Digital Imaging and Communications in Medicine (DICOM)

Supplement 33: Grayscale Softcopy Presentation State Storage

DICOM Standards Committee, Working Group 11 Display

1300 N. 17th Street

Rosslyn, Virginia 22209 USA

VERSION: Final Text

21 September 1999

Please send comments to Dave Snavely, NEMA (dav_snavely@nema.org)

TABLE OF CONTENTS

TABLE OF CONTENTS....i
Foreword....iii
Scope and Field of Application....iv
4 Symbols and abbreviations....2
A.1.2 IOD Entity-Relationship Model....5
A.X GRAYSCALE SOFTCOPY PRESENTATION STATE INFORMATION OBJECT DEFINITION 11
A.X.1 Grayscale Softcopy Presentation State IOD Description....11
A.X.2 Grayscale Softcopy Presentation State IOD Module Table....12
C.9.2 Overlay plane module....19
C.10.4 Displayed Area Module....20
C.10.5 Graphic Annotation Module....25
C.10.6 Spatial Transformation Module....32
C.10.7 Graphic Layer Module....32
C.11.1 Modality LUT module....34
C.11.2 VOI LUT module....37
C.11.6 Softcopy Presentation LUT Module....41
C.11.7 Overlay/Curve Activation Module....43
C.11.8 Softcopy VOI LUT module....45
C.12.2 Presentation Series Module....46
C.12.3 Presentation State Module....47
F.5.23 Presentation State Directory Record Definition....53
Annex B STORAGE SERVICE CLASS (Normative)....56
B.5 STANDARD SOP CLASSES....56
I.4 MEDIA STORAGE STANDARD SOP CLASSES....56
Annex N GRAYSCALE SOFTCOPY PRESENTATION STATE STORAGE SOP CLASS (Normative) 57
N.1. OVERVIEW....57
N.1.1 SCOPE....57
N.2 GRAYSCALE TRANSFORMATION SEQUENCE....58
N.2.1 Modality LUT....59
N.2.2 Mask....60
N.2.3 VOI LUT....60
N.2.4 Presentation LUT....61
N.2.5 Shutter....61
N.2.6 Pre-Spatial Transformation Annotation....62
N.2.7 Spatial Transformation....62
N.2.8 Post-Spatial Transformation Annotation....62
N.3 BEHAVIOR OF AN SCP....63
N.4 CONFORMANCE....63
N.4.1 Conformance Statement for An SCU....63
N.4.2 Conformance Statement for An SCP....63

Part 5 Addendum : Grayscale Softcopy Presentation State Encoding....65
7.3 BIG ENDIAN VERSUS LITTLE ENDIAN BYTE ORDERING....65
Section 8 Encoding of Pixel *and Overlay* Data....67
8.1 PIXEL AND OVERLAY DATA ENCODING OF RELATED DATA ELEMENTS....67
8.1.1 Pixel data encoding of related data elements....67
8.1.2 Overlay data encoding of related data elements....68
8.2 NATIVE OR ENCAPSULATED FORMAT ENCODING....68
8.2.1 JPEG Image Compression....69
8.2.2 Run Length Encoding Compression....69
A.1 DICOM IMPLICIT VR LITTLE ENDIAN TRANSFER SYNTAX....70
A.2 DICOM LITTLE ENDIAN TRANSFER SYNTAX (EXPLICIT VR)....70
A.3 DICOM BIG ENDIAN TRANSFER SYNTAX (EXPLICIT VR)....71
A.4 TRANSFER SYNTAXES FOR ENCAPSULATION OF ENCODED PIXEL DATA....71
D.2 VARIOUS ADDITIONAL EXAMPLES OF PIXEL AND OVERLAY DATA CELLS....72
Part 6 Addendum : Grayscale Softcopy Presentation State Data Dictionary....74
Index of Attribute Tags and UIDs....77

Foreword

The American College of Radiology (ACR) and the National Electrical Manufacturers Association (NEMA) formed a joint committee to develop a standard for Digital Imaging and Communications in Medicine (DICOM). This DICOM Standard and the corresponding Supplements to the DICOM Standard were developed according to the NEMA procedures.

This Supplement to the Standard is developed in liaison with other standardization organizations including CEN TC251 in Europe and JIRA in Japan, with review also by other organizations including IEEE, HL7 and ANSI in the USA. This Supplement has been prepared by the DICOM Working Group 11 (Display).

The DICOM Standard is structured as a multi-part document using the guidelines established in the following document:

- ISO/IEC Directives, 1989 Part 3 : Drafting and Presentation of International Standards.

This document is a Supplement to the DICOM Standard. It is an extension to PS 3.3, 3.4 and 3.6 of the published DICOM Standard which consists of the following parts:

PS 3.1	-	Introduction and Overview
PS 3.2	-	Conformance
PS 3.3	-	Information Object Definitions
PS 3.4	-	Service Class Specifications
PS 3.5	-	Data Structures and Encoding
PS 3.6	-	Data Dictionary
PS 3.7	-	Message Exchange
PS 3.8	-	Network Communication Support for Message Exchange
PS 3.9	-	Point-to-Point Communication Support for Message Exchange
PS 3.10	-	Media Storage and File Format for Data Interchange
PS 3.11	-	Media Storage Application Profiles
PS 3.12	-	Media Formats and Physical Media for Data Interchange
PS 3.13	-	Print Management Point-to-Point Communication Support
PS 3.14	-	Grayscale Standard Display Function

These parts are related but independent documents.

This Supplement includes the definition of the Grayscale Softcopy Presentation State Information Object and the Grayscale Softcopy Presentation State Storage SOP Class.

Scope and Field of Application

This supplement describes the Grayscale Softcopy Presentation State Storage SOP Class and defines the extensions to the Storage Service Class that support the usage of this SOP Class.

The purpose of the Grayscale Softcopy Presentation State Storage SOP Class is to allow a DICOM application entity to specify how stored pixel data values in a composite image object are to be translated to presentation values, called P-Values, that are independent of device or manufacturer. The display device converts P-Values to luminance in the case of a soft-copy display.

Support of the Grayscale Softcopy Presentation State Storage SOP Class requires that the display device support a specified Grayscale Standard Display Function that defines a linearly perceived response, as defined by the Grayscale Standard Display Function in PS 3.14.

If the originating and receiving display devices have been correctly calibrated to the Grayscale Standard Display Function, then the images displayed on each device will appear very similar in grayscale contrast once the transformations specified in the Presentation State have been applied. Note that the use of P-Values and the presentation state do not automatically improve the appearance of an image (in terms of grayscale contrast), rather they allow improvements in (or choices of) appearance that have been selected by a device or user to be reproduced repeatably.

In addition, the Grayscale Softcopy Presentation State Storage SOP Class allows image display characteristics to be separated from stored image object instances and updated independently.

The Grayscale Softcopy Presentation State Storage SOP Class is defined as a Composite object using the existing Storage Service Class. This allows presentation states to be stored and retrieved in a manner similar to Composite Image objects.

