

White Paper

The DICOM standard

What's inside?

- What is DICOM?
- How is DICOM used within a hospital environment?

Jonathan Whitby
Integration/Technical Support Manager
Barco
jonathan.whitby@barco.com



Copyright © 2007
BARCO n.v., Kortrijk, Belgium

All rights reserved. No part of this publication may be reproduced in any form or by any means without written permission from Barco.



v.02
29 Jan 2007

1	<i>Introduction</i>	4
2	<i>CT scanner console creates DICOM study</i>	5
3	<i>CT scanner console sends study to archive</i>	7
4	<i>Workstation retrieves study from archive</i>	10
5	<i>Workstation sends reconstructed data to archive</i>	13
6	<i>Additional DICOM concepts</i>	13

1 INTRODUCTION

DICOM, which stands for Digital Imaging and Communications in Medicine, is a standard that was developed to “promote communication of digital image information, regardless of device manufacturer”¹ within a radiological environment.

It defines:

- How the radiological objects (images, reports, measurements and other related patient documentation) should be structured
- How radiological objects should be exchanged

The standard is both a file format and a transfer protocol, and it is not tied to a particular manufacturer, hardware device, operating system or software application.

This is similar to other common formats such as JPEG image files, which are created and viewed by many applications with few interoperability problems. Another example is email servers that can send messages to any other email server regardless of the software or hardware being used.

Because the standard is open, anyone can write software that supports it.

However, while DICOM is an open standard, it is not a simple one. The standard itself is over two thousand pages long, and it makes very little use of commonly understood terminology.

Fortunately, what most people in a radiology department need to know about DICOM is not that complex. It is sufficient to know how it is used and some of the key concepts and terminology.

Some common types of DICOM devices are:

- Modalities (CT, MR, CR)
Current scanners follow the DICOM standard when they create, send and print images
- Picture Archiving and Communication System (PACS)
Most radiology departments have a PACS to handle the archiving and distribution of radiological images and image-related objects. A PACS will typically communicate with modalities and other third-party systems via DICOM but may use proprietary protocols when communicating with its own clients.
- Reporting / post-processing workstations
Many facilities have workstations capable of receiving DICOM images from a modality or PACS and producing additional radiological reports, measurements and images, which can then be sent another DICOM device, such as an archive. These workstations may be part of a PACS or may be third-party systems that work in conjunction with a PACS.

¹ Digital Imaging and Communications in Medicine (DICOM): Part 1: Introduction and Overview, p. iii, National Electrical Manufacturers Association, Rosslyn, VA, USA, 2000

-
- Printers / print servers
Most film printers, and some other medical printers, support the DICOM protocol for printing radiological images.

What follows is a simple example of how DICOM communication is used within the workflow of a radiology department:

1. A CT scan is performed.
2. The scanner console constructs a set of images from the raw data. This set of images is called a **study**.
3. The CT console sends the study to a PACS.
4. A workstation queries the PACS and retrieves the study.
5. Reconstructions or reformats (derived images created from the original study images) are done on the workstation and the derived images are sent back to the archive where they are merged with the rest of the study.




Figure 1 - Typical DICOM Network

2 CT SCANNER CONSOLE CREATES DICOM STUDY

A radiology network will contain information about a number of different patients. Each patient may have been scanned multiple times by different devices.

After a scan is completed, the software within the scanner will create a set of images from the raw data. This is called a **study**.

Sometimes a study consists of several different views or acquisitions, such as pre- and post-contrast CT scans, or T1- and T2-weighted MR views. Each view or acquisition is generally called a **series**.

Each series will contain a number of images or slices. Each slice is a **DICOM Information Object**.

When a sheet of film is printed, a certain amount of patient and study information is printed into the corners of each image. By looking at a sheet of film, even if it has been put in the wrong folder, you can determine what patient it belongs to and when the study was performed. With a reasonable filing system, the film can be restored to its proper place easily.

DICOM Information Objects follow a similar idea. While each image or report is its own Information Object, each contains enough information to specify the patient and study to which it belongs.

For each type of Information Object, there is an **Information Object Definition (IOD)** that specifies what information the object should contain and how the information should be stored. So, for each CT image, the scanner will use the CT Image IOD as a template for how the object should be created.

Each piece of information in a DICOM image is stored as a **DICOM tag**. A DICOM tag consists of a field name and a value. The field names are numerical, but the tags essentially contain information such as "*name=Doe, John*", "*Date of Birth=23 Oct 1996*", etc. Even the images themselves are stored as a DICOM tag.

