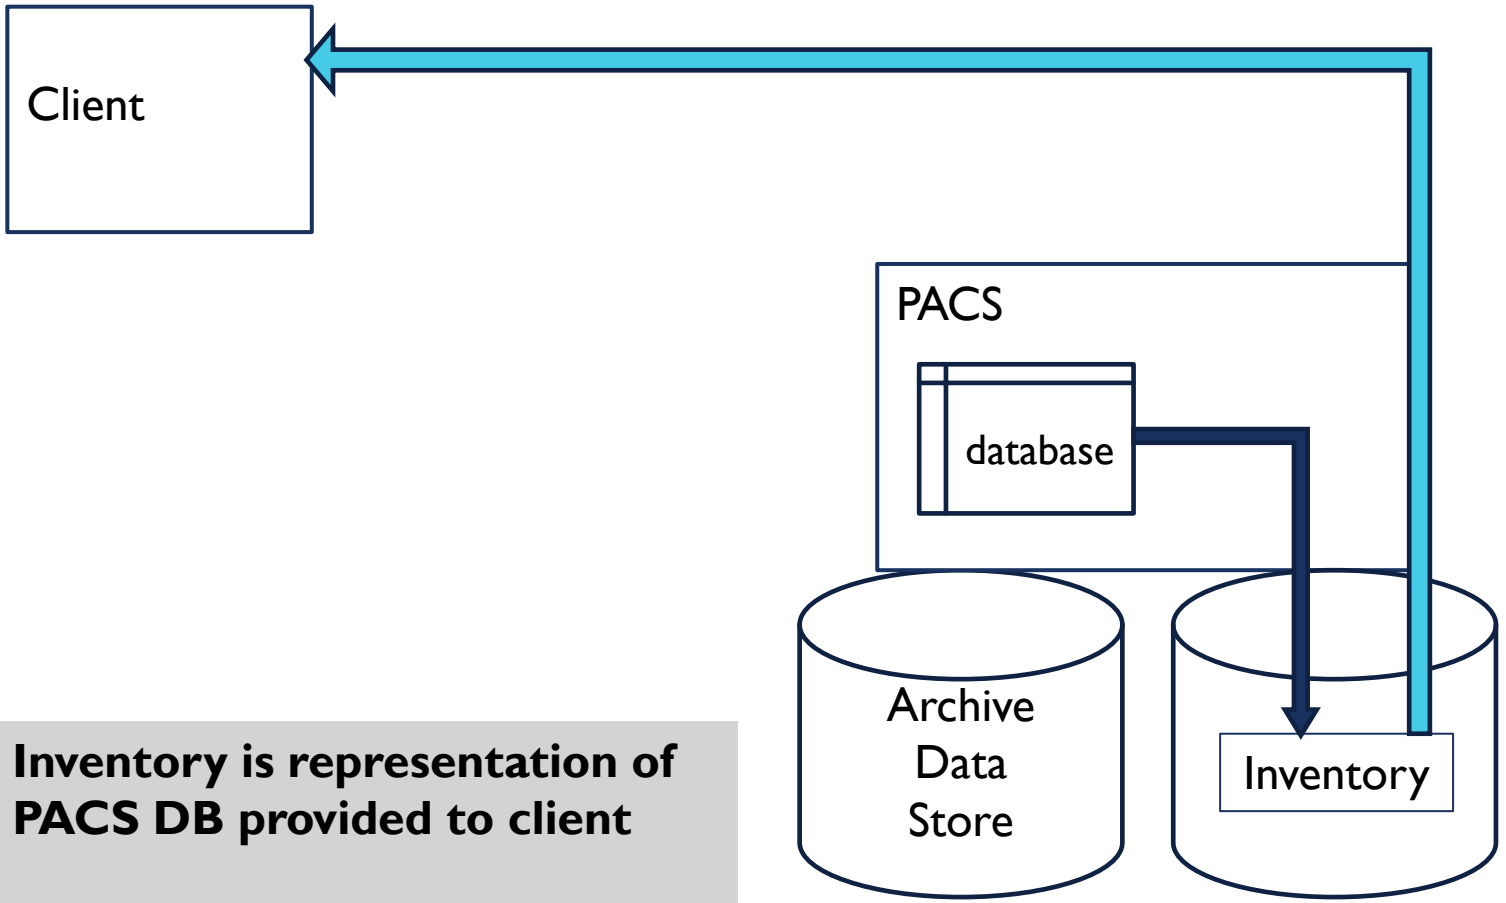

Inventory is representation of PACS DB provided to client

Equivalent of complete set of Repository Query responses

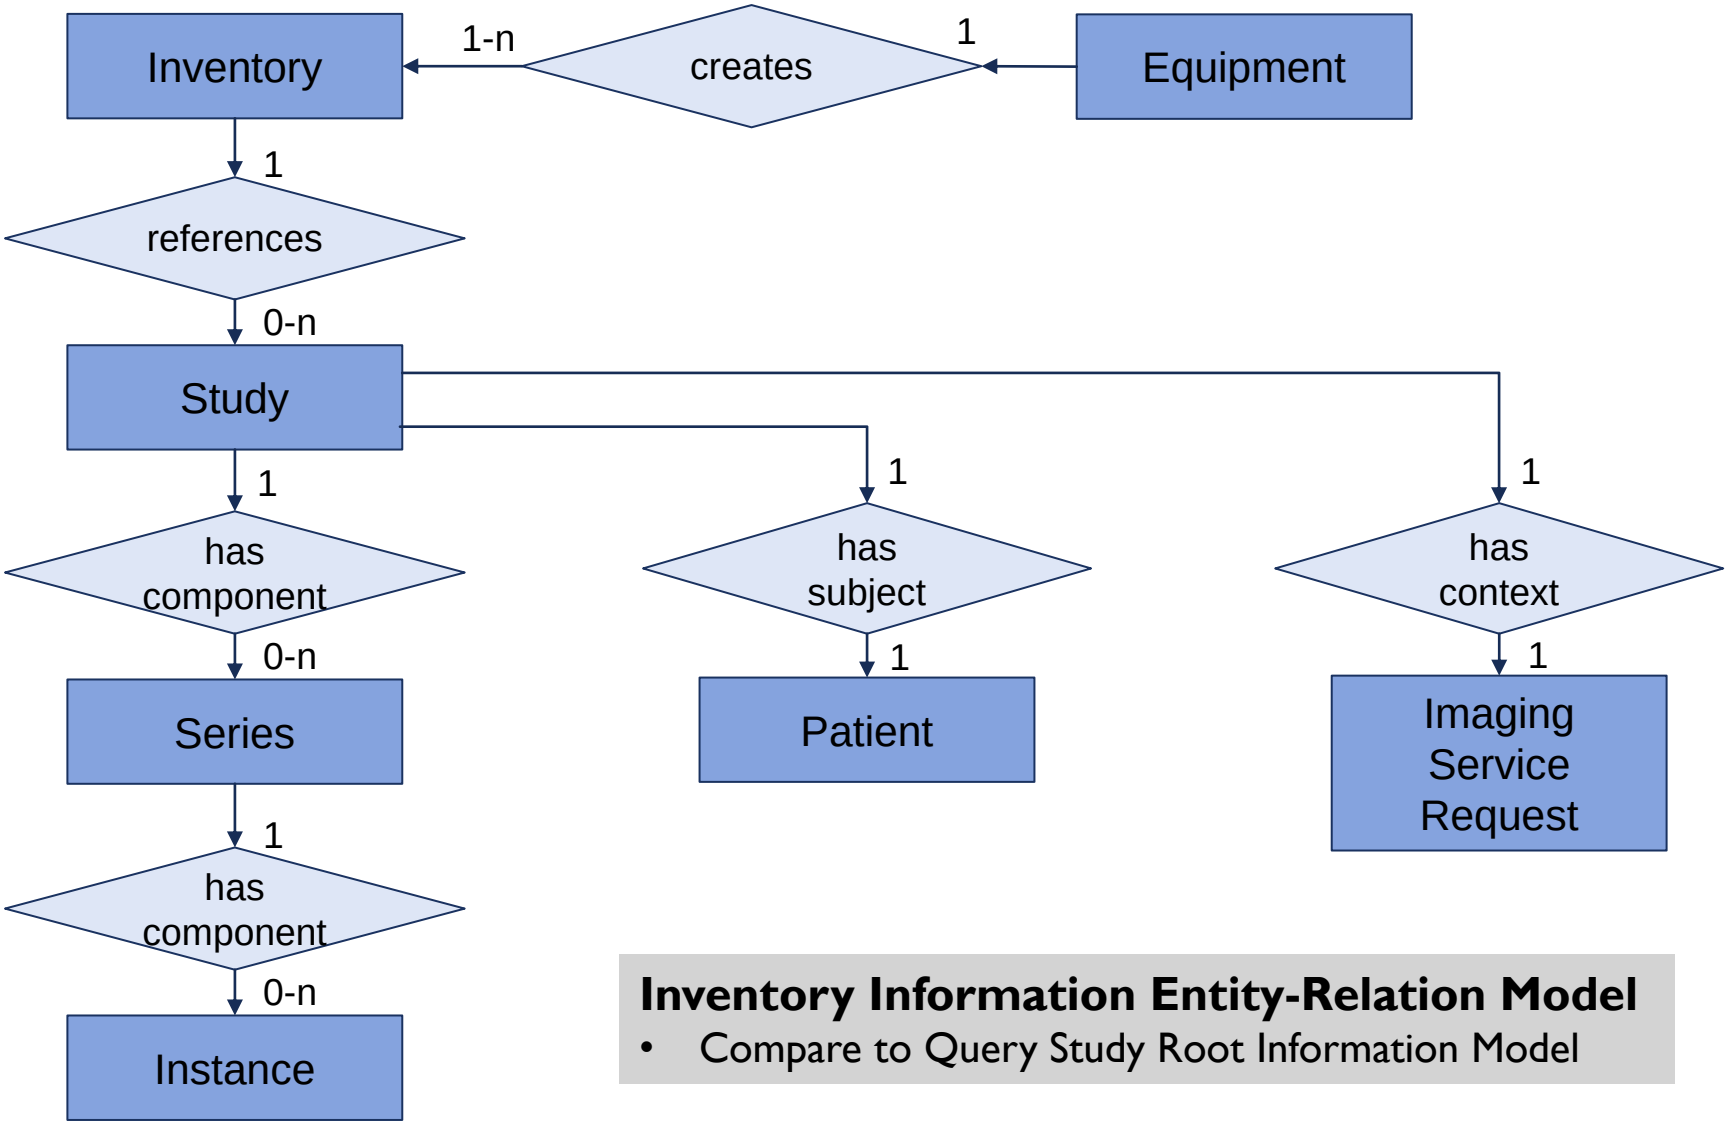

Inventory Information Entity-Relation Model

- Compare to Query Study Root Information Model
- Access mechanisms / links

Inventory IOD Structure

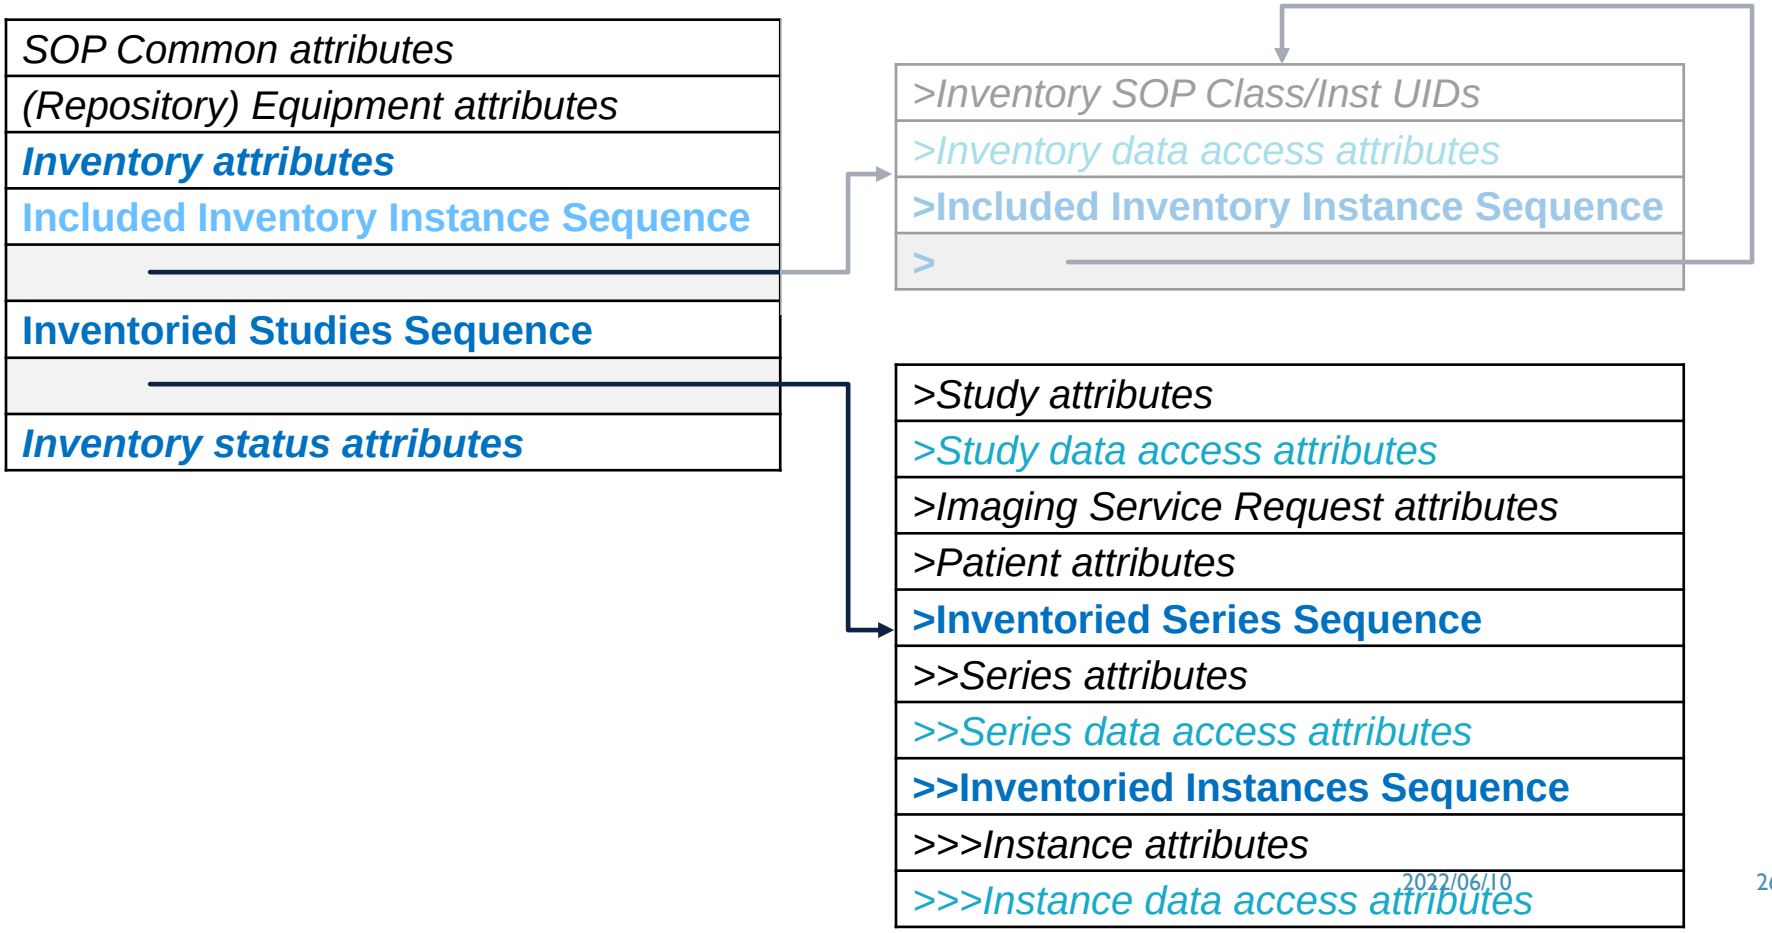

Inventoried Studies Sequence

- Hierarchical structure using Sequence attributes, following Inventory Information Entity-Relation Model (Study Root)
- Patient and Imaging Service Request attributes in Study level
- *Each level provides (optional) direct data access URIs (non-DICOM)*

Inventoried Studies Sequence
>Study attributes
>Study data access attributes
>Imaging Service Request attributes
>Patient attributes
>Inventoried Series Sequence
>>Series attributes
>>Series data access attributes
>>Inventoried Instances Sequence
>>>Instance attributes
>>>Instance data access attributes

Complex relationship of Study :: Imaging Service Request

- Potentially complex relationship between the Study and Imaging Service Requests (ISR) in the real world
  - See IHE RAD *Scheduled Workflow Integration Profile (n MSPS :: m MPPS)*
- Inventory information model follows the basic DICOM Study information model supporting a single Accession Number representing an ISR
  - See General Study Module PS3.3 Section C.7.2.1
- If a Study has multiple associated ISRs, the request attributes are encoded at the Series level in the Request Attributes Sequence
- Inventory IOD includes the Request Attributes Sequence to support this use

Level of Inventory

- Inventory may be produced at three levels - with Study records only, with Study and Series records, or Study, Series, and Instance records
- Supports different approaches to migration, or non-migration uses
- Production of Study-only Inventory expected to be significantly less resource-intensive than Instance-level Inventory