A SOP Class is defined that consolidates existing grayscale transformations, as already defined in Composite Image objects, with the Presentation LUT, and specifies stringent requirements on how a display SCP shall render them.

The SOP Class also defines minimal extensions to allow rotation, zooming and panning, and annotation with vector graphics and plain text. The requirements for an SCP to be able to render these annotations are less stringent than the requirements to replicate the grayscale transformations, in recognition of the fact that a similar meaning may be adequately conveyed without constraining an implementation to an unrealistic extent. In particular, though text and graphics are required to be rendered, the color or grayscale value are only recommended, not required. The font and style of text are not specified.

The images to which a presentation state applies must be part of the same study as the stored presentation state. Images may be selected individually or as a list of images. They may or may not all be from one series.

This supplement specifically does not support the specification of the layout of images on a screen. A future supplement may do so, but at the present time DICOM only defines the use of a conventional “screen save” to generate a secondary capture image that captures intended grayscale transformations as well as window layout and annotation.

This supplement does not support storing the state of color images. It may do so in future, particularly if a Color Display Function Standard is defined for DICOM, using for instance, the CIE LAB space.

This supplement does not address more complex presentation related issues that imply advanced display application functionality. A future supplement or efforts by other Working Groups may cover:

- image processing
- image segmentation
- image registration and fusion
- 3D display, surfaces representation, and rendering
- cine and time based presentation of multi-frame images, including preferred frame rate (beyond the mask subtraction that is defined in this supplement)
- query of Display Device characteristics (perhaps using SNMP)
- advanced graphic annotation including complex objects that may be edited

Since this document proposes changes to existing Parts of DICOM the reader should have a working understanding of the Standard. This proposed Supplement includes a number of Addenda to existing Parts of DICOM :

- PS 3.3 Addendum : Grayscale Softcopy Presentation State Information Object Definitions
- PS 3.4 Addendum : Grayscale Softcopy Presentation State Storage SOP Class
- PS 3.5 Addendum : Grayscale Softcopy Presentation State Storage Encoding
- PS 3.6 Addendum : Grayscale Softcopy Presentation State Data Dictionary

Changes to:

NEMA Standards Publication PS 3.3-1996

Digital Imaging and Communications in Medicine (DICOM)

Part 3: Information Object Definitions

Item: Add abbreviations to PS 3.3 Section 4

4 Symbols and abbreviations

The following symbols and abbreviations are used in this Part of the Standard.

...

BRHC **Bottom Right Hand Corner**

...

TLHC **Top Left Hand Corner**

Item: Amend PS 3.3 Figure 7.1a real world model:

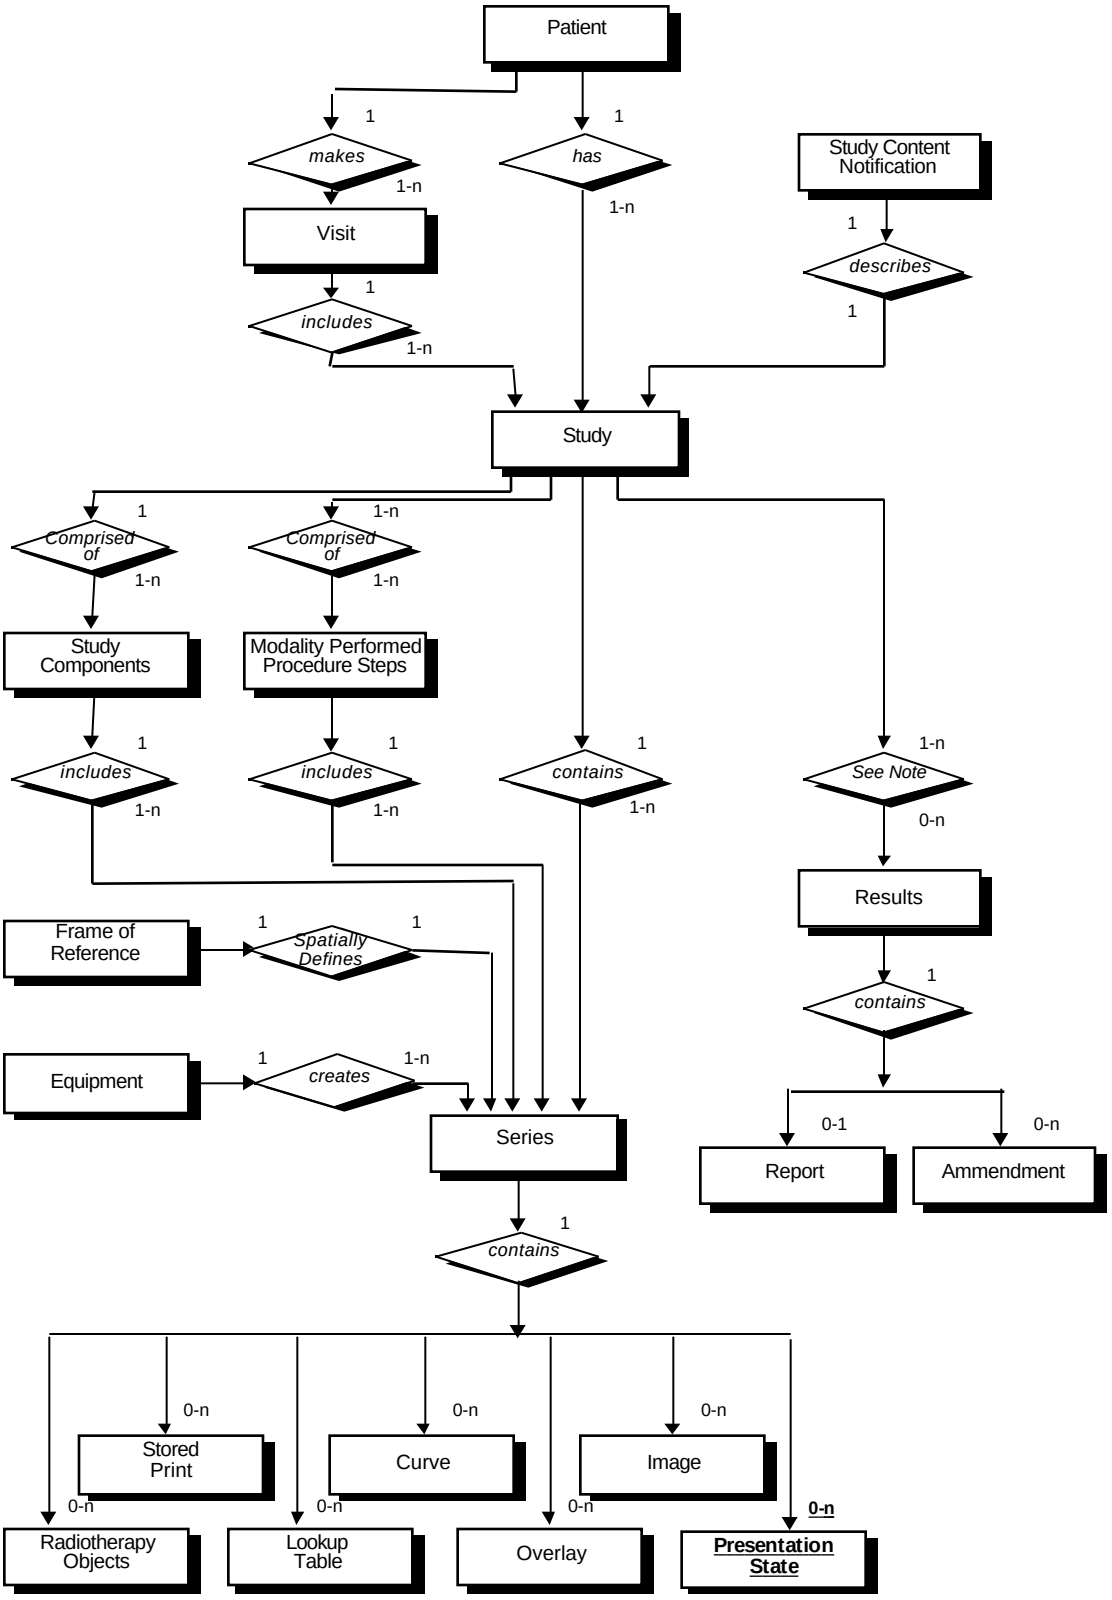

Figure 7-1a
DICOM MODEL OF THE REAL-WORLD

Item: Amend PS 3.3 Figure 7.2a information model:

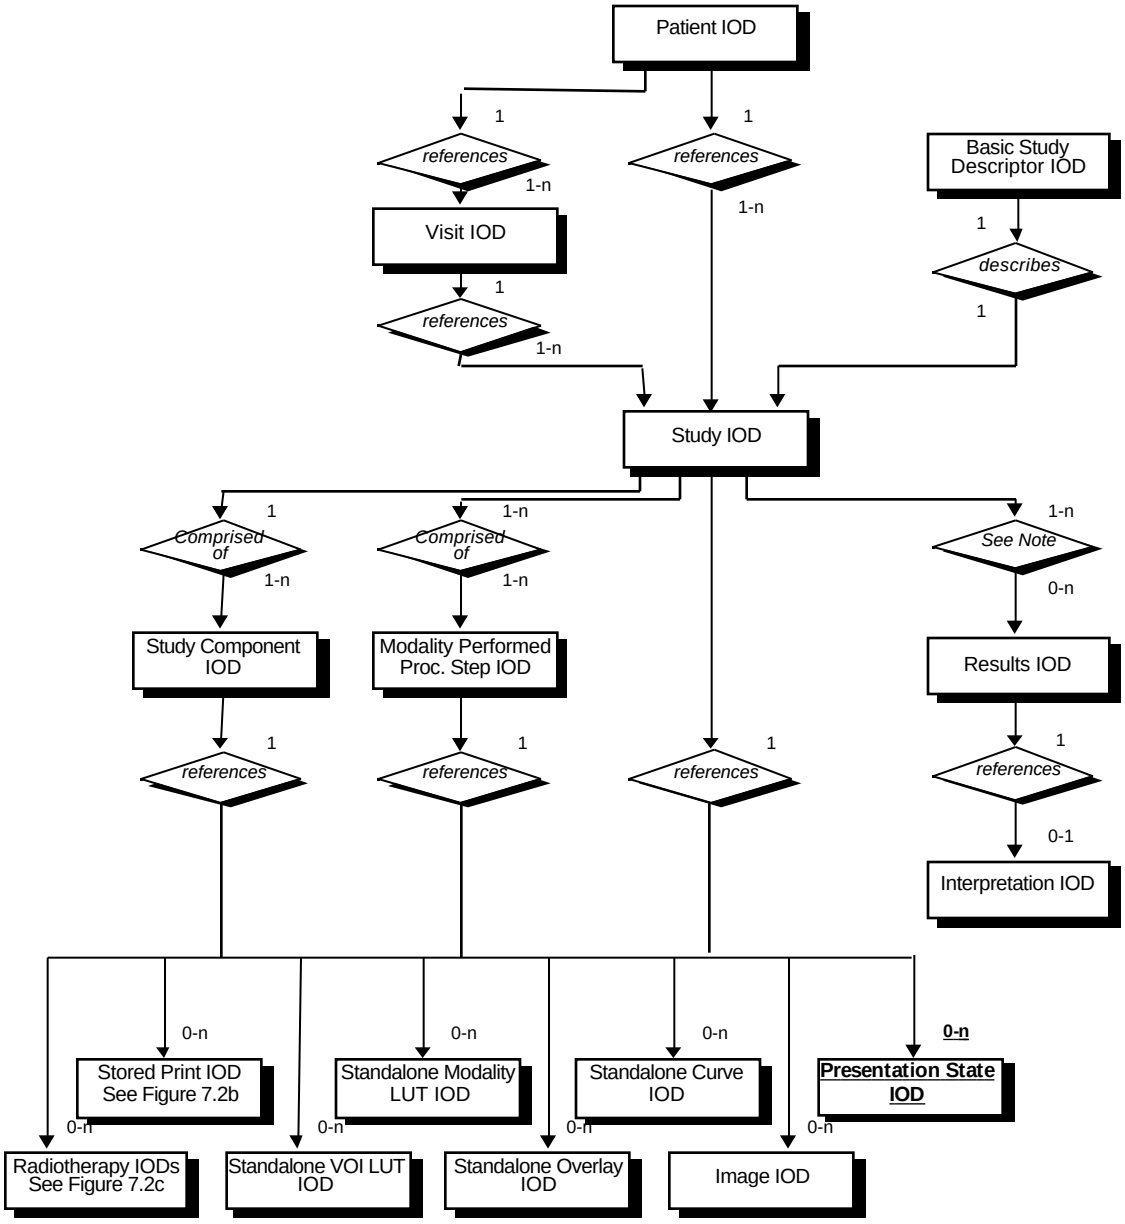

Figure 7-2a
DICOM INFORMATION MODEL

Item: Amend PS 3.3 A.1.2 Composite IOD E-R Model to add Presentation IE and refer to composite instances rather than images.

A.1.2 IOD Entity-Relationship Model

This Section of an IOD provides the Entity-Relationship (E-R) Model which depicts the relationships of the components or Information Entities (IE) of the specified IOD. It forms an IOD specific information model. This E-R model provides the complete context of how the ~~**image**~~ **composite instance** information shall be interpreted when a~~**n image**~~ **composite instance** is exchanged between two DICOM Application Entities.

Even though ~~**images**~~ **composite instances** are sent as discrete individual components, each Composite ~~**Image**~~ **Instance** IOD E-R Model requires that all ~~**images**~~ **composite instances** which are part of a specific study shall share the same context. That is, all ~~**images**~~ **composite instances** within a specific patient study share the same patient and study information; all ~~**images**~~ **composite instances** within the same series share the same series information; etc.