(0008,0020)	DA	[20000720]	#	8,	1	StudyDate
(0008,0021)	DA	[20000720]	#	8,	1	SeriesDate
(0008,0022)	DA	[20000720]	#	8,	1	AcquisitionDate
(0008,0023)	DA	[20000720]	#	8,	1	ContentDate
(0008,0030)	TM	[231035.000000]	#	14,	1	StudyTime
(0008,0031)	TM	[231258.000000]	#	14,	1	SeriesTime
(0008,0032)	TM	[231441.000000]	#	14,	1	AcquisitionTime
(0008,0033)	TM	[231528.000000]	#	14,	1	ContentTime
(0008,0050)	SH	[N188322]	#	8,	1	AccessionNumber
(0008,0060)	CS	[CT]	#	2,	1	Modality

Figure 2 - DICOM tags

The CT Image IOD specifies which tags should be present in each CT image, as well as what values are valid for some tags. For example, a CT image always has to have the Modality field set to CT.

Obviously all of the images in a study are going to contain a lot of the same information, but this ensures that image data from one patient or study is not confused with demographic information from another.

However, each image contains three tags that are particularly important:

- An **Instance UID** (unique identifier). This uniquely identifies a DICOM object. This allows software to easily determine if an image is already present, rather than having to compare a large number of tags.
- A **Study UID** that uniquely identifies all images in a study. A DICOM study is a set of images with the same Study UID.
- A **Series UID** that uniquely identifies all images in a series. A DICOM Series is a set of images with the same Series UID.

So, when the scanner finishes creating a DICOM study, it will have a set of DICOM CT images where:

- Each object contains all the patient and study information common to the whole study
- The patient information will be the same as any other study belonging to the same patient
- Each object contains the position, size, orientation, etc. of the individual image
- Each object contains the Study UID of the study
- The set of images is grouped into one or more series

- Each object contains the Series UID for its series within the study
- Each object contains its own UID
- Each object contains image data for the image

3 CT SCANNER CONSOLE SENDS STUDY TO ARCHIVE

After the study has been created, the scanner needs to send the study to a remote archive, using the DICOM Storage Service Class (also called a DICOM Push).




Figure 3 - DICOM Storage

Real-world Operation	DICOM service	Service user	Service provider
CT console sends study to PACS	Storage (CT)	CT Console	PACS

A DICOM Service, or Service Class, is a type of DICOM operation (or set of operations) such as Storage, Query/Retrieve, Print, etc. For each type of operation, one device acts as the **Service Class User (SCU)** and the other acts as the **Service Class Provider (SCP)**.

In the case of storage, the device sending the data (the scanner) is the SCU, while the device receiving the data (the PACS) is the SCP. The scanner uses the storage service provided by the archive.

However, a device can act as an SCU for some operations and an SCP for others. For example, the scanner might be able to receive data from other sources, in which case it would be the Storage SCP.

DICOM operations use **TCP/IP (Transmission Control Protocol/Internet Protocol)** for communication with remote devices. TCP/IP is the most common network protocol and is used by email programs, web browsers and most other network.

Every device on a TCP/IP network has an address, called an IP address. No two machines on the same network will have the same IP address, therefore DICOM operations can use this to identify and contact a specific machine.

Also, DICOM SCPs will *listen* for incoming communication on a particular port. It will only *hear* messages that are sent to that port; like a radio tuned to receive signals at a particular frequency. One machine can have many applications listening at different ports, but no two applications on the same machine can listen at the same port.

To connect to an SCP, an SCU connects to the IP address of the SCP and sends data to the particular port used for that DICOM operation. Once the network connection is established, the DICOM communication begins.

Before any data is sent, there is a negotiation phase. During this phase, the SCU (in this case the scanner console) will send:

- The **Application Entity (AE) Title** of the archive, referred to as the **called AE Title**
- Its own AE Title, referred to as the **calling AE Title**
- A list of the operations it wishes to perform (**Abstract Syntaxes**)
- A list of formats it supports (**Transfer Syntaxes**)

Every device in a DICOM network has an Application Entity title, which is used to identify a DICOM service on a device – some devices use a different AE Title for each DICOM operation. AE Titles are either configured by the manufacturer of a device or are determined by the administrator of a DICOM network. AE Titles should be unique within a network.