- Equivalent to Repository Query only at Study level, at Study and Series levels, or at Study, Series, and Instance levels

Inventory describes access to stored instances

- Data migration facilitated by direct filesystem access to stored SOP Instances
- Functionally equivalent (+/-) to network Retrieve (C-MOVE/WADO)
  - But Retrieve has performance issues for non-conventional access patterns and access through clinical operational pathways in PACS
- Inventory specifies available access mechanisms
  - Always default DIMSE (C-MOVE, C-GET) or DICOMweb
  - Non-DICOM file access protocol (e.g., NFS, SMB, HTTPS, …) to target Part10 files; Part10 files in ZIP, TAR[+GZIP], or BLOB containers; and Part 10 files in Study- or Series-based folders
- Inventory may record Message Authentication Code (cryptographic message digest) to ensure file integrity
  - Reader may independently compute MAC and compare to value in inventory

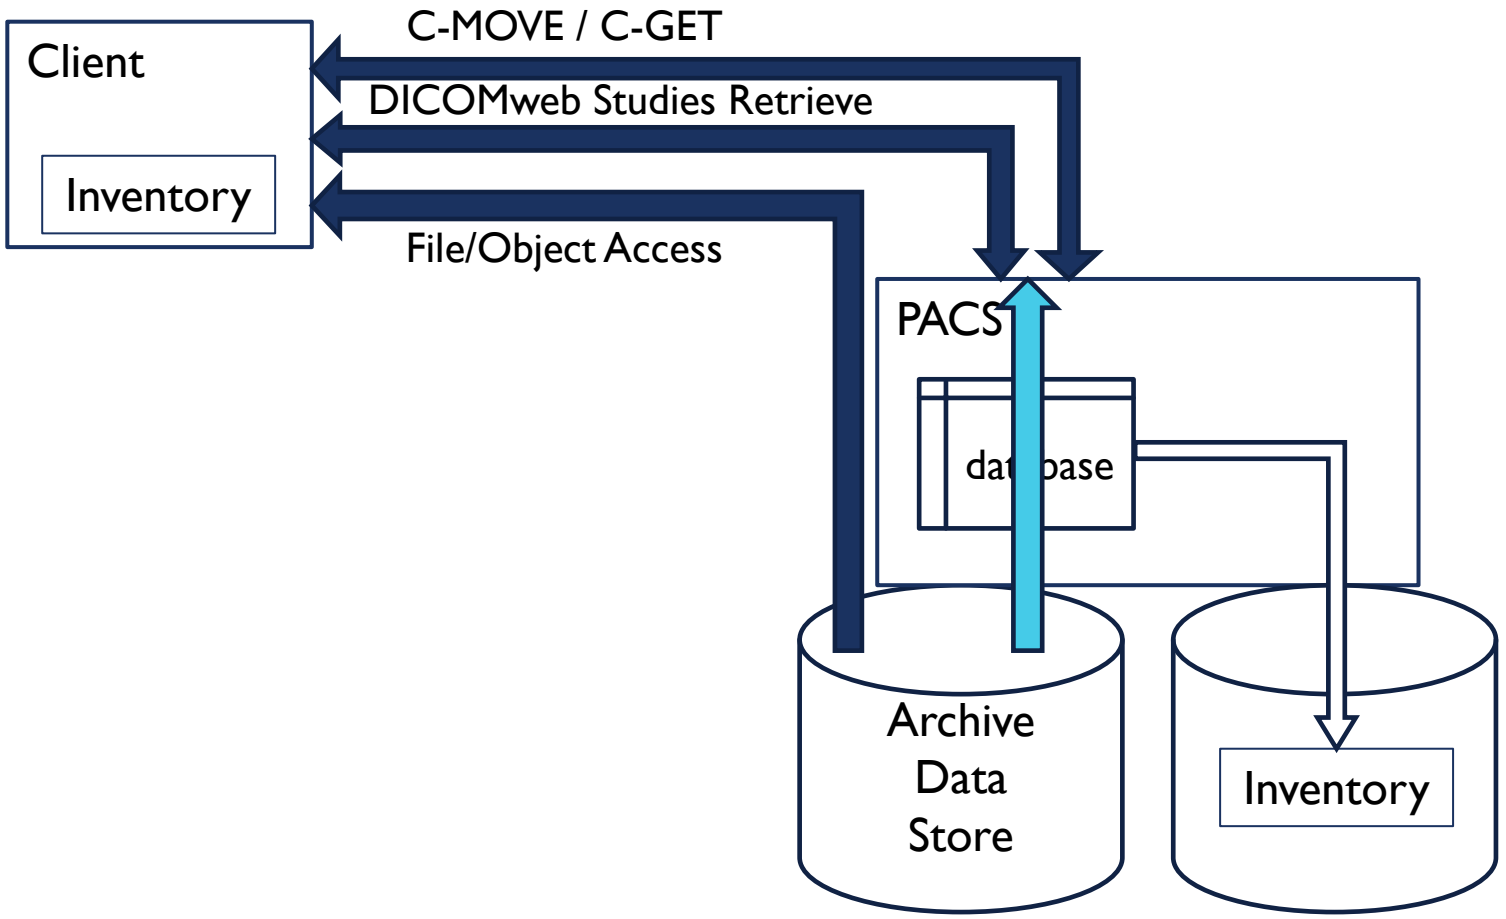

Inventory describes access to archive data: DIMSE, DICOMweb, non-DICOM protocol

Non-DICOM Protocol File Access

- **Study**
  - [Base URI +] Folder URI — all study files in folder
  - [Base URI +] File URI (Container) — all files in ZIP/TAR/BLOB
- **Series**
  - [Base URI +] Folder URI — all series files in folder
  - [Base URI +] File URI (Container) — all files in ZIP/TAR/BLOB
- **Instance**
  - [Base URI +] File URI — Pt10 file
  - [Base URI +] File URI (Container) + Filename in Container — Pt10 file in ZIP/TAR
  - [Base URI +] File URI (Container) + File Offset/Length — Pt10 file in BLOB

- Base URI is formally "Stored Instance Base URI"; File URI is formally "Stored Instance File URI"
- Base URI = <scheme>://<authority>/[<commonpath>/]
  - e.g., nfs://pacs.exampleinstitution.org/JZ0078555/
  - Set at Inventory, Study, and/or Series level (optional)
  - If Base URI used, Folder or File URI may be relative path beginning with ./