Figure A.1-1 is the DICOM Composite ~~**Image**~~ **Instance** IOD Information Model. It applies to all of the Composite ~~**Image**~~ **Instance** IODs defined in Annex A. However, a subset of this model may be specified by each individual Composite ~~**Image**~~ **Instance** IOD to accurately define the context for specific ~~**modality image**~~ **composite instance** exchange.

Sections A.1.2.1 through A.1.2.10 describe the Information Entities (IE) which comprise the Composite ~~**Image**~~ **Instance** IODs defined in this Annex.

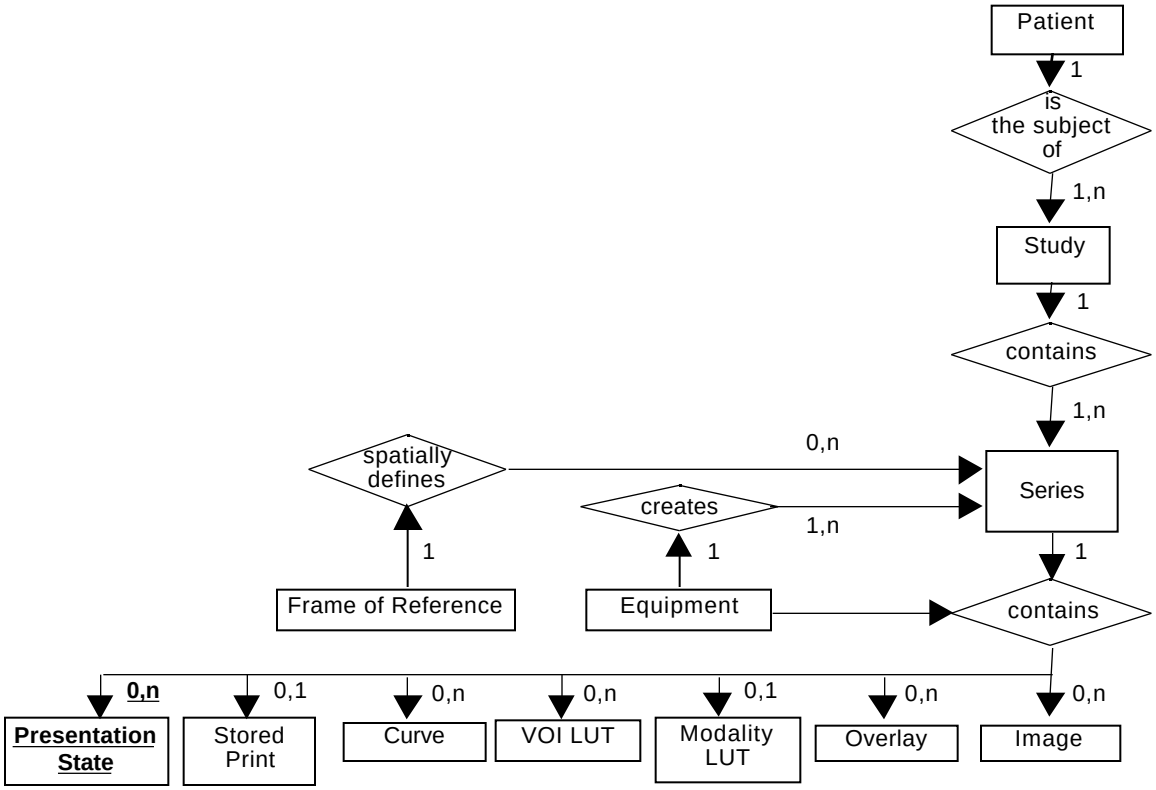

Figure A.1-1
DICOM COMPOSITE ~~IMAGE~~ INSTANCE IOD INFORMATION MODEL

Each Series shall contain at least one Curve IE, VOI Lookup Table **IE**, Overlay IE, **Modality LUT IE, Stored Print IE, Presentation State IE** or Image IE.

A.1.2.1 PATIENT IE

The Patient IE defines the characteristics of a patient who is the subject of one or more medical studies which produce medical images.

The Patient IE is modality independent.

A.1.2.2 STUDY IE

The Study IE defines the characteristics of a medical study performed on a patient. A study is a collection of one or more series of medical images, **presentation states,** overlays and/or curves which are logically related for the purpose of diagnosing a patient. Each study is associated with exactly one patient.

A study may include ~~**images**~~ **composite instances** that are created by a single modality, multiple modalities or by multiple devices of the same modality.

The Study IE is modality independent.

A.1.2.3 SERIES IE

The Series IE defines the Attributes which are used to group images, **presentation states,** overlays and/or curves into distinct logical sets. Each series is associated with exactly one Study.

The following criteria group~~s~~ ~~**images**~~ **composite instances** into a specific series:

a. All ~~**images**~~ **composite instances** within a series must be of the same modality

b. If a specific Composite ~~**Image**~~ **Instance** IOD specifies the support of a Frame of Reference IE, all ~~**images**~~ **composite instances** within the series shall be spatially related to each other; therefore, each series is associated with exactly one Frame of Reference IE

c. If a specific ~~**Image**~~ **Composite Instance** IOD specifies the support of the Equipment IE, all ~~**images**~~ **composite instances** within the series shall be created by the same equipment; therefore, each series is associated with exactly one Equipment IE

d. All ~~**images**~~ **composite instances** within a series shall have the same series information

Overlays and Curves may be grouped into a Series with or without Images. The Equipment IE and Frame of Reference IE are irrelevant to the Overlay IE and Curve IE.

Presentation States shall be grouped into Series without Images (i.e. in a different Series from the Series containing the Images to which they refer). The Frame of Reference IE is irrelevant to the Presentation State IE.

Note: The Series containing Presentation States and the Series containing the Images to which they refer are both contained within the same Study.

A.1.2.4 EQUIPMENT IE

The Equipment IE describes the particular ~~**imaging**~~ device which produced the series of ~~**images**~~ **composite instances**. A~~**n**~~ ~~**imaging**~~ device may produce one or more series within a study. The Equipment IE does not describe the data acquisition or image creation Attributes used to generate the ~~**images**~~ **composite instances** within a series. These Attributes are described in the ~~**Image**~~ **composite instance specific** IE**s (e.g. the Image IE)**.

A.1.2.5 FRAME OF REFERENCE IE

The Frame of Reference IE identifies the coordinate system which conveys spatial and/or temporal information of ~~**images**~~ **composite instances** in a series.

A Frame of Reference IE may be related to one or more series. In this case, it provides the ability to spatially or temporally relate multiple series to each other.