If the modality tries to send data to the PACS archive but does not give the archive's AE Title correctly, the archive will probably return an error and close the connection. Additionally, the archive may refuse the connection if the scanner's AE Title is not in its list of remote devices.

This is conceptually equivalent to the scanner phoning up the archive and saying, "This is CT5. Can I please speak with ARCHIVE1?" The archive will not talk to the scanner if ARCHIVE1 is not its *name*, and may not talk to the scanner if it does not know any device named CT5.

For this reason, most DICOM devices maintain a list of remote devices. For each remote device, it records information such as the AE Title, IP address and port of the device and potentially a list of services the device supports.

When technologists running the scanner console send the study to an archive, they will choose the device from the console's list, usually by AE Title or some form of alias, and the console software will use the IP address and port that it has on file.

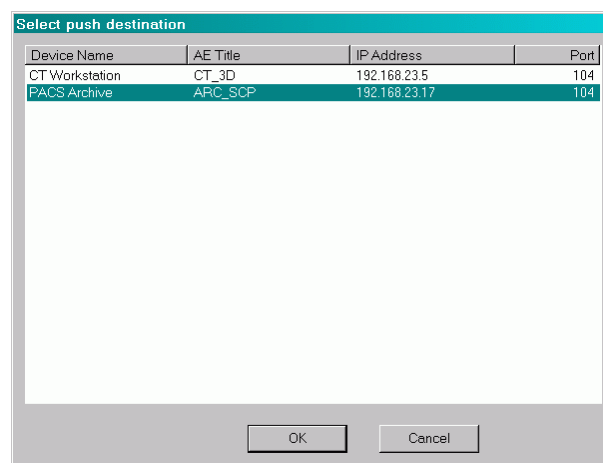


Figure 4 - selecting a DICOM destination

Similarly, when the archive receives the connection request, it matches the calling AE Title against its own list to verify that the scanner is a known device. Many systems will not accept any DICOM communication from unknown devices.

If there are no problems with the AE Titles, the archive will process the list of operations the scanner requested.

DICOM Storage is actually a Service Class, not a single service – it covers a number of different operations. The Information Object Definition (IOD) for CT images defines what a CT image object should look like (as opposed to an MR image object, which has its own IOD). Similarly, DICOM defines CT Image Storage as a specific type of storage (and MR Image Storage as another type). The definitions of how a service applies to an object are called **Service Object Pair (SOP) Classes**.

So if the scanner console wanted to send a study containing CT data, Secondary Capture data and a report, it would request (or **propose**):

- CT Image Storage
- SC Image Storage
- Basic Text SR (Structured Report)

The archive could reply that it accepts CT images and reports, but not SC images, and the console could only send the acceptable data.

If none of the proposed SOP Classes are acceptable to the archive, it will refuse them all and the connection will be closed.

Also, for each SOP class proposed by the console, it can propose a number of **Transfer Syntaxes**. Transfer Syntaxes are basically ways of encoding the data for transmission.

There is a default transfer syntax that all DICOM devices are required to support. In addition, there are many others, mostly types of compression.

The scanner console might propose:

- CT Image Storage with two Transfer Syntaxes:
 1. Implicit Little Endian (the default, uncompressed format)
 2. Lossless JPEG (though there are actually many forms of this).

If the archive accepted both syntaxes, the console would probably send the data using Lossless JPEG since it has lower transmission times. However, if the archive accepted only the default syntax, the archive would send the data using Implicit Little Endian syntax.

Either way, once the archive has replied, the scanner console will send the image data.

The images are sent one at a time, with the archive acknowledging transmission after each one. The archive does not know how many images are in the study, and it simply waits until the console stops sending images and closes the connection.

4 WORKSTATION RETRIEVES STUDY FROM ARCHIVE

A radiologist working on another station decides to pull up the study from the archive. To do this they perform a DICOM query.




Figure 5 - DICOM Query/Retrieve

Real-world operation	DICOM service	Service user	Service provider
1 Workstation queries archive	Query (Find)	Workstation	Archive
2 Workstation retrieves study from archive	Retrieve (Move)	Workstation	Archive
3 Archive sends study to workstation	Storage (CT)	Archive	Workstation

Within the workstation's software, they specify the attributes they want to search for. Most systems support queries by:

- patient name
- patient ID
- date of scan
- accession number
- referring physician

and may also support a number of other attributes.

