Liberate Your Image Data

Does Adherence to the DICOM Standard Guarantee Interoperability?

Table of Contents

Abstract 3

The DICOM Standard Does Not Guarantee Interoperability 4

DICOM for Dummies – A Quick Overview 5

DICOM for Dummies – Understanding a Typical DICOM Scenario 6

Why is DICOM So Complex? 7

Resolving DICOM Interoperability Issues 8

Abstract

The purpose of an IT industry standard is to:

1. Improve interoperability between disparate systems
2. Lower system integration costs
3. Reduce the work effort required to implement software from different vendors
4. Improve IT workflow

Wikipedia defines DICOM as follows:

> **D**igital **I**maging and **Co**mmunications in **M**edicine (**DICOM**) is a ***standard*** for handling, storing, printing, and transmitting information in medical imaging.

There is no doubt about the merits of DICOM compliance for healthcare enterprises and imaging technology vendors alike. Without DICOM, it would be difficult to envision the interoperability chaos that would be present in the medical imaging IT market. Having DICOM (and IHE) have made interoperability among PACS, post-processing software and modality software more efficient and more cost effective.

Why then are PACS administrators constantly challenged with complex interoperability issues? The reality is that, more often than not, image files are held captive by a vendor's interpretation of the DICOM standard or a healthcare organization's unique definitions held in the file metadata. The goal of this whitepaper is to explore the root causes of these DICOM challenges and examine the evidence behind the growing trend for PACS-neutral / vendor-neutral solutions.

The DICOM Standard Does Not Guarantee Interoperability

Widely adopted by hospitals, growing in acceptance, and use by other healthcare providers, DICOM (Digital Imaging in Communications in Medicine) is the standard developed by the American College of Radiology and the National Electrical Manufacturers Association for managing medical imaging data. DICOM enables the integration of medical image data between modalities, PACS and post processing workstations for the purpose of storing, retrieving, sharing and displaying image data.

Though DICOM has reached a level of ubiquitous acceptance among medical imaging technology providers and healthcare enterprise organizations, understanding its subtle implications is critical for users evaluating so-called "DICOM compliant" solutions. Questions that need to be addressed:

> How can a PACS administrator find out whether or not two DICOM "compliant" systems interoperate as promised by both vendors *before* purchasing a PACS, a post-processing application or a new modality?
>
> What should a user do prior to purchasing the system to ensure the new system can be interconnected with existing systems?

Having a DICOM-compatible device does not necessarily mean that the image data is stored in a DICOM-conformant format. It also does not mean that when the data is imported or exported from a device that the DICOM file will conform to another system's interpretation of DICOM. DICOM data is supposed to be populated in attributes as specified by the DICOM standard. However, the ability to interpret the attribute data depends on what the sending DICOM device has decided to include in each of the recommend attributes. The same goes for the receiving DICOM device, as it may not process all of the data in each of the attributes.

It's important to note that the DICOM standard is a voluntary standard. It is not administered by any DICOM policing organization that provides formal conformance validation. For this reason, vendors developing solutions should test their applications against software implementations which are available in the public domain. Users must consider how to go about testing DICOM interoperability between two vendors who both claim DICOM conformance. Vendors may have tested interoperability internally at their respective facilities, but users still need to ask vendors whether or not a configuration that is identical to theirs has been implemented successfully. And then follow up on these answers with direct verification.

The wide array of possible configurations gets to the heart of the matter. DICOM is a large specification. Claiming "DICOM Compliance" is for all practical purposes meaningless as there is no easy way to verify conformance claims. And, given that most deployments of DICOM-compatible devices are unique, there is no practical way to verify claims by checking with other users. The real challenge for a user is how to determine compliance for a specific implementation. To do so, users must determine what aspects of DICOM are being claimed, and ensure that the claims are in fact compatible.

DICOM for Dummies – A Quick Overview

Understanding the DICOM standard requires more in-depth explanation than is provided in this whitepaper. However, in order to understand why the root causes of DICOM compliance challenges point to the viability of PACS-neutral or vendor-neutral solutions, one must understand the basics of DICOM. DICOM service classes are a good starting point.

DICOM defines several broad categories of capabilities called "service classes". Among these are "Storage", "Query/Retrieve", "Print", "Worklist Management" and "Media Storage". For example, a scanner that only supports the DICOM Print Service class will not be able to communicate with a PACS that only supports Storage Service class, even though both are advertised as being "DICOM compliant".

Virtually all services operate over a network. For these services, the specific type of service is called a "SOP Class". The possible roles are a) requesting the service or b) providing the service. Common IT terminology refers to the user of the service as the "client" or "requester" and the provider of the service the "server". But DICOM shuns those terms to reinforce the point that DICOM considers both sides to be peers. DICOM refers to the user as an "SCU" (for "service class user") and the provider of the service as an "SCP" (for "service class provider"). Each SOP class has its own unique identifier in the form of a string of digits and periods. For DICOM communications to occur one device must be an SCU (user) of the desired type of service and the other device must be an SCP (provider) of exactly the same service.