Accommodation for metadata updates

- Concern - updates to metadata attributes may not (yet) have been propagated into stored SOP Instance files retrieved through file access protocol
  - Recognition of common PACS implementation (descriptive, not prescriptive)
  - Important technical mechanism supporting efficient archive operations when update of substantial numbers of objects is invoked, or objects have long access time
- Inventory includes current metadata attributes that have not yet been propagated into stored SOP Instance files
  - Client directly accessing files must obtain current metadata from inventory
- Equivalent in Repository Query using request for Metadata Sequence or Updated Metadata Sequence

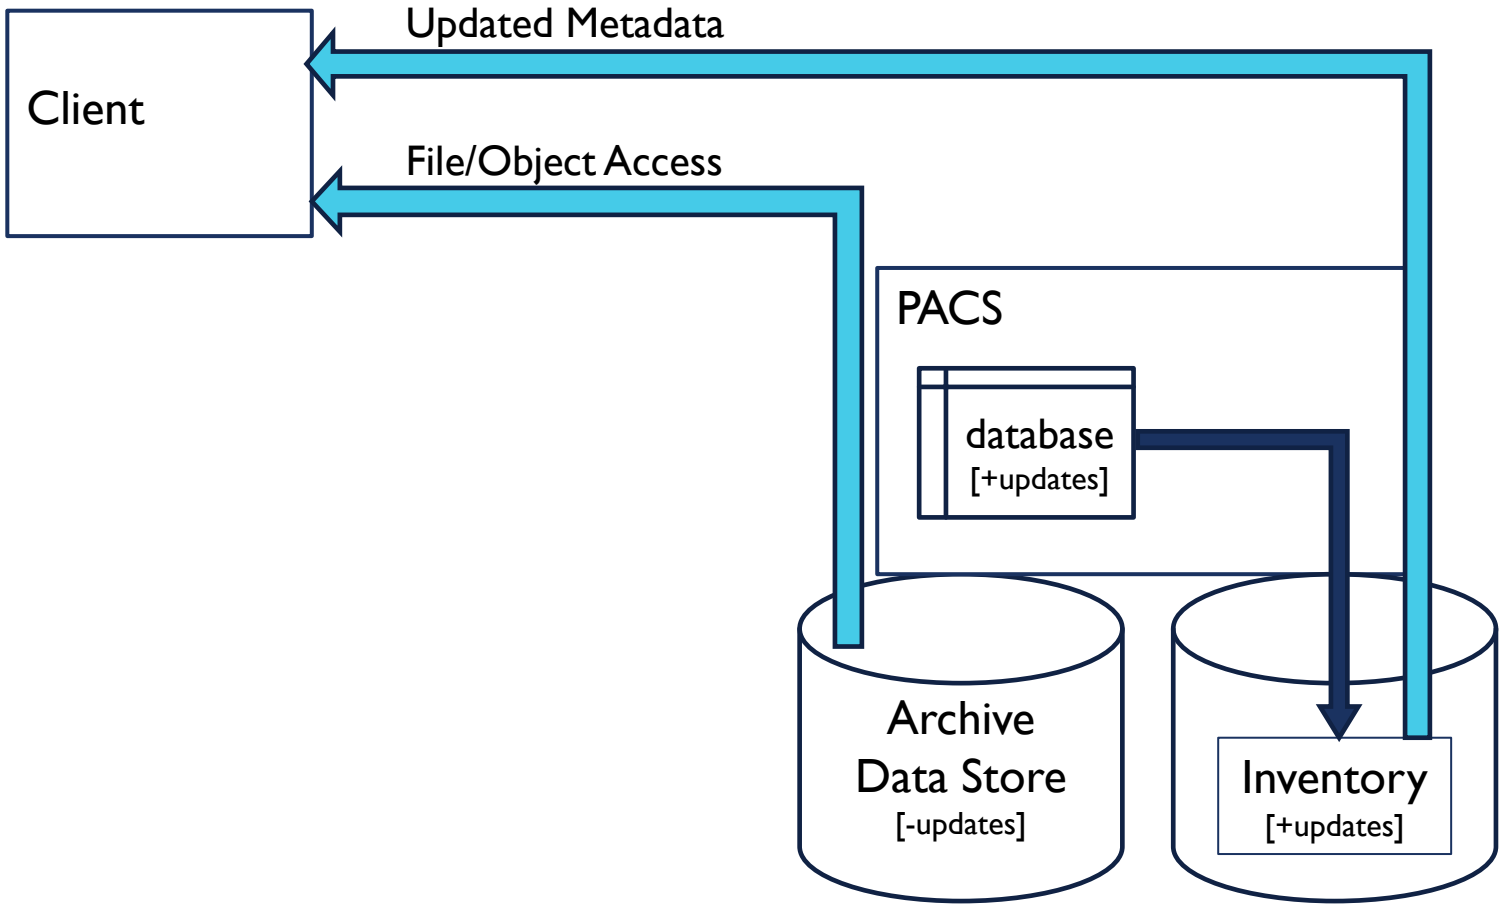

Inventory provides current metadata for directly accessed stored objects

Inventory may be for subset of complete archive

- "Scope of Inventory" in Inventory header describes study selection for inclusion in Inventory
- Similar to C-FIND key attribute matching
  - All key value matching supported – date /time range, single value matching, wildcard matching, list of UID matching
  - Two new optional matching mechanisms – multiple value matching, empty value matching
- Matching Key Attributes in Sequences by matching method
  - Avoids use of variance to VR and VM allowed in Query Key Attributes
- Scope declares use of extended matching mechanisms (equivalent to agreement by Extended Negotiation in C-FIND)
  - Relational query, combined date-time, fuzzy semantic names, multiple and empty value matching

Support for deprecated objects

- Many PACS support deprecated objects ('hard deleted', as well as 'soft deleted' or 'hidden')
  - Including images rejected for quality or patient safety reasons through PACS UI or IHE Imaging Object Change Management (KO objects with specific title concept codes)
  - But SOP Instances may be required to be archived for regulatory reasons (e.g., patient X-ray exposure)
- Inventory includes Studies/Series/SOP Instances identified as Removed from Operational Use, with specific attribute/reason code
  - 'Operational use' includes diagnostic, clinical, therapeutic, and administrative uses
  - Such SOP Instances might not appear in normal C-FIND responses, but are included in inventory
- Inventory includes IOCM KO objects if included in archive
  - PACS simply reports what is has in archive
  - Up to the receiving app (migration client) to determine what to do about them – different organizations may have different policies

Multiple inventory objects

- Enterprise inventories may be enormous
- For implementation specific reasons content of an inventory may need to be split into more than one SOP Instance
  - Practical limits on the maximum size of an individual SOP Instance - tractable object sizes
  - Archive inventoried by multiple parallel processes
- Logical inventory structured as tree of Inventory SOP Instances
  - Inventoried studies in subsidiary nodes included by reference (Incorporated Inventory Instance Sequence)
  - Root instance provides scope, final production status, and links to all instances
  - *More on this later*

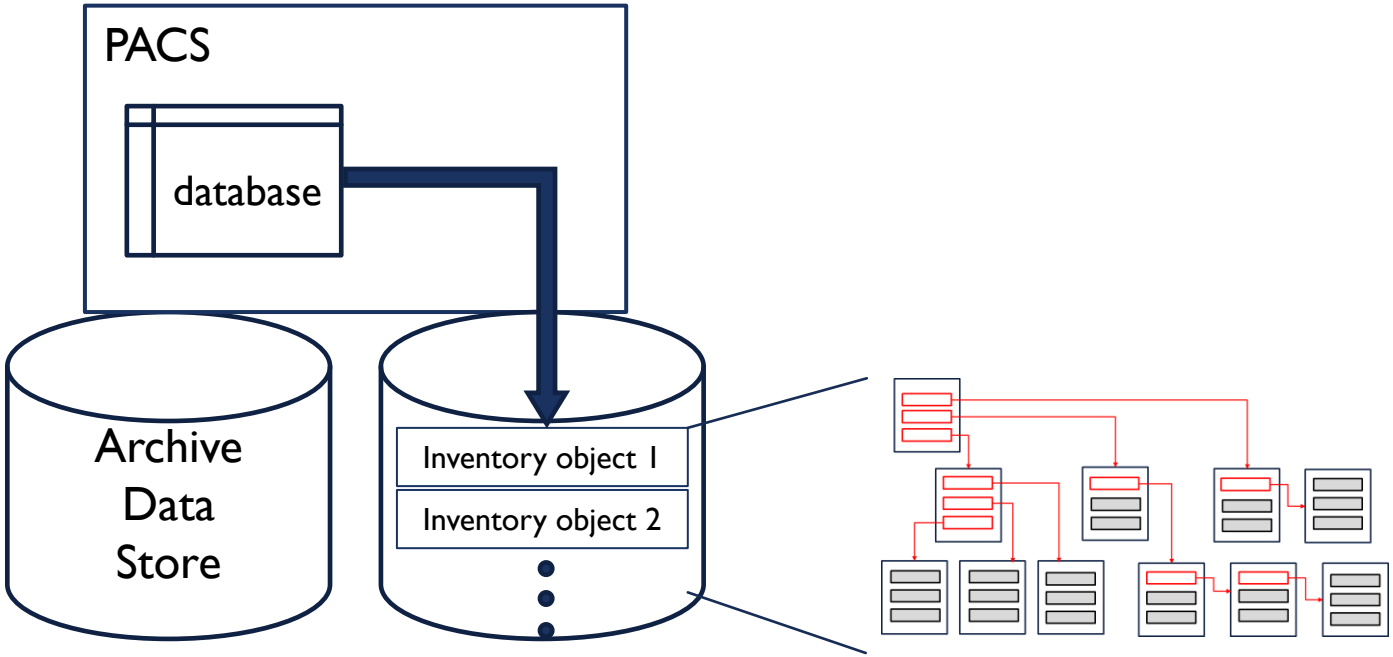

Inventory may be divided into multiple instances organized in tree

Associated Services for Inventory Instances

Inventory as DICOM Non-Patient Object for Query / Retrieve / Store

- Inventory defined within category of DICOM Non-Patient Objects
  - Color Palette, Hanging Protocol, Defined Protocol, etc.
- Simple addition to Part 4 Annex GG defines Storage Service (C-STORE transfer)
- New Part 4 Annex for Inventory Query/Retrieve (C-FIND, C-MOVE, C-GET)
  - Inventory Q/R Info Model
  - Based on other NPO Q/R Services
- Simple addition to Part 18 Section 12 defines equivalent DICOMweb operations
  - Standard transcoding to XML and JSON for DICOMweb retrieve
- Optionally stored in a Part 10 conformant file - transferred via Media Service, or accessed via non-DICOM file access protocol
  - Similar to accessing Part 10 files from archive