A.1.2.6 IMAGE IE

The Image IE defines the Attributes which describe the pixel data of an image. The pixel data may be generated as a direct result of patient scanning (termed an Original Image) or the pixel data may be derived from the pixel data of one or more other images (termed a Derived Image). An image is defined by its image plane, pixel data characteristics, gray scale and/or color mapping characteristics, overlay planes and modality specific characteristics (acquisition parameters and image creation information).

An image is related to a single series within a single study.

The pixel data within an Image IE may be represented as a single frame of pixels or as multiple frames of pixel data. The frames of a Multi-frame image (a cine run or the slices of a volume) are sequentially ordered and share a number of common properties. A few Attributes may vary between frames (eg.-Time, Angular Displacement, Slice Increment). All common Image IE Attributes refer to the first frame of a multiple frame image.

Overlay, Lookup Table and Curve data may be included within an Image IE only if this information is directly associated with the image.

A.1.2.7 OVERLAY IE

The Overlay IE defines the Attributes which describe an independent set of Overlay Planes. The Overlay IE may represent in a bit-map format, graphics or text and is used to indicate such items as region of interest, reference marks and annotations. These Overlay Planes may or may not be coincident with an image. If the Overlay Plane is coincident with an image, sufficient information shall be available to allow an overlay to be presented at a display station superimposed on a particular image with which it is associated. An Overlay IE shall be related to only one Series IE.

An Overlay Plane may be represented as a single frame (when associated with a single frame image) or as multiple frames of overlay planes (when associated with a Multi-frame image).

Notes: 1. Examples of independent overlay planes are:

a) line drawings which illustrate the equipment and patient setup prescribed

b) line drawings which represent anatomy, pointers and text

c) drawings showing the layout of images and text fields for filming formats

2. The Overlay IE is similar in concept to the 'Graphics Data Set' defined by earlier versions of this Standard.

A.1.2.8 CURVE IE

A Curve is used to represent graphical data that can be specified as a series of connected points. Curve data may or may not be superimposed on a coincident image. An independent Curve, like an independent Overlay, can exist as would an image without any Pixel Data. Curves can be used to specify multi-dimensional graphs, regions of interest, and annotation. Curve Data is not compressed in any of the DICOM Standard Transfer Syntaxes specified in PS 3.5.

Each curve is specified as a series of connected points. One or more Curves shall be described by using one or more even numbered Repeating Groups (5000-501E,eeee) whose attributes are

described in the Curve Module. The Type of Data (50xx,0020) contained in the Curve shall be specified. For independent Curves, the Curve Identification Module is used to identify the Curve.

A.1.2.9 MODALITY LUT IE

The Modality LUT IE defines the Attributes which describe the transformation of manufacturer dependent pixel values into pixel values which are manufacturer independent (e.g. Hounsfield units for CT, Optical Density for film digitizers, etc.). The Modality LUT may be contained within an image, **or a presentation state which references an image,** or as a Standalone Modality LUT which references an image. When the transformation is linear, the Modality LUT is described by Rescale Slope (0028,1053) and Rescale Intercept (0028,1052). When the transformation is non-linear, the Modality LUT is described by Modality LUT Sequence (0028,3000).

A.1.2.10 VOI LUT IE

The VOI LUT IE defines the Attributes which describe the transformation of the modality pixel values into pixel values which are meaningful for print, display, etc. This transformation is applied after any Modality LUT. The VOI LUT may be contained within an image, **or a presentation state which references an image,** or as a Standalone VOI LUT which references an image. When the transformation is linear, the VOI LUT is described by the Window Center (0028,1050) and Window Width (0028,1051). When the transformation is non-linear, the VOI LUT is described by VOI LUT Sequence (0028,3010).

A.1.2.11 PRESENTATION STATE IE

The Presentation State IE defines how a referenced image (or images) will be presented (e.g. displayed) in a device independent grayscale space (i.e. in P-Values), and what graphical annotations and spatial and grayscale contrast transformations will be applied to the referenced image pixel data.

Item: Add to PS 3.3 Table A.1-1

Table A.1-1
COMPOSITE INFORMATION OBJECT MODULES OVERVIEW

IODs Modules	Presentation State
Patient	M
Patient Summary	
General Study	M
Patient Study	U
Study Content	
General Series	M
CR Series	
NM Series	
Presentation Series	M

Frame Of Reference	
U S Frame of Ref.	
General Equipment	M
SC Equipment	
General Image	
Image Plane	
Image Pixel	
NM Image Pixel	
Palette Color Lookup Table	
Contrast/ Bolus	
Cine	
Multi-frame	
NM Multi-frame	
Frame Pointers	
Mask	C
Display Shutter	C
Bitmap Display Shutter	C
Device	
Therapy	
CR Image	
CT Image	
MR Image	
NM Image	
NM Isotope	
NM Detector	
NM TOMO Acquisition	
NM Multi-Gated Acquisition	
NM Phase	
NM Reconstruction	
US Region Calibration	

US Image	
SC Image	
X-Ray Image	
X-Ray Acquisition	
X-Ray Collimator	
X-Ray Table	
XRF Positioner	
XRF Tomo Acquisition	
XA Positioner	
Bi-Plane Sequence	
Bi-Plane Image	
Overlay Identification	
Overlay Plane	C
Multi-frame Overlay	
Bi-Plane Overlay	
Curve Identification	
Curve	
Audio	
Displayed Area	M
Overlay/Curve Activation	C
Graphic Annotation	C
Spatial Transformation	C
Graphic Layer	C
Modality LUT	C
VOI LUT	
Softcopy VOI LUT	C
Softcopy Presentation LUT	M

LUT Identification	
Presentation State	M
SOP Common	M

Item: Add to PS 3.3 new Section A.X

A.X GRAYSCALE SOFTCOPY PRESENTATION STATE INFORMATION OBJECT DEFINITION

A.X.1 Grayscale Softcopy Presentation State IOD Description

The Grayscale Softcopy Presentation State Information Object Definition (IOD) specifies information that may be used to present (display) images that are referenced from within the IOD.

It includes capabilities for specifying:

a. the output grayscale space in P-Values
b. grayscale contrast transformations including modality and VOI LUT
c. mask subtraction for multi-frame images
d. selection of the area of the image to display and whether to rotate or flip it
e. image and display relative annotations, including graphics, text and overlays