The DICOM standard clearly defines service class requirements. Systems that claim support for a particular service class must specifically support the service class as outlined by the DICOM standard. To be more specific, SCP and SCU support must be embedded for Verification, Storage, Storage Commitment, and Query/Retrieve. SCU support is required for Print, Worklist Management, and Detached Study Management.

Similarly, DICOM identifies specific types of "services" with each service class. Devices cannot simply claim to support a service class. Vendors must specifically identify in their conformance statement the role(s) that is supported for specific types of services in a service class.

The SCU will present a number of abstract syntaxes and transfer syntaxes. The SCP in turn will respond for each abstract/transfer syntax pair whether or not it will accept the association. The SCU then selects which abstract/transfer syntax from the list of accepted pairs that it will use to transfer. It is the SCU that determines the format (e.g., compressed, lossy 8bit JPEG, lossless) the images will be stored in. If the software in use provides an auto-synch engine which supports a wide range of abstract and transfer syntaxes, it is likely that the SCU will find an acceptable pair. A flexible approach would also be to configure the system to accept the SOP class in the event the solution doesn't have an acceptable pair.

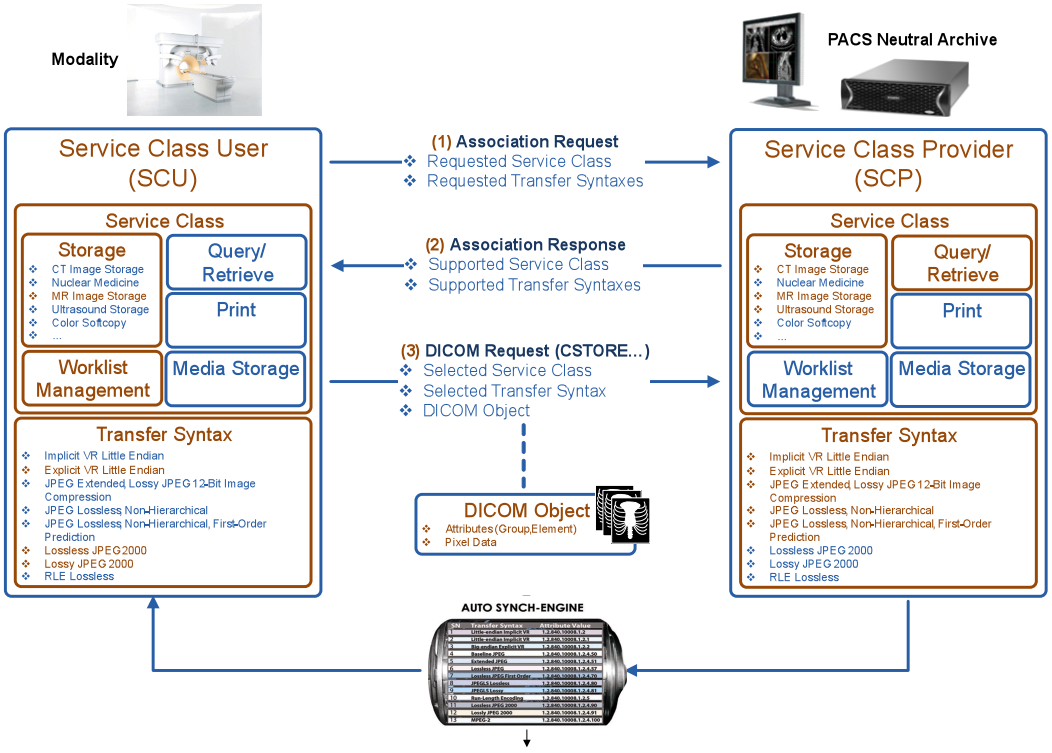

* Highlighted Service Classes and Transfer Syntaxes represented supported components of the SCU and SCP

DICOM for Dummies – Understanding a Typical DICOM Scenario

In the Emergency Room, a digital CT modality scans an unidentified patient. The images are sent from the CT in ER to a PACS viewing station in Radiology. The patient name in the DICOM attributes is “John Doe”. When a patient is unidentified at the time of image generation it is common for these images to be missing key patient identifying data such as the name, patient ID, gender, DoB, etc., that a PACS viewing station often requires to properly retrieve and review prior exams for the patient. After the radiologist performs the initial read on the patient, the PACS forwards the images to the long-term archive, including the missing patient data. As such, the archive now contains patient image data without sufficient data in the attributes to identify the patient in the study. Why is this scenario allowed under the DICOM standard?

The answer is that the smallest elements of DICOM objects are called attributes, examples of which include the patient name, gender, DoB, study description, modality type, etc. Each attribute has a defined “type” in the DICOM standard. There are three “type” definitions. Type 1 means that the attribute is mandatory and must be populated with DICOM compliant data. A Type 2 attribute is also mandatory, but can be sent as "unknown". A Type 3 attribute is optional which means it may or may not have data included.