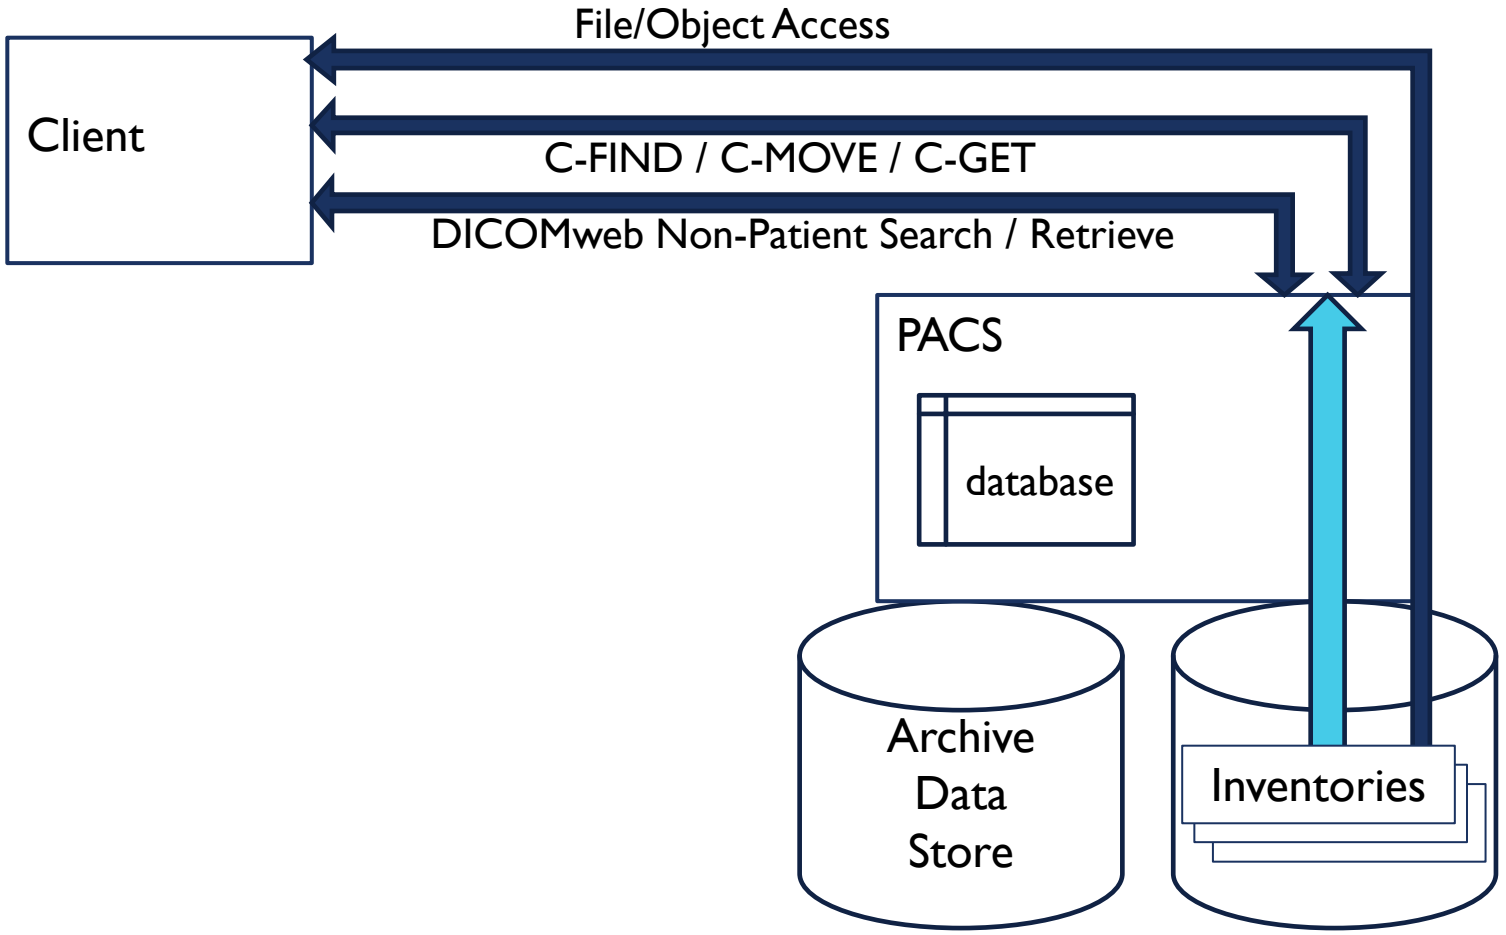

Inventory accessed by DIMSE, DICOMweb, non-DICOM protocol

Network service to initiate inventory creation

- App external to PACS/VNA may request production of inventory, especially for dynamically determined subsets of archive
  - Migration, research, or QA clients
- New service modeled on Storage Commitment (already implemented in PACS)
  - Inventory Creation IOD for DIMSE-N services conveys requested Scope of Inventory and Content Level
  - Initiate, Cancel, Pause, Resume: N-ACTION from SCU
  - Progress indication, Completion: N-EVENT REPORT push from SCP
  - Asynchronous mechanism – supports extended period to create inventory
- Periodic progress indication with reporting interval requested by SCU
- 'Initiate' response includes errors and warnings to indicate unsupported parameters

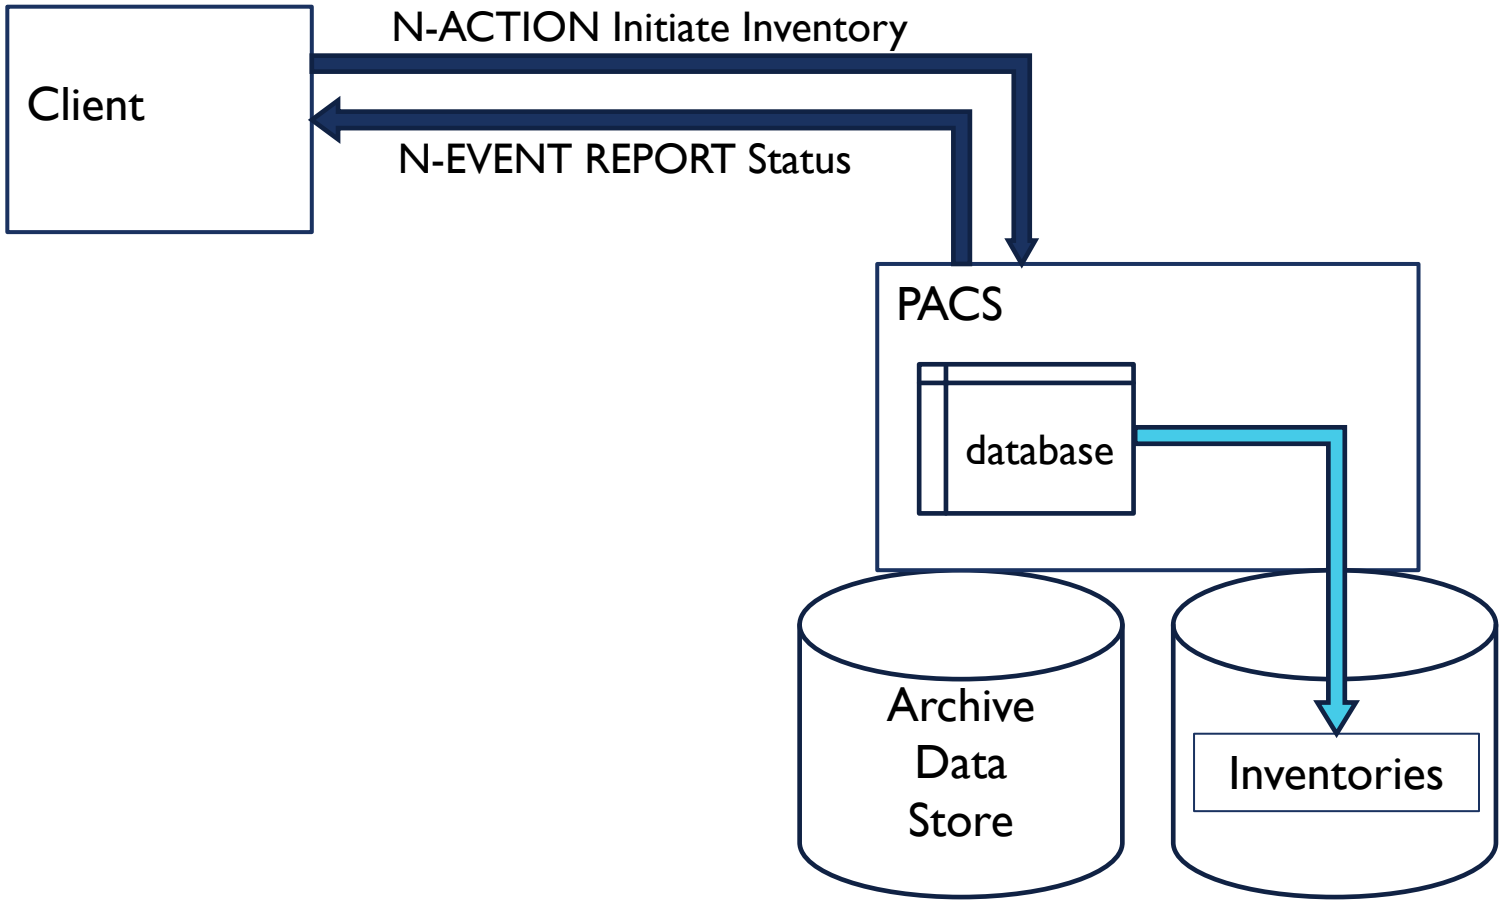

Client initiates inventory production through DIMSE service

Inventory Production States

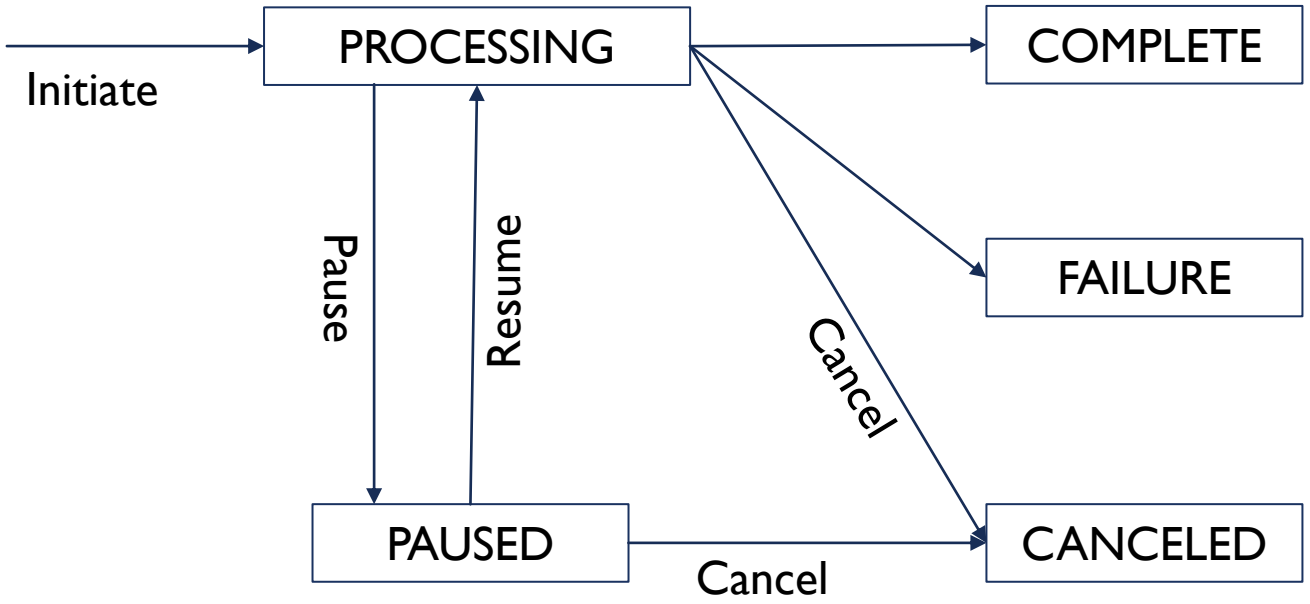

Some state changes may be due to N-ACTION requests (labeled arrows)

Each defined service is separable from others

- Inventory Creation service not required - repository may implement creation of inventory from its administrative UI
- Access to Inventory instances at discretion of repository – could use DIMSE C-STORE (with or without C-MOVE or C-GET), DICOMweb, or non-DICOM file access protocol
- ID and location of Inventory instances may be supported by C-FIND or DICOMweb, or may be done out of band by non-DICOM means (e.g., email notification of filename)
- But there is no DICOM Conformance to just an IOD – at least one instance exchange service must be supported to assert DICOM Conformance