A.X.2 Grayscale Softcopy Presentation State IOD Module Table

Table A.X-1
Grayscale Softcopy Presentation State IOD MODULES

IE	Module	Reference	Usage
Patient	Patient	C.7.1.1	M
Study	General Study	C.7.2.1	M
	Patient Study	C.7.2.2	U
Series	General Series	C.7.3.1	M
	Presentation Series	C.12.2	M
Equipment	General Equipment	C.7.5.1	M
Presentation State	Presentation State	C.12.3	M
	Mask	C.7.6.10	C - Required if the referenced image(s) are multi-frame and are to be subtracted
	Display Shutter	C.7.6.11	C - Required if a Display Shutter is to be applied to referenced image(s) and the Bitmap Display Shutter Module is not present
	Bitmap Display Shutter	C.7.6.14	C - Required if a Display Shutter is to be applied to referenced image(s) and the Display Shutter Module is not present
	Overlay Plane	C.9.2	C - Required if Overlay is to be applied to referenced image(s) or the Bitmap Display Shutter Module is present
	Overlay/Curve Activation	C.11.7	C- Required if referenced image contains curve or overlay data which is to be displayed
	Displayed Area	C.10.4	M
	Graphic Annotation	C.10.5	C - Required if Graphic Annotations are to be applied to referenced image(s)
	Spatial Transformation	C.10.6	C - Required if rotation, flipping or magnification are to be applied to referenced image(s)
	Graphic Layer	C.10.7	C - Required if Graphic Annotations or Overlays or Curves are to be applied to referenced image(s)
	Modality LUT	C.11.1	C - Required if a Modality LUT

		is to be applied to referenced image(s)
Softcopy VOI LUT	C.11.8	C - Required if a VOI LUT is to be applied to referenced image(s)
Softcopy Presentation LUT	C.11.6	M
SOP Common	C.12.1	M

In the Grayscale Softcopy Presentation State IOD, the Presentation Series Module specializes some Attributes of the General Series Module, and the Presentation State Module specializes some Attributes of the Mask and Display Shutter Modules.

Notes: 1. Subtraction between different images is not supported.

2. The Mask Module condition implies that it need not be supported by an SCP that supports presentation states only for single frame image storage SOP Classes, or instances of multi-frame image Storage SOP Classes that contain only one frame.

3. The Display Shutter may be used to darken image areas that surround important information and exclude extraneous bright areas that increase glare and ambient lighting impairing image interpretation. For example, unexposed areas in a CR image might be obscured using the Display Shutter, rather than permanently replacing image pixels in those areas.

4. This IOD does not support the storage of a multi-frame overlay in the IOD itself, but does support selective activation of multi-frame overlays within the referenced images via the Overlay/Curve Activation Module.

Item: Add to PS 3.3 Section C.7.3.1 new Modality Defined Term

C.7.3.1.1 General Series Attribute Descriptions

C.7.3.1.1.1 Modality

Defined Terms for the Modality (0008,0060) are:

CR	= Computed Radiography	CT	= Computed Tomography
MR	= Magnetic Resonance	NM	= Nuclear Medicine
US	= Ultrasound	OT	= Other
AS	= Angioscopy	BI	= Biomagnetic imaging
CD	= Color flow Doppler	CP	= Culposcopy
CS	= Cystoscopy	DD	= Duplex Doppler
DG	= Diaphanography	DM	= Digital microscopy
EC	= Echocardiography	ES	= Endoscopy
FA	= Fluorescein angiography	FS	= Fundoscopy
LP	= Laparoscopy	LS	= Laser surface scan
MA	= Magnetic resonance angiography	MS	= Magnetic resonance spectroscopy
PT	= Positron emission tomography (PET)	RG	= Radiographic imaging (conventional film/screen)
ST	= Single-photon emission computed tomography (SPECT)	TG	= Thermography
XA	= X-Ray Angiography	RF	= Radio Fluoroscopy
RTIMAGE	= Radiotherapy Image	RTDOSE	= Radiotherapy Dose
RTSTRUCT	= Radiotherapy Structure Set	RTPLAN	= Radiotherapy Plan
HC	= Hard Copy	PR	= Presentation State

Item: PS 3.3 Section C.7.6.10 Mask Module is amended to correct some typographic errors

C.7.6.10 Mask Module

Table C.7-16 specifies the Attributes that describe mask operations for a Multi-frame image.

Table C.7-16
MASK MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Mask Subtraction Sequence	(0028,6100)	1	Defines a sequence which describe**s** mask subtraction operations for a Multi-frame Image.
>Mask Operation	(0028,6101)	1	Defined Term identifying the type of mask operation to be performed. See C.7.6.10.1 for further explanation.

>Applicable Frame Range	(0028,6102)	3	Each pair of numbers in this multi-valued attribute specify a beginning and ending frame number inclusive of a range where this particular mask operation is valid. Discontinuous ranges are represented by multiple pairs of numbers. Frames in a Multi-frame Image~~s~~ are specified by sequentially increasing number values beginning with 1. If this Attribute is missing in this particular sequence item, then the mask operation is applicable throughout the entire Multi-frame ~~i~~Image, subject to certain limits as described in C.7.6.10.1.1.
>Mask Frame Numbers	(0028,6110)	1C	Specifies the frame numbers of the pixel data used to generate this mask. Frame**s** in a Multi-frame ~~i~~Image are specified by sequentially increasing number values beginning with 1. Required if the Mask Operation (0028,6101) is AVG_SUB.
>Contrast Frame Averaging	(0028,6112)	3	Specifies the number of contrast frames to average together before performing the mask operation. If the Attribute is missing, no averaging is performed.
>Mask Sub-pixel Shift	(0028,6114)	3	A pair of floating point numbers specifying the fractional vertical [adjacent row spacing] and horizontal [adjacent column spacing] pixel shift applied to the mask before subtracting it from the contrast frame. See Section C.7.6.10.1.2.
>TID Offset	(0028,6120)	2C	Specifies the offset to be subtracted from the current frame number in order to locate the mask frame in TID mode. If ~~**omitted**~~**zero length**, TID Offset defaults to 1. ~~**Used when**~~ **Required if** Mask Operation (0028,6101) is TID.
>Mask Operation~~s~~ Explanation~~s~~	(0028,6190)	3	Free form explanation of this particular mask operation.
Recommended Viewing Mode	(0028,1090)	2	Specifies the recommended viewing protocol(s). Defined terms: SUB = for subtraction with mask images; NAT = native viewing of image as sent. Note: If an implementation does not recognize the defined term for Recommended Viewing Mode (0028,1090) , reverting to native display mode is recommended.

Note: Frame numbers begin at 1.

C.7.6.10.1 Mask Subtraction Attribute Descriptions

C.7.6.10.1.1 Mask Operation

Mask Operation (0028,610~~0~~**1**) specifies a type of mask operation to be performed. The Defined Terms identifying the mask operation to be performed are as follows:

NONE (No Subtraction) No mask subtraction operation is specified;

AVG_SUB (Average Subtraction) The frames specified by the Mask Frame Numbers (0028,6110) are averaged together, shifted by the amount specified in the Mask Sub-pixel Shift (0028,6114), then subtracted from the contrast frames in the range specified in the Applicable Frame Range (0028,6102) . Contrast Frame **Averaging** (0028,6112) number of frames starting with the current frame are averaged together before the subtraction. If the Applicable Frame Range is not present in this sequence item, the Applicable Frame Range is assumed to end at the last frame number of the image minus Contrast Frame Averaging (0028,6112) plus one;

TID (Time Interval Differencing) The mask for each frame within the Applicable Frame Range (0028,6102) is selected by subtracting TID Offset (0028,6120) from the respective frame number. If the Applicable Frame Range is not present in this sequence item, the Applicable Frame Range is assumed to be a range where TID offset subtracted from any frame number with the range results in a valid frame number within the Multi-frame image.