The screenshot shows a 'Perform Query' dialog box with the following fields and options:

- Device Name: PACS Archive (ARC_SCP)
- Patient Name: DOE
- Study Date:
 - Any Date
 - Today
 - Yesterday
 -
- Modality: CT
- Patient ID: [Empty text box]
- Accession: [Empty text box]
- Study ID: [Empty text box]
- Study Description: [Empty text box]
- Buttons: Query, Cancel

Figure 6 - specifying query attributes

Generally a combination of attributes can be specified with the query, resulting in the list of studies that match all attributes. For attributes such as the Study Date, often a range of values can be specified.

After the radiologist starts the query, the workstation will connect to the archive and open DICOM negotiation. As with storage, the archive is acting as an SCP, this time of the Query (or Find) service, and is therefore listening on a particular TCP/IP port. The workstation is acting as a user (Query SCU) and initiates the connection to the archive's address and port.

The two systems will go through the DICOM negotiation phase and then the workstation will send the query to the archive.

When a study is sent to an archive that supports Query/Retrieve, a certain amount of information about the study is added to a database. This database includes a number of key DICOM tags. When the workstation queries the archive, it sends a list of DICOM tags to the archive, and the archive searches its database and returns a list of studies that match all of the tags it has been sent.

Patient Name	Patient ID	Date of Birth	Sex	Study	Accession	Study Des...	Referring Physician	Date	Time	Modality	Images
HSA Aorta Angio			M			ABD	FREISCHLAG	1993/05/13	09:08	CT	199
HSA Circle of Willis		MR	M			CIRCLE O...	WEISS	1993/04/21	09:55	MR	93
HSA Facial Bone 3D			M			HEAD	DR.NEILFELDST...	1994/10/12	12:00	CT	55

Figure 7 - query results

Many systems support sending a partial attribute through the use of *wildcards*. A wildcard is usually the character *, and if a patient's full name is not known, querying for patient name "Doe*" will return all patients whose last names start with "Doe," such as "Doe,John" and "Doer,John." Many systems will automatically append these wildcards to the end of text attributes (like patient name or study description) so that querying for "Doe*" and "Doe" are equivalent – both will return all patients whose last name starts with "Doe."

In order to successfully use DICOM Query/Retrieve, there are a few things to keep in mind:

- Some systems may not support the same query attributes as others. Also, some are case sensitive, meaning that a query for patient name "Doe_John" will not return studies filed under the patient name "DOE_JOHN." Typically patient names will be stored in all capital letters.
- If no attributes are specified (a blank query), the archive will generally return a list of all studies it has in storage. However, blank queries of large archives can take a long time to complete.
- The list of studies returned by the archive will include a certain amount of information about each study (again, in the form of key DICOM tags) including the Study UID. This is important for the workstation to then retrieve the study.
- A DICOM query is a completely self-contained operation – after the archive has sent the workstation the list of matching studies, it will close the connection.

In order to retrieve the study, the workstation must re-connect to the archive. This time, the workstation acts as an SCU of the Move (or Retrieve) operation, and the archive once again acts as SCP.

After the DICOM negotiation, the workstation sends the archive the UID of the study it wishes to retrieve, and the AE Title of the device to which the study should be sent: the **Move Destination**.

When the archive receives the send destination's AE Title, it checks against its list of remote devices. Unless the archive has the AE Title in its list, along with an IP Address and a Port, it cannot send the study and the retrieve will fail.

If the archive does have an address and port for the send destination, and the UID matches a study in its database, it will accept the query and push the study back to the workstation. So, the archive will start its own Storage SCU and try and connect back to the workstation's Storage SCP.