DIMSE

Media

Web

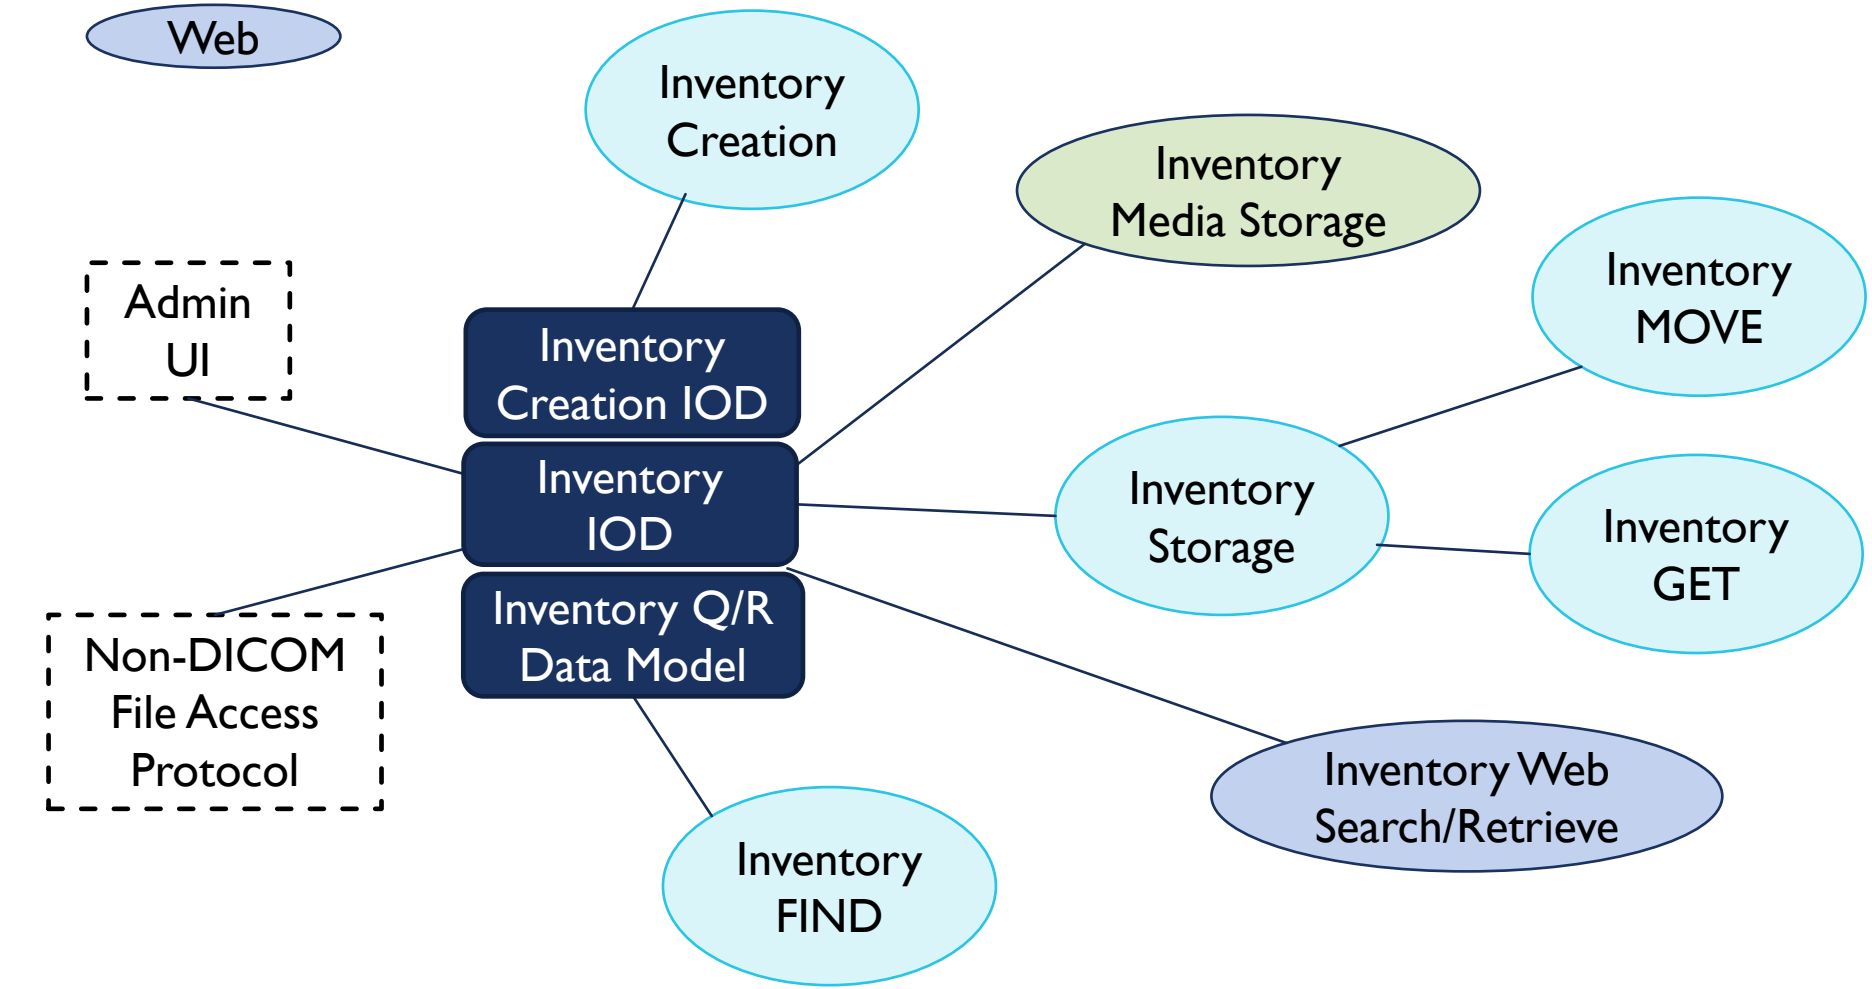

Inventory SOP Instance-related IODs and associated services

Summary of Sup 223

- Repository Query SOP Class with partial results and continuation
- Inventory IOD
- Specifies (if available) direct filesystem access to stored instances
- Inventory has current metadata (possibly not propagated to stored instances)
- Inventory may be produced for subsets of full archive
- Supports instances not for operational use (including IOCM controls)
- Inventory in class of DICOM Non-Patient Objects with typical network services (C-STORE, C-FIND, C-MOVE, C-GET, and DICOMweb equivalents)
- Service to initiate inventory creation

Level and Scope of Inventory Examples

Complete inventory

- Zero Length Scope of Inventory Sequence indicates "universal match" – all studies, with just study records, or all instances

Inventory Level	STUDY
Scope of Inventory Sequence	

Inventory Level	INSTANCE
Scope of Inventory Sequence	

Inventory for moving to deep archive

- All studies with Study Date prior to and including 2010

Inventory Level	INSTANCE
Scope of Inventory Sequence	
>Range Matching Sequence	
>>Item 1 (beginning of range)	
>>Study Date	
>>Item 2 (end of range)	
>>Study Date	2010

Incremental inventory

- All studies with Study Update DateTime on or after July 1, 2022

Inventory Level	INSTANCE
Scope of Inventory Sequence	
>Range Matching Sequence	
>>Item 1 (beginning of range)	
>>Study Update DateTime	20220701
>>Item 2 (end of range)	
>>Study Update DateTime	

Quality Control inventory

- Studies missing Patient ID – Scope of Inventory needs to declare use of empty value matching

Inventory Level	STUDY
Scope of Inventory Sequence	
>Extended Matching Mechanisms	EMPTY_VALUE
>Empty Value Matching Sequence	
>>Patient ID	

Patient Safety Audit inventory

- CT Studies with a Structured Report; inventory may be further searched by client, e.g., for X-Ray Radiation Dose Report objects in the SR Series that will be retrieved

Inventory Level	INSTANCE
Scope of Inventory Sequence	
>Extended Matching Mechanisms	MULTIPLE_VALUE
>General Matching Sequence	
>>Modalities in Study	CT\SR

CT Radiation Dose inventory with Relational Search

- All CT Radiation Dose Report objects – repository must support Relational Search and matching on Template Identifier

Inventory Level	INSTANCE
Scope of Inventory Sequence	
>Extended Matching Mechanisms	RELATIONAL
>Query/Retrieve Level	INSTANCE
>General Matching Sequence	
>>Content Template Sequence	
>>>Template Identifier	10011
>>>Mapping Resource	DCMR

Research inventory

- MR-PET head Studies in 2019 – assumes Study Description includes the string "HEAD"

Inventory Level	INSTANCE
Scope of Inventory Sequence	
>Extended Matching Mechanisms	MULTIPLE_VALUE
>Range Matching Sequence	
>>Item 1 (beginning of range)	
>>Study Date	2019
>>Item 2 (end of range)	
>>Study Date	2019
>General Matching Sequence	
>>Study Description	*HEAD*
>>Modalities in Study	MR\PT