C.7.6.10.1.2 Mask Sub-pixel Shift

A pair of floating point numbers ~~**specifying the fractional pixel shift**~~ specifying the fractional vertical [adjacent row spacing] and horizontal [adjacent column spacing] pixel shift applied to the mask before subtracting it from the contrast frame. The row offset results in a shift of the pixels along the column axis. The column offset results in a shift of the pixels along the row axis. A positive row offset is a shift toward~~s~~ the pixels of the lower row of the pixel plane. A positive column offset is a shift toward~~s~~ the pixels of the left hand side column of the pixel plane.

Item: Amend PS 3.3 Section C.7.6.11 Display Shutter Module

C.7.6.11 Display Shutter Module

The Display shutter is a geometric mask which may be applied on the image for presentation purposes in order to neutralize the display of any of the pixels located outside of the shutter shape. Geometry of the shutter is specified with respect to a row and column coordinate system where the origin is the upper left hand pixel. This origin is specified by the values 1,1 for row/column. A row coordinate represent**s** a ~~**number of**~~ r~~a~~**o**w spacing (vertical) and a column coordinate represents a column spacing (horizontal). Up to three different shutter shapes may be used and superimposed.

The manner in which the display area is neutralized (black-out, gray, or other means) is~~ beyond the scope of this Standard~~ **defined by the Attribute Shutter Presentation Value (0018,1622), or undefined if this Attribute is absent or empty**.

Table C.7-17
DISPLAY SHUTTER MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Shutter Shape	(0018,1600)	1	Shape(s) of the shutter defined for display. Enumerated Values are: RECTANGULAR CIRCULAR POLYGONAL This multi-valued Attribute shall contain at most one of each Enumerated Value.
Shutter Left Vertical Edge	(0018,1602)	1C	Required if Shutter Shape (0018,1600) is RECTANGULAR. Location of the left edge of the rectangular shutter with respect to pixels in the image given as column.
Shutter Right Vertical Edge	(0018,1604)	1C	Required if Shutter Shape (0018,15~~5~~**6**00) is RECTANGULAR. Location of the right edge of the rectangular shutter with respect to pixels in the image given as column.
Shutter Upper Horizontal Edge	(0018,1606)	1C	Required if Shutter Shape (0018,15~~5~~**6**00) is RECTANGULAR. Location of the upper edge of the rectangular shutter with respect to pixels in the image given as row.
Shutter Lower Horizontal Edge	(0018,1608)	1C	Required if Shutter Shape (0018,15~~5~~**6**00) is RECTANGULAR. Location of the lower edge of the rectangular shutter with respect to pixels in the image given as row.
Center of Circular Shutter	(0018,1610)	1C	Required if Shutter Shape (0018,15~~5~~**6**00) is CIRCULAR. Location of the center of the circular shutter with respect to pixels in the image given as row and column.
Radius of Circular Shutter	(0018,1612)	1C	Required if Shutter Shape (0018,15~~5~~**6**00) is CIRCULAR. Radius of the circular shutter with respect to pixels in the image given as a number of pixels along the row direction.
Vertices of the Polygonal Shutter	(0018,1620)	1C	Required if Shutter Shape (0018,1600) is POLYGONAL. Multiple Values where the first set of two values are: row of the origin vertex column of the origin vertex Two or more pairs of values follow and are the row and column coordinates of the other vertices of the polygon shutter. Polygon shutters are implicitly closed from the last vertex to the origin vertex ~~are~~ **and** shall be non-intersecting polygons.

Shutter Presentation Value	**(0018,1622)**	**3**	**The value used to replace those parts of the image occluded by the shutter, in P-Values, from a minimum of 0000H (black) up to a maximum of FFFFH (white).** **Note: The maximum P-Value for this Attribute may be different from the maximum P-Value from the output of the Presentation LUT, which may be less than 16 bits in depth.**

Item: Add new PS 3.3 Section C.7.6.14 Bitmap Display Shutter Module

C.7.6.14 Bitmap Display Shutter Module

The Bitmap Display Shutter is a bitmap that defines an arbitrary shape which may be applied on the image for presentation purposes in order to neutralize the display of any of the pixels defined in the bitmap.

The manner in which the display area is neutralized (black-out, gray, or other means) is defined by the Attribute Shutter Presentation Value (0018,1622).

The bitmap is specified as a reference to an instance of the Overlay Plane Module C.9.2. The referenced Overlay is specialized such that:

- Overlay Type (60xx,0040) shall be “G”,

- Overlay Bits Allocated (60xx,0100) shall be 1,

- Overlay Bit Position (60xx,0102) shall be 0 and

- Overlay Origin (60xx,0050) shall be 1\1.

Overlay Rows (60xx,0010) and Overlay Columns (60xx,0011) shall be the same as Rows (0028,0010) and Columns (0028,0011) in the image respectively.

A value of 1 in the Overlay Data (60xx,3000) shall indicate a pixel to which the shutter is applied, i.e. replaced with Shutter Presentation Value (0018,1622).

The Overlay specified in this Attribute shall not be activated (used as a conventional overlay) by the Overlay/Curve Activation Module C.11.6.

Table C.7-xx
BITMAP DISPLAY SHUTTER MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Shutter Shape	(0018,1600)	1	Shape of the shutter defined for display. Enumerated Values are: BITMAP This Attribute shall contain one Value.
Shutter Overlay Group	(0018,1623)	1	Specifies the Group (60xx) of an Overlay stored within the Presentation State IOD that contains the bitmap data, as defined in the Overlay Plane Module C.9.2.
Shutter Presentation Value	(0018,1622)	1	The value used to replace those parts of the image occluded by the shutter, in P-Values, from a minimum of 0000H (black) up to a maximum of FFFFH (white). Note: The maximum P-Value for this Attribute may be different from the maximum P-Value from the output of the Presentation LUT, which may be less than 16 bits in depth.

Item: Amend to PS 3.3 Section C.9.2 Overlay Plane Module

C.9.2 Overlay plane module

...

Table C.9-2
OVERLAY PLANE MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
...	...	...	...
Overlay Bits Allocated	(60xx,0100)	1	Number of Bits Allocated in the Overlay. If the overlay data are embedded in the Image Pixel Data (7FE0,0010), the value of this Attribute shall be the same as Bits Allocated (0028,0100). **If the overlay data are stored in the Overlay Data (60xx,3000) Attribute, the value of this Attribute shall be 1.**
Overlay Bit Position	(60xx,0102)	1	Bit in which Overlay is stored. See PS 3.5 for further explanation. **If the overlay data are stored in the Overlay Data (60xx,3000) Attribute, the value of this Attribute shall be 0.**

...	...	...	...

...

Item: Add to PS 3.3 new Section C.10.4

C.10.4 Displayed Area Module

This Module describes Attributes required to define a Specified Displayed Area space.