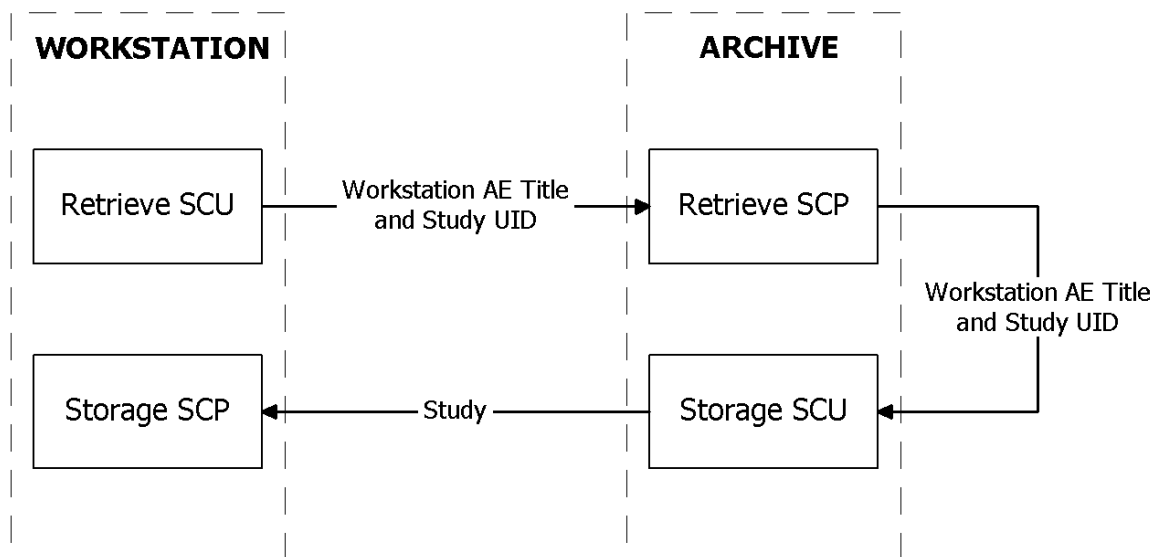


Figure 8 - DICOM Retrieve operation

The archive will connect to the workstation's address and listening port, and go through the DICOM negotiation before sending the study across. The only difference is that both the archive's Retrieve (Move) SCP and the workstation's Retrieve SCU are still running, with the archive updating the workstation on how many images have been sent and how many are remaining. Once the archive finishes sending the data, both DICOM connections will be closed.

DICOM Query/Retrieve is actually more complex than this as it is not always done at the study level. Also, it is often possible to do a relative query; having received the list of studies on an archive, you can select a particular study and re-query on that study, getting information about the various series contained within. However, queries strictly at the study level are the most commonly supported type.

5 WORKSTATION SENDS RECONSTRUCTED DATA TO ARCHIVE

After receiving the data, the radiologist uses the workstation to review the study and perform multi-planar reconstructions and to annotate key images. The additional data is then sent, or pushed, back to the archive, where it is merged into the study.




Figure 9 - DICOM Storage

Real-world operation	DICOM service	Service user	Service provider
Workstation sends reconstructions to archive	Storage (CT, SC)	Workstation	Archive

The transmission of data from one DICOM device to another has been described at length, but it may not be clear how the archive knows these additional objects are part of the same study.

When a workstation creates additional study objects, they are typically given the same patient demographic information as the images from which they are derived. Also, certain information is common to all images within a study, the most important element being the Study UID. By giving the new images the same Study UID as the source images, they should be identified by all DICOM devices as belonging to the same study as the source images.

Also, because each new image is given its own instance UID (each of which must be globally unique) the archive can see that none of these images are already contained in the study, and it will therefore not overwrite any original images with derived ones.

Finally, the derived images are created as one or more new series, each with its own series-level attributes such as Series UID, which again will not match any Series UIDs in the original study.

So, when the archive receives the study, it can easily group it into the appropriate patient, study and series. If the workstation pushed the source images along with the derived images, the archive would be able to tell that it already had images with the same UID and either ignore the incoming images or overwrite the existing images.

6 ADDITIONAL DICOM CONCEPTS

DICOM is still an evolving standard. Already many other services have been incorporated, such as:

- *Work lists and work list management*
By linking into the Hospital/Radiology Information System (HIS/RIS) DICOM devices can automate the input of patient information and the flow of data. The work list manager will send the list of scheduled patients to a scanner, where they can be selected instead of being manually entered. Once the scan is complete

this information is fed back into the work list and any other devices that need access to the data are notified.

- *Structured Reporting*
Reports can be stored in DICOM in such a way that the text and the relevant source images are linked. Additionally, the text in these reports can be extracted and passed to a HIS / RIS system.
- *Security*
Services have been added to secure DICOM networks from unauthorized access to data, and even allow radiologists to digitally sign off on reports or exams.
- *Presentation States*
Scanners and other image creators can send a presentation state along with an image to ensure that the brightness and contrast is consistent between different computer monitors and printed copies.
Additionally, presentation states can be used to define multiple billable acquisitions from a single multislice CT acquisition.