DICOM Secretariat/MITA
1300 North 17th Street, Suite 900
Arlington, VA 22209, USA
http://www.dicomstandard.org/
dicom@dicomstandard.org

MINUTES

MEETING NAME	09-WG33: WG-33 Data Archive and Management	
MEETING PLACE/DIAL IN		
DATE & TIME	Wednesday, February 3, 2021 | 11:00 am – 12:30 pm US ET	
PRESIDING OFFICERS	Matthew Bishop, UnityPoint Health	
	Keith Eklund, Healthcare Tech Solutions	
VOTING MEMBERS PRESENT	Ambra Health	Ostrow, Daniel
	American College of Radiology	Maldonado, Josh
	Argentix Informatix	Silver, Elliot
	Change Healthcare	Ho, Kinson
	DesAcc EMEA	King, Graham
	Healthcare Tech Solutions	Eklund, Keith
	Laitek, Inc.	Behlen, Fred
	Laitek, Inc.	Brown, Barry
	Laitek, Inc.	Costea-Barlutiu, Razvan
	Laitek, Inc.	Solomon, Harry
	Mayo Clinic Rochester	Persons, Kenneth
	Society for Imaging Informatics In Medicine	Bishop, Matthew
	Society for Imaging Informatics In Medicine	Carey, Cheryl
OTHERS	Experies Consulting	McCloskey, Thomas
	London Health Science Center	Aizawa, Luiz

VOTING MEMBERS ABSENT	AAPM	Bevins, Nicholas
	Canon Medical Research USA	O'Donnell, Kevin
	Canon/Vital Images	Dawson, Tim
	Canon/Vital Images	Whitby, Jonathan
	European Society of Radiology	Mildenberger, Peter
	GE Healthcare	Nichols, Steven
	GE Healthcare	Numan, Jouke
	Grafimedia	Georgiadis, Pantelis
	PixelMed Publishing	Clunie, David
	Society for Imaging Informatics In Medicine	Henson, Kyle
	Varian Medical Systems, Inc.	Schwere, Thomas

DICOM SECRETARIAT Anna Zawacki, SIIM

1 CALL TO ORDER AND REVIEW OF ANTI-TRUST RULES AND DICOM PATENT POLICY (Co-Chairs, Secretariat)

The meeting was called to order. Guidelines for Conducting NEMA Meetings were read and attendance was recorded.

2 REVIEW AND APPROVE AGENDA (Co-Chairs)

The agenda was reviewed and approved.

3 REVIEW MINUTES (Co-Chairs)

The minutes of the previous meeting were reviewed and approved.

4 TOPIC ITEMS TO BE DISCUSSED (All)

- Discuss the 1st Open Issue – Meta-inventory / Inventory split into multiple objects

  ftp://d9-workgrps@medical.nema.org/MEDICAL/Private/Dicom/WORKGRPS/WG33/2021/2021-02-03/meta-inventory.pptx

Current draft specification (rev 18)

- Content of an inventory may be split into more than one SOP Instance
- Objects are implicitly linked associatively by a Transaction UID in each object identifying the Inventory Group
- Objects in Group are monotonically numbered, last object specifies total number of objects
  - Ensures that all objects are accounted under implicit associative link
- Complete list of objects in Group appears only in completion notification for the Inventory Initiation service (if that was in fact used to create the inventory), or in the responses to an Inventory FIND query (if that service is in fact implemented)

Must find all inventory objects – how does it know it has them all?
Directory with bunch of files, open them all, presume they are DICOM files and find the last one, and how many are there supposed to be?
Is there a more robust way to get the list than some ad hoc method?

Issue for user apps

- For the client app to know it has all the objects in the Inventory Group (without N-EVENT REPORT or C-FIND), it must find all objects, open them all, ensure it has objects numbered 1 to n with identical Transaction UID, and that object n has the Number of Instances in Group attribute with value n
  - Presumably, the client has information about the location of Inventory objects (e.g., in a specific filesystem folder)
  - Process could be facilitated by additional file naming conventions or filesystem capabilities (symbolic file links)

Seems to be more of an issue for the creator than the user.

Issues for creator apps

- Monotonic serialization of objects problematic with distributed inventory creation process
  - E.g., across multiple PACS/VNAs
  - Alternative: explicit listing of group component instances
- Transaction UID problematic for reuse of stable subsets of inventory
  - E.g., existing deep archive inventory must be replicated with different Transaction UID, rather than simply be referenced
  - Alternative: Group established by explicitly identified objects

If you have multiple subsystems that have pieces of the archive and you distribute the process to inventory that? How do you monotonically serialize those objects? Those subsystems may already have inventories, and if they do, now you have to go in and mock around with it to change a couple of header attributes.

