Overview and Major Concepts

DICOM consists of a protocol (or in layman's terms, an agreement on how to communicate) which provides certain functionality and offers an integrated environment which will facilitate an efficient workflow.

DICOM Services and Protocol

The DICOM services use the DICOM protocol, which consist of a set of rules necessary for communication. These rules include commands and an agreement about the meaning of the respective messages (syntax).

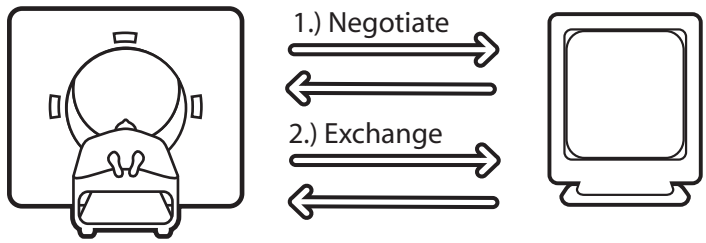

Part 7 of the DICOM standard includes the protocol. A protocol is a set of rules that all communicating parties abide by so that by following these rules communication can occur effectively and efficiently. Part 7 is, in many ways, the "grammar" of the DICOM language.

From a high level (see figure), the DICOM protocol consists of a two-step process. The first step is to communicate what information will be exchanged. This negotiation is critical to ensure that the receiver can indeed understand and interpret the information. DICOM is, therefore, fundamentally different from a broadcast, which is only a one-way transmission. All aspects of the communication are negotiated and agreed upon, including what information is going to be exchanged (i.e. CT, MR, US images, etc.), and how it is encoded (compressed, uncompressed, etc.).

After the negotiation has been completed successfully, step two takes place, during which the actual information exchange occurs. Each interaction is executed by specific commands, which are always being acknowledged. For example, the receiver will always respond back with an indication that the command was received, a guarantee that the message has reached the receiver. (If the message has not been received, an application level time-out warning will typically be generated.)

Also, since DICOM uses a reliable communication protocol as it's "carrier," (i.e., TCP/IP), which has its own error checking, results will be very reliable and have robust end-to-end communication protocol.

With the DICOM commands, there are messages exchanged, such as images. And there are certain rules about the information in those messages. First, there is the location of all the data elements, so that if someone is looking, for example, for the patient name, or the pixel data in order to display the image, he or she knows where to find it. Another way to look at it is to consider the DICOM message a template or form that a device will fill in. Unlike some other protocols, such as HL7, where the data elements are always at a fixed location, and separated by delimiters, the DICOM message is self-describing. It has a tag that identifies what the element represents, which always precedes each data element.

There are also certain rules with regard to the information in the data elements themselves. For example, how do you know where to find the first and last name in the Patient Name field, and where to put the middle initial? This seems trivial, but in certain cultures, the surname (family name) is used first, others always use the given name (first name) first. And, of course, there are even more variations and cultural deviations about a woman's maiden name or even a mother's maiden name.

The data elements require a very concise definition. For example, most people know what a Patient Name is and what it represents. However, does everyone know the difference between a Filler number, Order number, Request number, and Accession number? Some institutions use different names for slightly different entities. The same applies for the actual procedure to be performed; is it a procedure, exam, or study?

After we've agreed upon the correct terminology and meaning, we need to agree upon the relationship between the entities. Can a single study have multiple procedures? Imagine a Radiology Information System (RIS) that can schedule two procedures, each being part of the same study (such as a combined CT/PET procedure). There could be two different diagnostic reports generated by the PACS, one for the CT and one for the PET procedure. If the PACS sends these two reports back to the information system the RIS should not get confused because it could potentially only identify one single report belonging to a single study. The RIS might have a database scheme that can allocate only a single report, identified by a single Accession number for that study. If the relationship of one to many studies and procedures is supported by the PACS, but not by the RIS, there might be an issue. There are also often "work-arounds", like the creation of dummy procedures, or adding Accession numbers (numbers with an extension), and it is important to identify these issues up front so that one can put the appropriate measures in place prior to installing a PACS system.

Relationships between entities are defined in the information model, which might, for example, define that one should be able to support multiple procedures as part of a single study. Adherence

to that information model by all parties involved with the DICOM communication is essential to guarantee true compatibility.