Tree of Inventory Instances Included by Reference

Inventory IOD Structure

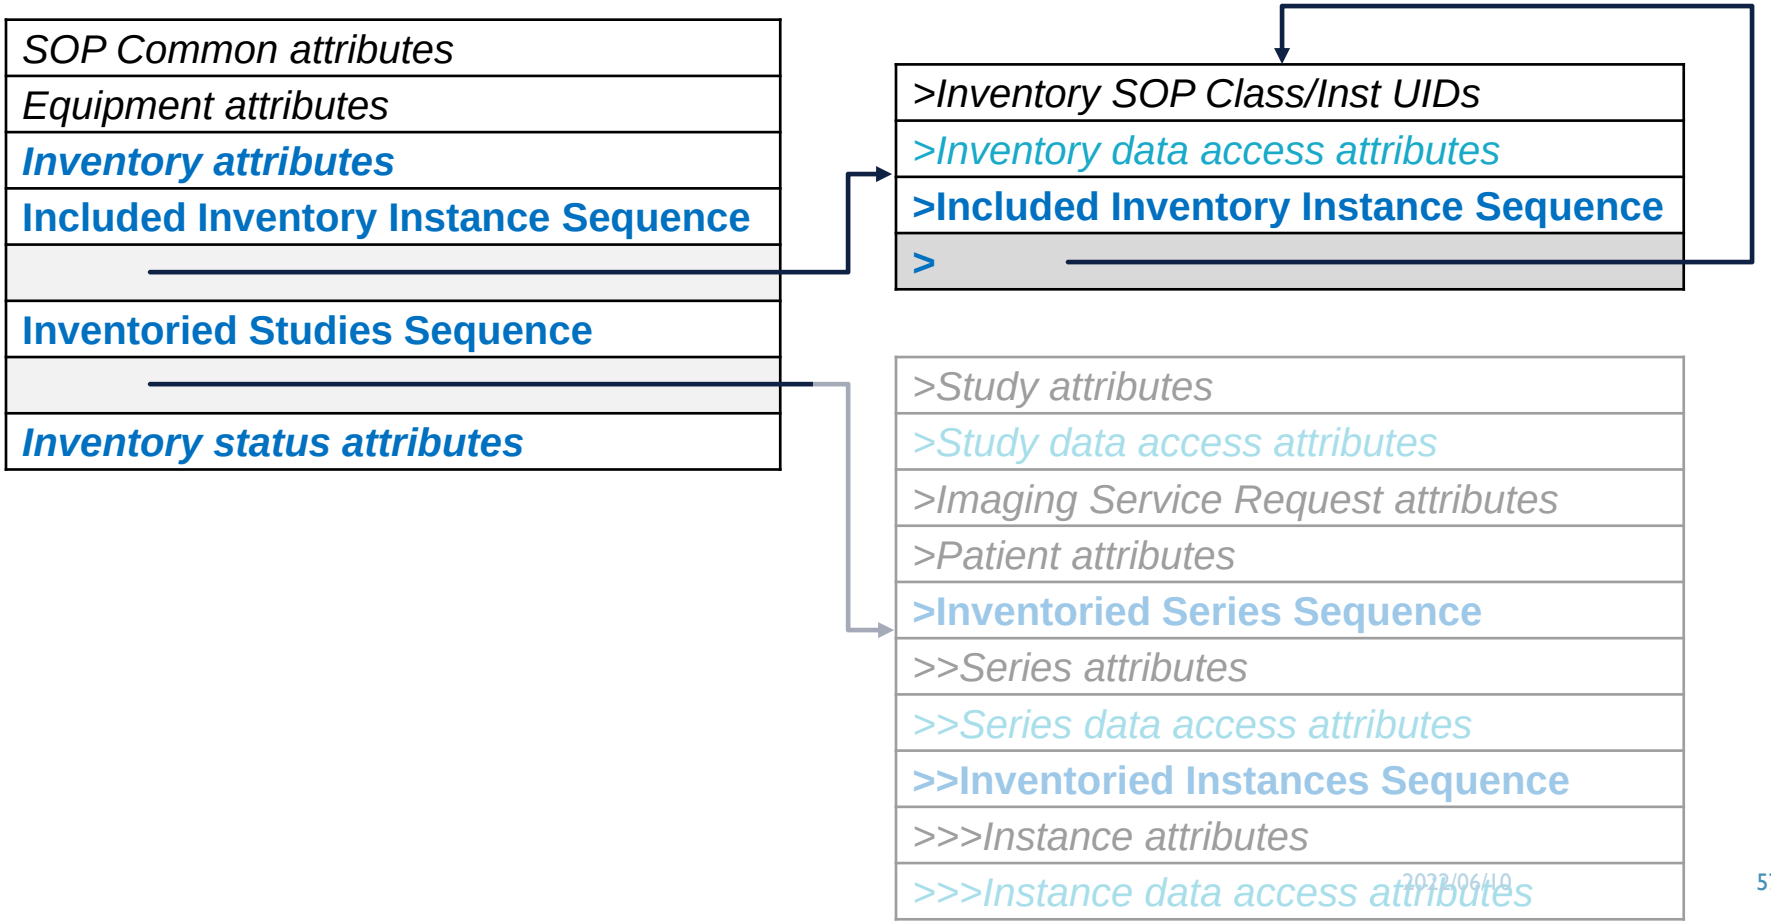

Inventory Content by Reference

- Inventory object includes:
  - Content Date/Time (Date/Time of inventory initiation)
  - Equipment producing inventory
  - Level and Scope of Inventory
  - Links to other inventory objects included by reference, and/or
  - Studies Inventory
  - Completion status (with respect to Scope of Inventory at Content Date/Time on Equipment)
- Linked objects form a tree from a root object
  - Root specifies Content Date/Time, Scope of Inventory, and Completion status *for tree*
  - Those attributes in non-root objects ignored (but are valid for their subtree)
- Pre-existing inventory objects can be linked – may appear under multiple roots

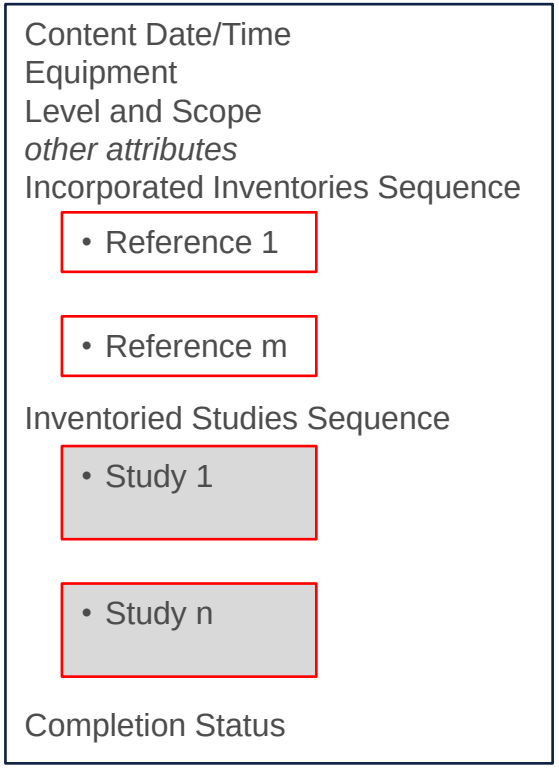

Tree of Instances Incorporated by Reference

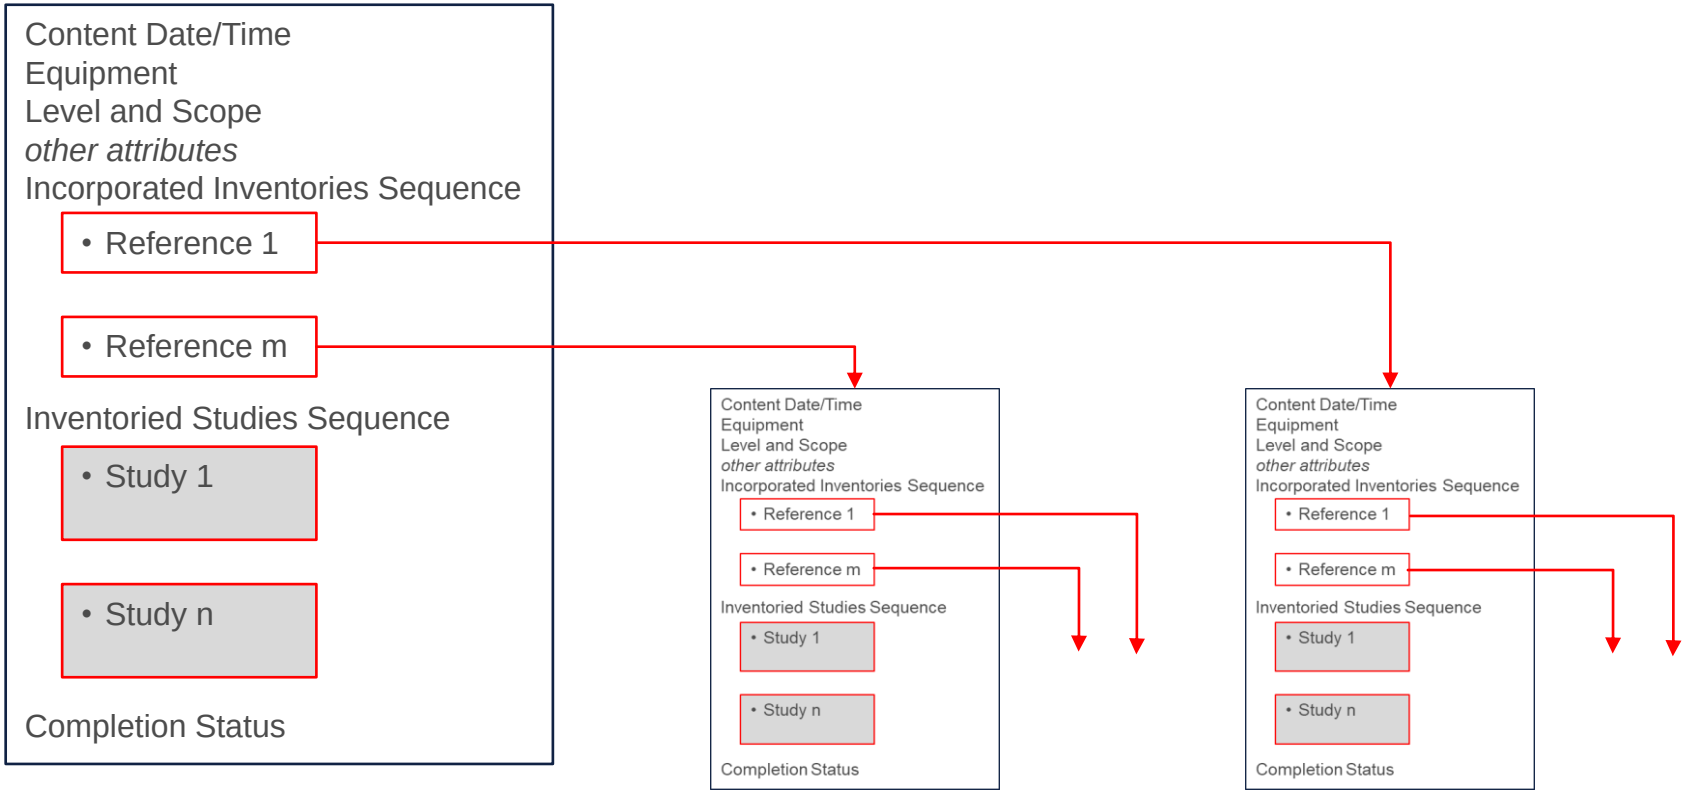

Example Approach: Serial Production

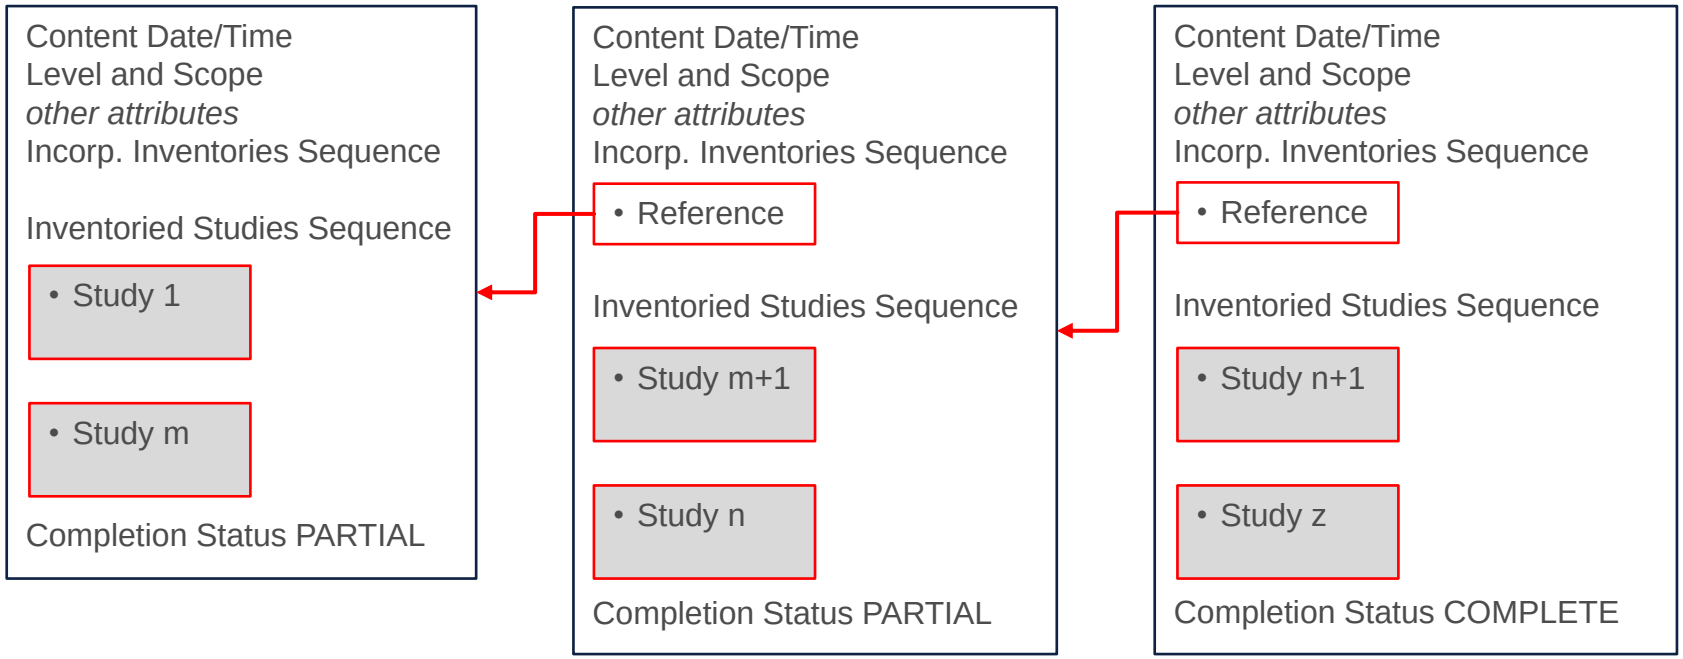

Objects filled with inventoried studies until max size, next object links to prior, last becomes root