Example

- Archive includes several subsystems, some of which have stable data (e.g., one with historical studies prior to 2010)
- Each subsystem separately creates or maintains an Inventory
  - For historical subsystems, the set of Studies is static, and new objects are not allowed to be added to them, so a semi-fixed Inventory object is feasible (updated only with demographic changes)
- An Inventory of the complete archive could simply assemble the subsystem inventories by reference, without needing to replicate them just to change some header fields

Additional related specifications

- Scope of Inventory and Content Date/Time in all objects must be identical
  - Could be relaxed with rules on Scope and Date/Time specified in first object of Group
- Study may only appear once (and complete) in Inventory Group
  - Issue for creator where different series of a study are on different storage devices independently inventoried – e.g., post-processing analytic result series
  - Relaxing spec simplifies creator, complicates user (maybe)

It can also come up theoretically in cases when you have different tiers of storage; depending on the heat of the data, when you re-heat the data to make it more

performant, you don't necessarily remove it form the colder tier, in which case both tiers would end up with a full copy of a study in the inventory if we are allowing inventories by reference.

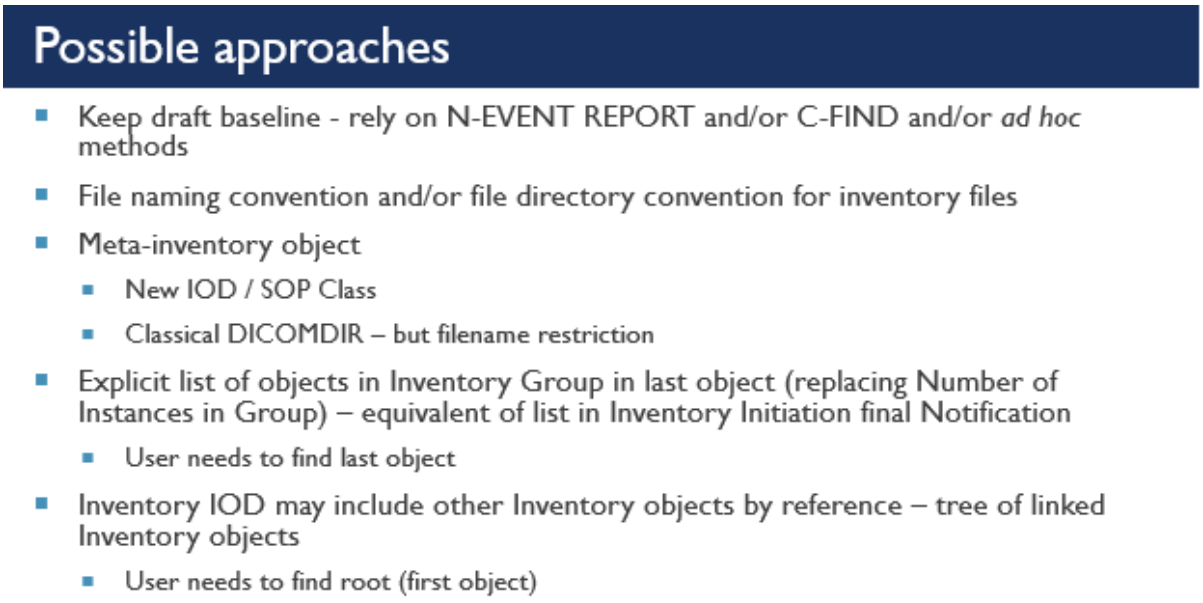

2 real options – last 2 bullets
Problem when you are doing a federated inventory, you don't necessarily know which will be the last one coming in, some process will have to decide that it's done, so it creates the last instance.

The other question – if you are doing federation, is each system that you are forming out to responsible for serializing their own inventory objects or passing back to the main system to serialize them altogether?
The specific use case here is to be able to deliver inventory on media to a new system. That's the part of this that will actually be implemented. People will invest in a tool for importing bulk datasets, not likely to invest in a tool to export them to a successor system.
Preference for bullet # 4 – easier to make it happen.

Could you do an inventory object which contents is just the list of other inventory objects? We aren't creating a new SOP class we are using the same SOP class in two different ways, or in a hybrid approach.
Doesn't necessitate that you create a separate instance that references others, but it allows you to reuse the last instance or create a separate one. Depends on how we spec it out, but that would be the intention.