The Many Dimensions of DICOM

The DICOM standard provides many different services, each for a very specific purpose. These services can be considered different levels or "dimensions" which when used in concert can provide an integrated digital imaging environment.

In referring to the many facets of the DICOM standard, some people use the terms "levels" or "dimensions" of DICOM (see figure).

1. The first level is the mechanism to exchange information by checking the presence of the other device (Verification) and exchange objects, such as images, using the DICOM Storage Services (Store). New imaging acquisition techniques require that all of the appropriate information be exchanged specific to those modalities; this causes additional DICOM Storage services to be defined continuously. Examples of recent additions are ophthalmology and radiation therapy, to name just two. When thinking about this first level of DICOM, one could compare it with a delivery service such as the mail. What is in the envelope does not matter – it could be a CT, MR, Ultrasound or other image – what matters is that it gets to where it is supposed to, and does so in a timely fashion.

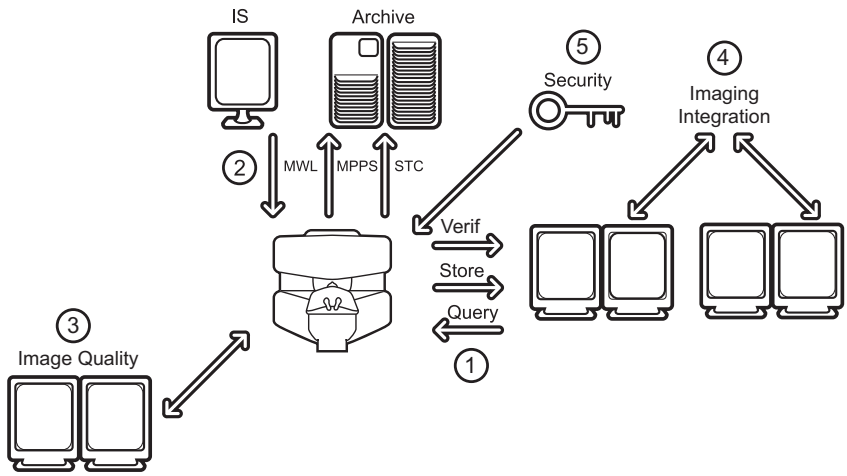

After sending an image using Store (or push), one might want to use Query to retrieve it (or pull). Image retrieval is a common practice for viewing stations to look not only at images stored on an archive, but also for modalities. This is especially true when a technologist wants to use the

same viewing technique for the images acquired as well as used for the previous exam, so that they can be easily compared by a radiologist. This is not uncommon, for example, with a MRI, where techniques often change when new scan protocols are developed and implemented.

For a modality that has very well defined protocols (and where the images do not exactly look alike anyway, such as Ultrasound) the image retrieval at the modality is not as critical. A good rule to follow to try to determine whether one would want the previous images is to determine if, in the film operation, a technologist might retrieve the patient folder to examine the prior exams. If so, this is an indication that the softcopy images should be available as well. If the images must be available and the modality does not provide the "pull" feature, one might need a PACS workstation and/or a web browser to be located close by to view the images. Printing could be considered part of this level as well.

2. The next level of DICOM deals with Image and . As soon as people started to exchange images, the users found out they needed a PACS system administrator, who then spent an overwhelming portion of his time dealing with "orphan" images. Orphan images are ones that are lost, or being temporarily stored in what is euphemistically called the "broken" studies area or "penalty box" at the archive or QA station.

   Why would these mismatches occur? Imagine that I personally might have 3 exams done in a radiology department, an US, CT and MR. More often than not, my images would end up under three different folders, one possibly spelled Osterwik, one maybe spelled Oosterwik, and - who knows, if I'm lucky maybe even an exam filed under the correct spelling, Oosterwijk. There have been no scientific studies conducted that I'm aware of, but I have seen reports that say generally 20-40% of the data entered at a modality by technologists have spelling errors.

   The solution is central data entry and management. DICOM Modality Work List (MWL), which allows for the scheduling, ordering and patient demographic information to be retrieved at a modality, is therefore, a must - and almost all vendors now provide this service at their modalities and the PACS. A technologist merely scans a bar code or selects the patient information from a menu, and the information is copied on the screen. This is a classic no-brainer; people don't want to have to fix misspellings anymore and therefore, MWL is a widely implemented requirement for all new acquisition modality purchases.