In the “John Doe” scenario, where the Patient Name, ID, gender, DoB, etc., are not entered into the corresponding attributes, the DICOM image object from the CT in ER is compliant with the DICOM

standard because both of these attributes are defined as Type 2. Notwithstanding this missing patient data, this particular object is uniquely identified with the mandatory Type 1 attributes such as Study, Series, and Image SOP class UID. Using these unique keys, a PACS is capable of storing and displaying these objects. From an archive perspective, these objects either are updated with the correct patient data using an HL7 patient data update transaction or they end up in an "exception queue" which requires reconciliation at a later time by a PACS administrator.

One of the key issues at play here is that the DICOM standard is written like an engineering specification. It is hard to read, and there is no "DICOM for DUMMIES" guide available that can help a user, or even an implementation specialist, through an evaluation of DICOM conformity and its many unique but conforming variances.

Why is DICOM So Complex?

The DICOM specification is rather extensive. Many defined services have yet to be implemented. On the other hand, new technologies such as wavelet compression and new media types are designed outside of DICOM. In one sense, the standard specifies too much, but from another perspective, not enough. Why is that?

The DICOM standard provides a set of tools but does not define a specific architectural solution. Applications that use DICOM are complex and require significant time to implement newly-defined DICOM services. As a result, most vendors take a “piece meal” approach because they are not able to introduce changes to all aspects of their software solution to achieve new DICOM capabilities in one particular release. They adopt the new DICOM standard (or IHE profile) over a series of new software releases.

In reality, some manufacturers put more effort into their DICOM conformance statements than others. Software products vary in their support of DICOM SOP classes. For example, an MR scanner that only supports sending DICOM compliant image objects only supports the DICOM MR Image Storage SOP class. Whereas a PACS workstation must support much more functionality and requires support for many more SOP class specifications.

Resolving DICOM Interoperability Issues

This whitepaper has provided a basic understanding of the DICOM standard. It is safe to say that "the DICOM standard is anything but standard". For the majority of cases, DICOM works effectively. But it is always the "exception use case" that causes the most disruption.

When confronted with a DICOM interoperability issue, customers typically take the following approaches:

1. **Talk to the Vendors** – This is the first place to start. After all, the problem could be resolved by simply making a change to a DICOM parameter by one or both vendors involved. There are challenges with this approach:

    a. Problem definition: You must be able to clearly define the problem. This sounds easy but DICOM interoperability problems are manifested in many different ways. Recognizing a DICOM issue may not be easy.
    b. Getting to the right vendor representative: Most first and second-line support representatives do not fully understand DICOM issues. The first reaction will most likely be "it is the other vendor's issues as our product is DICOM-conformant". Remember, both vendors might take this approach.
    c. Getting both vendors to collaborate on the problem: The customer is in the middle of this rock and hard place. Getting the right person on the phone from both vendors for a productive discussion on how to resolve the interoperability problem can be a real challenge.

    Vendor collaboration on the problem at hand with a focus on problem resolution is a good start. If the vendors can agree on a solution, the solution will often be packaged in a periodic update to the product. Unfortunately, this typically means waiting for the problem to be resolved until the next software update or upgrade from each vendor - a development that could be months away.

2. **Implement a Manual Procedure** – Many DICOM interoperability issues are manifested as a workflow issue. To address the issue at hand, special manual procedures to resolve the issue are defined. This might mean that a technologist or a PACS administrator is called upon to perform special resolution procedures. This will result in unwanted work and time-delays for getting the study to its target destination. In many cases, the manual workaround will resolve the DICOM issue, possibly creating a technologist or PACS administrator productivity (and job satisfaction) issue.

3. **Engage a Professional Services Firm** – Approaches 1 and 2 are the typical first course of action. Only after dealing with these first two approaches will management "consider" engaging a professional services firm to implement a solution to the issue at hand. These solutions can be costly and they typically result in a custom solution that might have to be re-implemented in the future if new vendor software or new product is introduced.

An alternative approach is to install a DICOM "engine" for attribute modification. It is common that DICOM interoperability challenges number more than one. Users can liberate DICOM image data using a middleware application which provides a PACS Administrator or Radiology IT Analyst with the capability to discover the problem, program a rule-based attribute modification or translation table and resolve the issues automatically going forward. Even with new software upgrades from existing PACS vendors, the analyst will be able to implement all required changes to ensure DICOM compatibility between all DICOM systems.

A further step can be taken to migrate the imaging data from proprietary PACS vendor storage to a PACS-neutral or vendor-neutral archive (VNA) with a web-based, zero-footprint, clinical viewer. This type of migration offers distinct benefits. First, it releases and consolidates the image data held captive in disparate departmental and organizational PACS archives. Second, it improves the lifecycle management of image data and reduces the need for extensive, costly data migrations to move to new PACS providers in the future. Third, it eliminates the need for multiple integrations required to provide clinical viewing access to enterprise-wide imaging data. Fourth, when a PACS-neutral / vendor-neutral archive is used in combination with a DICOM engine, all patient images can be shared across any number of disparate systems for improved image workflow, clinician efficiency and patient outcomes.

For more information about Vendor-neutral archives (VNA), see our whitepaper entitled: Achieving a Vendor-Neutral Archive (VNA).