Example Approach:
Baseline and incremental update

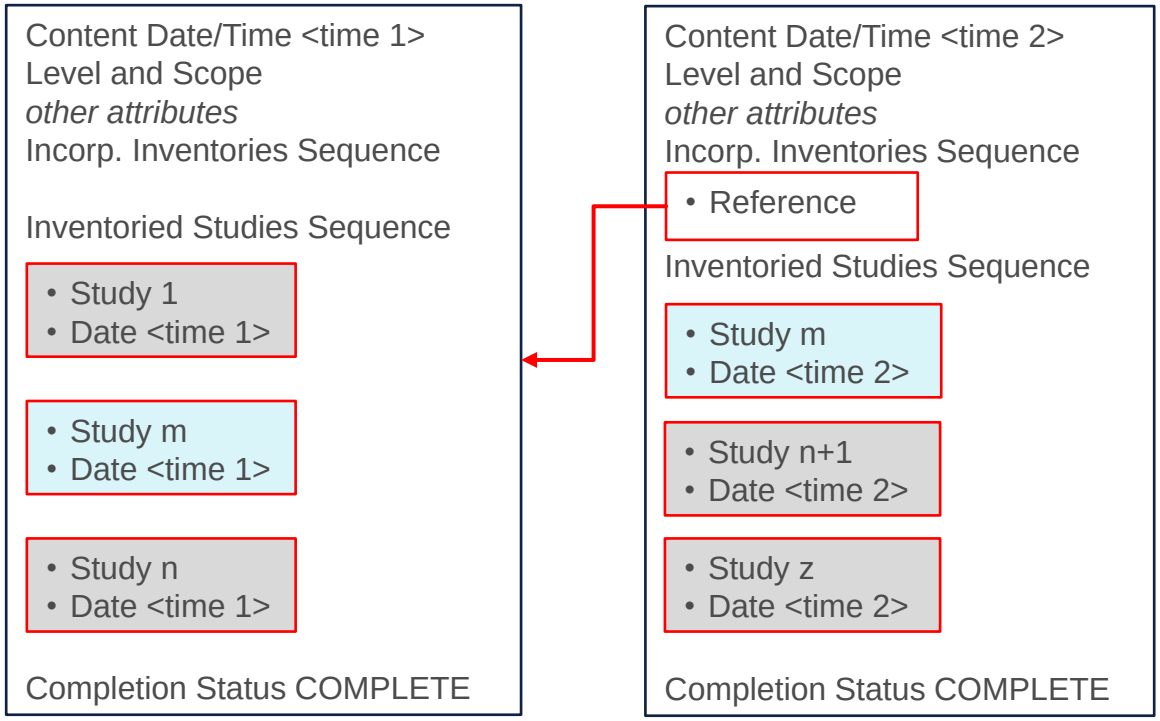

Baseline included by reference in incremental update inventory, changed study simply has new (complete) record

Example Approach:
Parallel inventory (multiple subsystems)

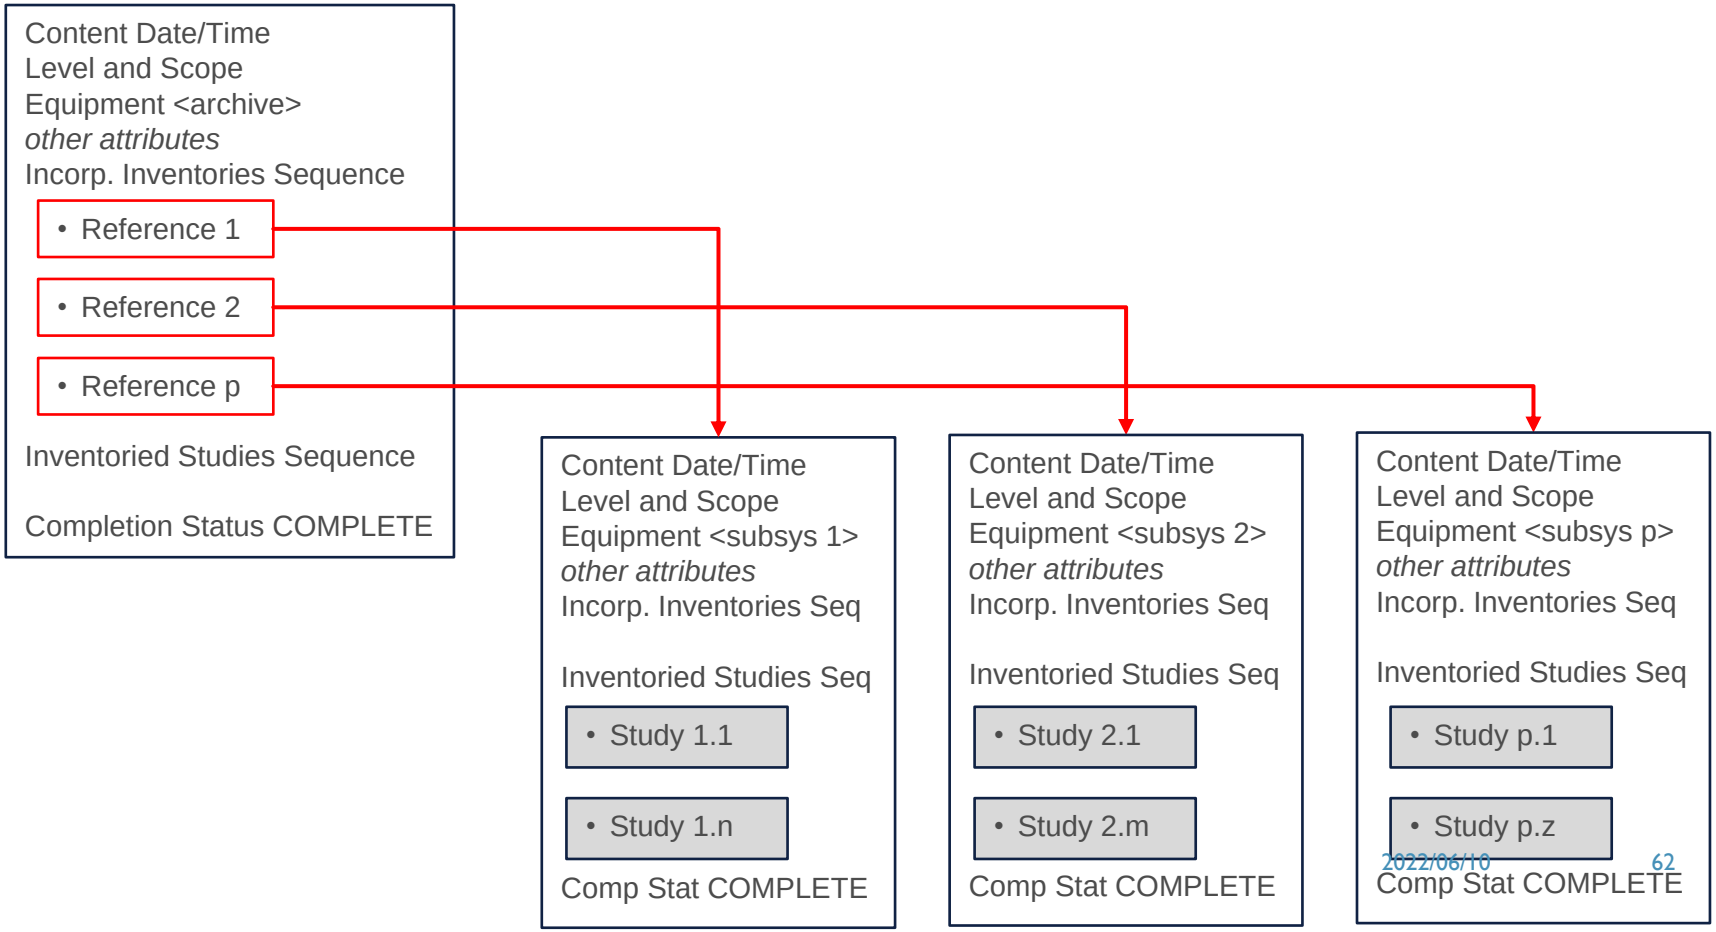

Arbitrary tree

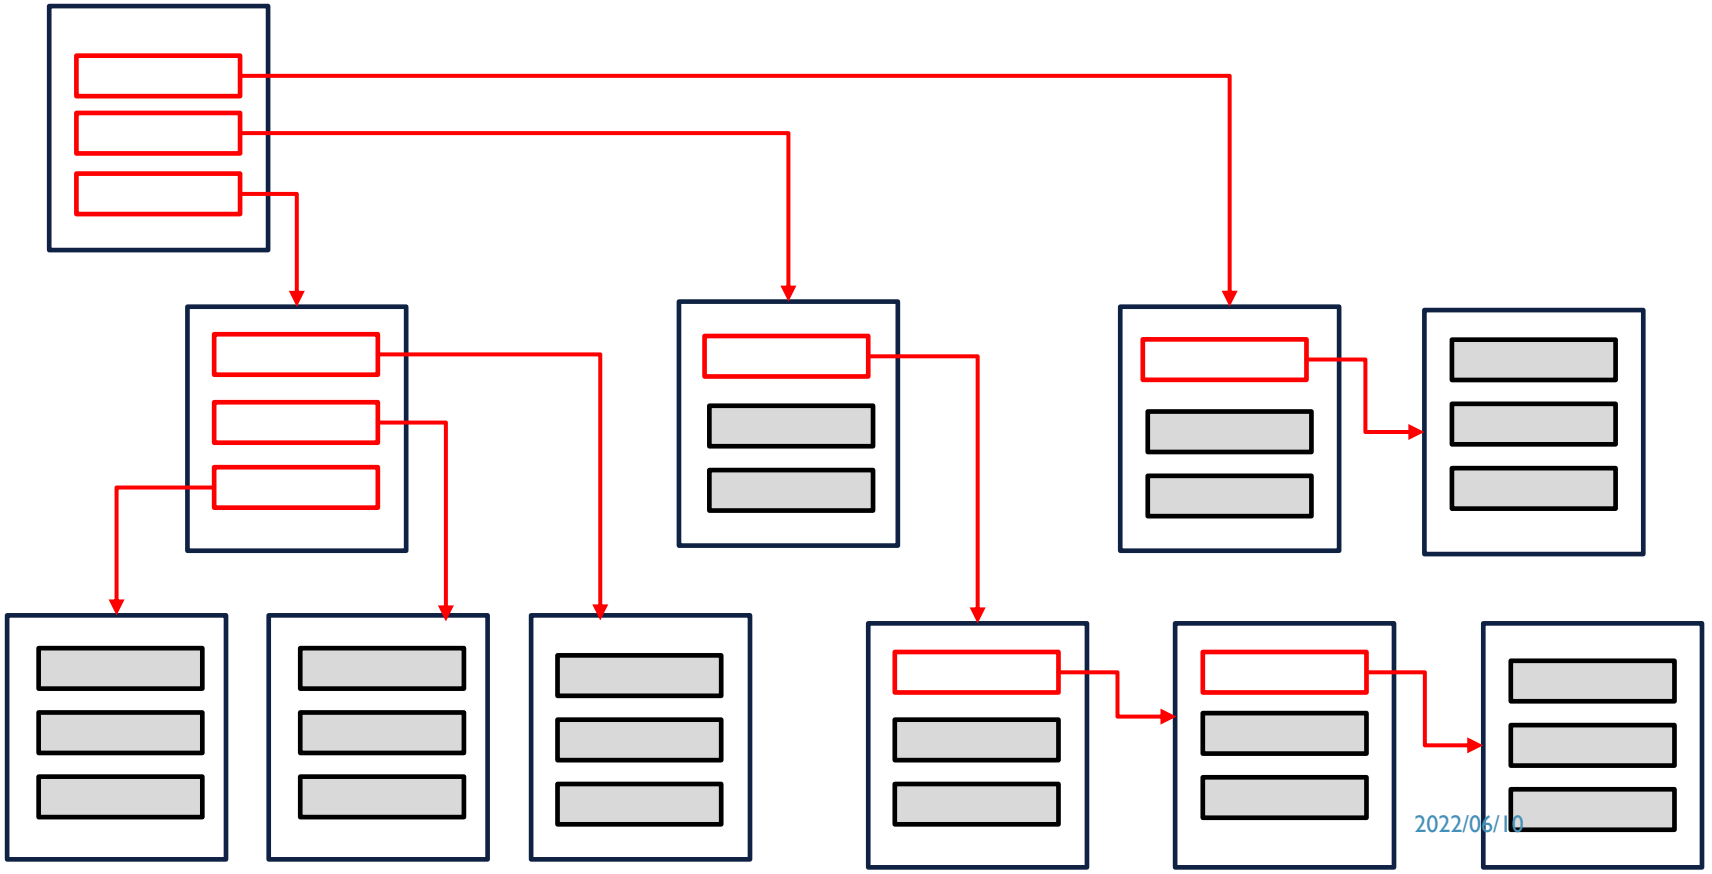

Empty inventory

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

Affirmative declaration that at Content Date/Time, identified equipment had no studies matching Scope