2 separate issues – 1 whether we can reuse and the other – whether we get the list.

If we reuse – we can't have transaction UID as definitive. So, you'd have a set of inventory objects that have lists of studies in them, and another object. The 1st object that sets the scope of inventory and tells you which inventory objects comprise that inventory - it can be fragile to do it, if somehow the inventory object that provides the link is corrupted, then you lose the link.
In our current draft if you lose any of the objects then the whole object is suspect.

Ability to re-use ties in with ability to have the same study show up more than once – as part of the same inventory doing a delta.
So, here is the inventory auto generated at the beginning of the month plus here is one from the last 4 days – so it doesn't have to do all the work to generate the entire one history. Want to keep those in a merged inventory or keep them separate?

At the end of the day a new IOD is going to be easier to manage – but we need a way to manage it or abdicate responsibility to manage it - forces it down on the client to find a way to figure it out. Vote for a separate IOD if we are going to talk about and include standards around how to update meta-inventory as object availability changes.

Action: Harry to do 2 sketches of how this would work.

Next issue - Non Dicom file paths

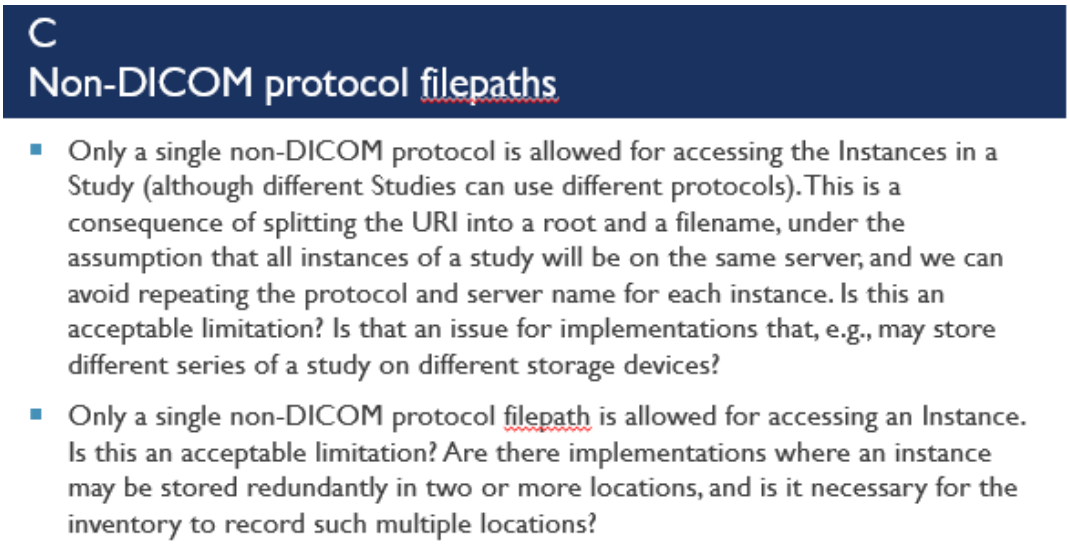

Coming down to an approach that – splitting the URI into the root and filepath was an optimization because in most cases all the instances at least in the series, probably in the study, are all going to have the same roots to a certain extent, just a different file name.

G – Studies Inventoried Once
Consensus - studies may appear multiple times in the inventory and the client will have to deal with it.

L – Tape Management
Tapes are a good example of a non-standard file protocol and we said we exclude it. Attribute at the instance level – a media identifier – if the PACS knows something it can put it there, if not, then no. If the PACS knows, it can put the media ID into the attribute, and you can parse through the inventory to find it.

On another topic – SIIM is working on the SIIM21 Annual Meeting and Matt would like to submit a proposal for a session dedicated to the work of WG-33 - talk about the relationship between DICOM and SIIM (Cheryl), and then talk about the work on the supplement and pitch public comment. Would like Harry to participate in the panel. Harry agreed. More info to come.

5 OLD BUSINESS

6 NEW BUSINESS

7 DATE AND TIME OF NEXT MEETINGS (Secretariat)

- Continue T-con meetings bi-weekly
- Next call is February 17, 2021 between 11:00 am and 12:30 pm ET

NEMALINK CODE	09-WG33
SUBMITTED BY	Hull, Carolyn
SUBMITTED ON	2/16/21
LEGAL APPROVAL	2/25/21
UPLOAD LOCATION	Enter upload location.