   The companion service of the Modality Work List, the Modality Performed Procedure Step (MPPS) allows a device to communicate what the exam performed actually was (vs. the exam that was scheduled) to allow changes in the procedure to be communicated. In addition, it tells when a procedure has been started, indicating this on the scheduling list, and provides information about the number of images generated. And in the case of an X-Ray, a report of the level of radiation dose is communicated with this service. The other service in this category is the

DICOM Storage Commitment (STC), which transfers the responsibility for the images to the receiver, so they can be safely removed from the local disc.

3. Image quality is another dimension in which a lot of improvements have been made, thanks to DICOM standardization. The problem had always been how to achieve consistency in the image presentation of different monitors as well as on a film, independent from the make, type, and characteristics of the media. A "gold standard" for every monitor and hardcopy device has been defined, specifying exactly what luminance or density level would be produced for specific input values. This standard maps these input values into a domain that is perceptually linear. In other words, input values are mapped into a space that is perceived linear for a human observer.

   What is the impact of this standard? The result is that the images look all alike, so that when an image is sent from a radiologist to a physician, he looks at the same grayscale presentation. The same applies for the technologist. He or she wants to make sure that what is seen on the monitor is equivalent of what is presented on the film, even if the printer is in the basement or in the main radiology instead of the outpatient clinic. In order to achieve this consistency, devices have to support the DICOM grayscale or color standard, potentially implementing a so-called presentation Look Up Table to map the values.

   There is one other component of the presentation consistency, and that occurs when a physician does things like zoom an image, add annotation, or change the Window width and level. This information should be preserved in a standard manner - not, as still happens often in a proprietary way, so that it is impossible to view exactly the same presentation on a workstation from another vendor. This issue is addressed by the DICOM Presentation State.

4. The next dimension deals with what one could consider *imaging integration*. The term "integration" is used to indicate information about the images, such as measurements, computer aided diagnostic data for significant images. DICOM identifies the key images of a study so that a physician does not have to review every image in the study (which is especially important when the initial study contains thousands of images such as for CT and MR). The measurements and findings, either by a human or computer, are generally encoded as a Structured Report (SR). For example, an OB/GYN Ultrasound exam of an unborn baby always requires certain measurements to be done by the technologist or physician. This information is currently routinely generated on a form to be filled in, which easily could be replaced with an electronic template. (Note: the images themselves are not stored in these templates, just pointers to them, or if so desired, their thumbnails.)

5. The last important dimension level are the security mechanisms. Remember, the DICOM standard facilitates only the exchange of information. It is merely a part of the overall information chain. It is, therefore, also only a relatively small part of the effort that an institution has to go through to become compliant with federal and local requirements (such as those set forth by the U.S. federal HIPAA[5] regulation). Before a user accesses an image there are procedures and rules about many aspects of that operation, such as the placement of the monitor, for example. There are access controls and authorization rules that are addressed by the application level software, using passwords or even biometric access controls. There is an audit and logging mechanism required that logs any unauthorized data access.

After we have accessed the information and we want to retrieve it using a non-secure communication channel (such as the Internet), DICOM addresses the issue of confidentiality. This is a relatively easy extension; the data can just be encrypted using standard mechanisms and utilities. Then there is the safeguard provided by electronic signatures. This takes care of the data integrity, i.e. prevents anyone from changing information without the receiver noticing it.

How DICOM Facilitates Workflow

DICOM is crucial to an integrated environment. The interplay of people, information and processes can be all summed up in one magic word – workflow. All these levels/dimensions are designed to facilitate an efficient workflow, which is the primary objective of the DICOM standard:

> "The goals of DICOM are to achieve interoperability and to improve workflow efficiency between imaging systems and other information systems in healthcare environments worldwide.[6]"

Let's look at one scenario (see figure opposite page).

1. **Select patient information**: When a patient is registered and a procedure is ordered, this is typically communicated to a Radiology Information System (RIS), which does the scheduling and creates information about the procedure, scheduling, and patient demographics available (using the DICOM Modality Worklist) to the technologist at the modality. This info is then used to perform the acquisition, because the scheduled procedure information (such as *CT-Abdomen*) is available for the technologist. In addition, the patient demographic information is copied automatically and used to identify the images that are sent out.