Multiple records for study – reconciliation

- Study may appear multiple times in tree of inventory SOP instances
  - Parallel production of inventories with replicated storage, or with different series on different subsystems
  - Sequential production (baseline + increment) with change to study content
- Each Study record complete as known by the creator of that Inventory instance
- Consumer required to reconcile
  - No consolidation/reconciliation requirement on producer – simplify production
  - Consumer (migration client) typically needs to do reconciliation anyway due to merge with EHR data, errors in DICOM data, or idiosyncratic/poor PACS implementation
- Inventory records identify equipment and time stamp to support reconciliation

Inventory Links

- Challenge for inventory user to locate and retrieve all referenced inventories in the tree
  - No requirement for producer to support Inventory Query/Retrieve as SCP
- Inventory includes access info with each link to referenced Inventory instance
  - DIMSE Inventory Query/Retrieve Service, DICOMweb Non-Patient Instance Service, and/or non-DICOM file access protocol
- Incorporated Inventory Instance Sequence defined recursively
  - Tree structure specifies all subsidiary Inventory instances and their links
- Incorporated Inventory Instance Sequence from referenced instance copied into referencing instance

Links with hierarchical structure of subtree links

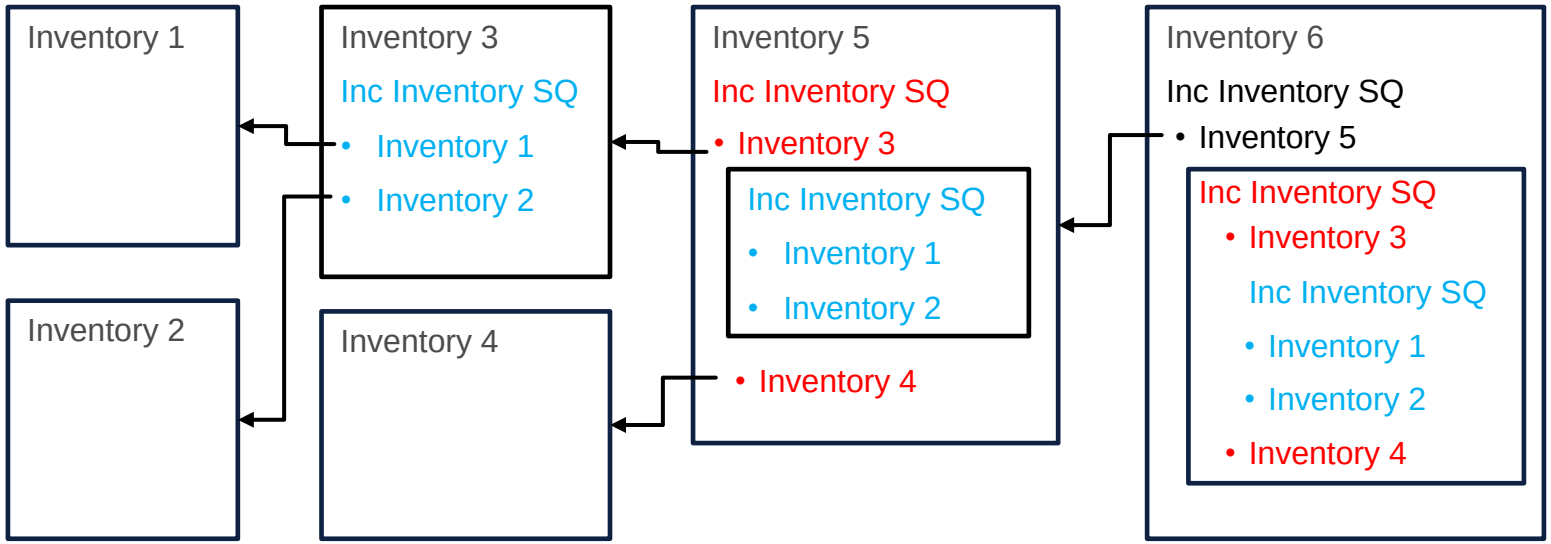

Links structure from referenced object added *in toto* as subsidiary attribute

Inventory links in related services

- Related Services need to reference an inventory
  - Inventory Creation complete N-EVENT REPORT
  - Inventory Query Model C-FIND
- With tree of Inventory Instances, do we give entire tree, or just the root?
- Mechanics of specifying recursive structure in DIMSE-N services is problematic
- Approach is to identify only root inventory object in N-EVENT REPORT and C-FIND
  - With full tree and access links in Included Inventory Instances SQ, app only needs to open one object (root)

What's Not in Sup 223 Normative Specs (but in Informative Annex)

Inventory instance produced directly, or through intermediary using Repository Query

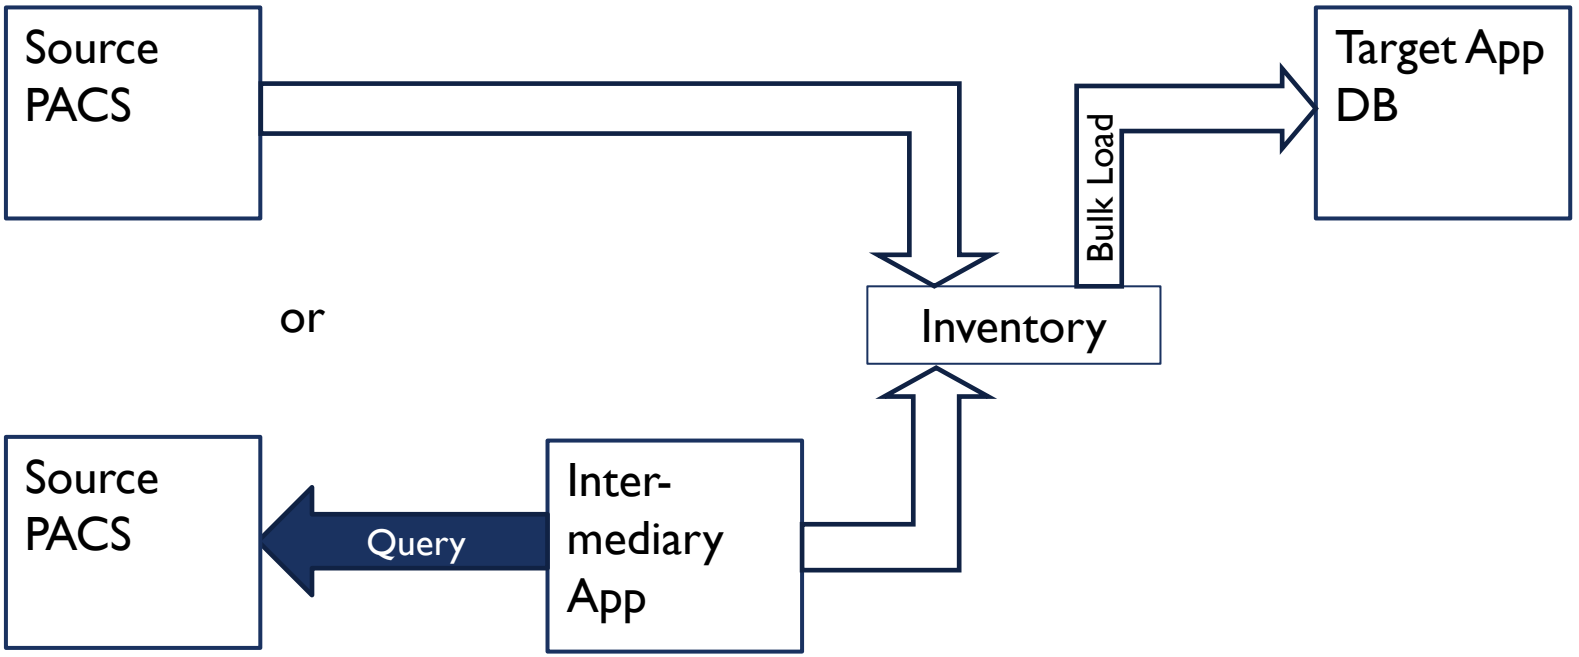

Considerations for Inventory produced through intermediary

- Informing intermediary of desired Scope of Inventory
  - Local configuration, or Inventory Creation SOP Class
- Requesting all Type 1 and 2 Attributes
- Managing Key Attributes not supported for matching by SCP
- Differences in encoding of Key Attributes
- Handling Updated Metadata
- Production of tree of Inventory SOP Instances
- Handling partial results/continuation, errors

Security issues (1)

- Services associated with inventories have no specifications regarding access control or transport security, but those are necessary for any real-world deployment, and are implementation- and site-specific
  - Same as all DICOM network services
  - Site must also address access control and transport security for non-DICOM protocol
  - Storing access credentials in Inventory is NOT appropriate
  - Implementation should consider interaction of user identity associated with inventory production application and access controls for specially protected information
- Files (Inventory and stored SOP Instances) have inherent security vulnerabilities
  - Part 10 header
  - ZIP and TAR executable content

Security issues (2)

- No specifications regarding encryption of data at rest (e.g., SOP Instances in the repository, or Inventory objects)
  - DICOM does specify methods for secure (encrypted) files, but issues such as key management and distribution again are implementation- and site-specific
  - Storing keys in Inventory is NOT appropriate
- Applications using Inventories need robust validation and controlled introduction to operational environment
- Resource over-use by inventory application may result in denial of service to critical functions

Inventory object lifecycle management

- No specifications for explicit management of lifecycle of Inventory objects
  - At discretion of system managing inventory objects
  - No method specified for client to request deletion, except when canceling inventory production
  - Creator may specify expiration datetime in Inventory Creation Complete notification, and in Inventory C-FIND response

Incremental access to Inventory SOP Instance content

- Inventory objects may be very large; DIMSE C-STORE transfers whole object only
- Nothing comparable to frame-based retrieve specified in DIMSE services
- DICOMweb has inherent byte range retrieval (optional support)
- Direct network filesystem object retrievals may allow interactive retrieval (protocol typically integrated into operating system filesystem I/O)
- Low information entropy means compression is effective; compressed whole object transfer may be feasible
  - Deflated Transfer Syntax
  - ZIP or GZIP of Inventory in Part 10 file

Use of multiple AE Titles

- A PACS may support multiple AE Titles to provide different "views" of the repository
  - Different primary Patient ID based on caller's organization
  - Specific repository subsystem
  - Studies with research consent