2. (**Optional**) **retrieve prior images**: In certain cases, prior images could be helpful to ensure that the acquisition uses the same techniques, and that those techniques can be easily compared with the previous study. In such cases, a technologist might *pull* them over by using DICOM Query/Retrieve.

5 Health Insurance Portability and Accountability Act, addressing security and privacy requirements for healthcare institutions.
6 Quote from DICOM Strategic document: http://medical.nema.org.

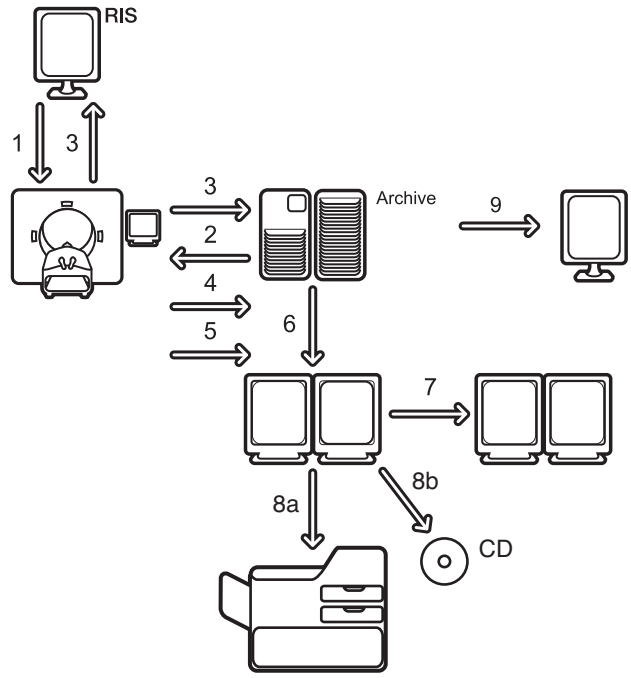

3. **Start procedure**: As soon as the acquisition starts, the PACS and Information System are notified using the DICOM Modality Performed Procedure Step. The Image Manager in the PACS knows that images will be forthcoming and it could remove the procedure from the scheduling list so that other devices will not attempt to perform the same procedure.

4. **Send images**: Images are acquired and sent or "pushed", using the DICOM Storage Service, to a PACS archive or QA station.

5. **Update with exam complete status**: When completed, the DICOM Modality Performed Procedure Step will communicate what the actual performed procedure was (should it have been changed from the originally scheduled procedure), and how many images were created. The number of images is important for the Image Manager, which inventories them to make sure it has them all available in the PACS archive. Knowing that the exam was completed is important for the PACS scheduler so that it can alert a radiologist to read the study. (It does this by adding it to the interpretation Worklist.) The actual performed procedure information will be used by the RIS to ensure that the billing is done correctly. Radiation dose information might be exchanged as well, in the case of an X-Ray exam.

6. **Read study**: A radiologist opens the *to-be-read* folder on a workstation, which was created using the DICOM General Purpose Worklist services. Prior exams might be pulled over using the DICOM Query/Retrieve. Depending on the system architecture, these prior images could

have been pushed to the workstation prior to the reading in order to increase performance (by using the DICOM Storage).

7. **Make images available for physician**: The images could be made available for the Referring Physician as well. They could access the images using a workstation in their office. The DICOM grayscale standard display function is used to make sure that the images look almost identical on the radiologist workstation and the physician workstation. In addition, any additional annotation and changes in the image appearance, such as window width and level, are exchanged between the radiologist workstation and the physician workstation.

8. **Create hardcopy/softcopy**: The patient might need an additional hardcopy, such as film or a CD to take to the physician or specialist. A file room clerk can retrieve the images using DICOM Query/Retrieve and either print the significant images using DICOM Print or burn a DICOM-compatible disc (either CD or DVD).

9. **Make images available for Primary Care Physician**: Using the Internet, the images could also be made available to the Primary Care Physician. DICOM security will make sure that the information delivered over the Internet to the physician is encrypted and that the proper authorization has taken place prior to the image retrieval. The DICOM Key Object Note identifies which images are significant so that the complete study will not have to be reviewed.

This is just one scenario detailing how all these DICOM services work together. Missing one or more of these services will require accommodation in some other way, either by exchanging pieces of paper, manual entry in another information system, or time-consuming additional procedures. The DICOM services as described above are specified in greater detail later in this book.