The Specified Displayed Area is that portion of the image displayed on the device.

If Presentation Size Mode (0070,0100) is specified as SCALE TO FIT, then the specified area shall be displayed as large as possible within the available area on the display or window, i.e. magnified or minified if necessary to fit the display or window space available.

If Presentation Size Mode (0070,0100) is specified as TRUE SIZE, then the physical size of the rendered image pixels shall be the same on the screen as specified in Presentation Pixel Spacing (0070,0101).

If Presentation Size Mode (0070,0100) is specified as MAGNIFY, then the factor that shall be used to spatially interpolate image pixels to create pixels on the display is defined.

Note: If this factor is specified as 1.0, then one image pixel will correspond to one displayed pixel, and if the Specified Displayed Area is the entire image, and it fits on the display, then the number of displayed pixels will equal the number of image pixels.

In all modes, the actual area rendered on a display device may be greater than the Specified Display Area, if the ratio of rows and columns of the Specified Display Area differs from the ratio of rows and columns of the display device or window. The Displayed Area relative annotations specified in C.10.5 Graphic Annotation Module are rendered relative to the Specified Displayed Area, not the actual rendered displayed area.

Notes: 1. The content of a display outside the Specified Display Area is not defined. In particular no padding value (such as black) is specified.

2. In the TRUE SIZE and MAGNIFY modes, if the entire Specified Displayed Area is not visible, then display relative graphic annotations may be obscured.

This Module explicitly specifies the aspect ratio to be used to display the image, even if it is 1:1, and it may be different from that specified in the referenced image.

Notes: 1.Depending on the mode, the aspect ratio is either specified using the Presentation Pixel Aspect Ratio (0070,0102), or derived from the Presentation Pixel Spacing (0070,0101).

2. This explicit definition of aspect ratio implies that graphic objects that are specified relative to the Specified Display Area will not change their shape regardless of the size or shape of the presentation device (e.g. whether a landscape or portrait monitor is used).

3. The mechanism of interpolation, if necessary, is not specified.

4. The image may need to be cropped and scroll bars or a panning mechanism provided in order to provide access to sections of the image that do not fit within the available area on the display or window.

Table C.10-4
DISPLAYED AREA MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Displayed Area Selection Sequence	(0070,005A)	1	A sequence of Items each of which describes the displayed area selection for a group of images or frames. Sufficient Items shall be present to describe every image and frame listed in the Presentation State Module. One or more Items shall be present.
>Referenced Image Sequence	(0008,1140)	1C	Sequence of Repeating Items where each Item provides reference to a selected set of Image SOP Class/SOP Instance pairs that are defined in the Presentation State Module. Required if a sequence item is present, and if the displayed area selection in this Item does not apply to all the images listed in the Presentation State Module.
>>Referenced SOP Class UID	(0008,1150)	1C	Uniquely identifies the referenced SOP Class. Required if sequence item is present.
>>Referenced SOP Instance UID	(0008,1155)	1C	Uniquely identifies the referenced SOP Instance. Required if sequence item is present.
>>Referenced Frame Number	(0008,1160)	1C	Identifies the frame numbers within the referenced SOP Instance to which this displayed area selection applies. Required if sequence item is present and the referenced image is a multi-frame image and the displayed area selection does not apply to all frames.
>Displayed Area Top Left Hand Corner	(0070,0052)	1	The top left (after spatial transformation) pixel in the referenced image to be displayed, given as column\row. Column is the horizontal offset (X) and row is the vertical offset (Y) relative to the origin of the pixel data before spatial transformation, which is 1\1.
>Displayed Area Bottom Right Hand Corner	(0070,0053)	1	The bottom right (after spatial transformation) pixel in the referenced image to be displayed, given as column\row. Column is the horizontal offset (X) and row is the vertical offset (Y) relative to the origin of the pixel data before spatial transformation, which is 1\1.

>Presentation Size Mode	(0070,0100)	1	Manner of selection of display size. Enumerated Values: SCALE TO FIT TRUE SIZE MAGNIFY See C.10.4 for further explanation.
>Presentation Pixel Spacing	(0070,0101)	1C	Physical distance between the center of each pixel in the referenced image (before spatial transformation), specified by a numeric pair – adjacent row spacing (delimiter) adjacent column spacing in mm. Notes: 1. This value may be different from Pixel Spacing (0028,0030) or Imager Pixel Spacing (0018,1164) specified in the referenced image, which are ignored, since some form of calibration may have been performed (for example by reference to an object of known size in the image). 2. If the row and column spacing are different, then the pixel aspect ratio of the image is not 1:1. Required if Presentation Size Mode (0070,0100) is TRUE SIZE, in which case the values will correspond to the physical distance between the center of each pixel on the display device. May be present if Presentation Size Mode (0070,0100) is SCALE TO FIT or MAGNIFY, in which case the values are used to compute the aspect ratio of the image pixels.

>Presentation Pixel Aspect Ratio	(0070,0102)	1C	Ratio of the vertical size and the horizontal size of the pixels in the referenced image, to be used to display the referenced image, specified by a pair of integer values where the first value is the vertical pixel size and the second value is the horizontal pixel size. See C.7.6.3.1.7. Required if Presentation Pixel Spacing (0070,0101) is not present. Notes: 1. This value may be different from the aspect ratio specified by Pixel Aspect Ratio (0028,0034) in the referenced image, or implied by the values of Pixel Spacing (0028,0030) or Imager Pixel Spacing (0018,1164) specified in the referenced image, which are ignored. 2. This value must be specified even if the aspect ratio is 1:1.
>Presentation Pixel Magnification Ratio	(0070,0103)	1C	Ratio of displayed pixels to source pixels, specified in one dimension. Required if Presentation Size Mode (0070,0100) is MAGNIFY. Notes: 1. A value of 1.0 would imply that one pixel in the referenced image would be displayed as one pixel on the display (i.e. it would not be interpolated if the aspect ratio of the image pixels is 1:1). 2. A value of 2.0 would imply that one pixel in the referenced image would be displayed as 4 pixels on the display (i.e. up-sampled by a factor of 2 in each of the row and column directions). 3. A value of 0.5 would imply that 4 pixels in the referenced image would be displayed as 1 pixel on the display (i.e. down-sampled by a factor of 2 in each of the row and column directions). 4. If the source pixels have an aspect ratio of other than 1:1, then they are assumed to have been interpolated to a display pixel aspect ratio of 1:1 prior to magnification.

Notes: 1. In scale to fit mode, the Displayed Area Top Left Hand Corner (TLHC)and Bottom Right Hand Corner (BRHC) have the effect of defining how any zoom or magnification and/or pan has been applied to select a region of an image to be displayed (the Specified Displayed Area), without assuming anything about the size of the actual display.

2. The TLHC and BRHC may be outside the boundaries of the image pixel data (e.g. the TLHC may be 0 or negative, or the BRHC may be greater than Rows or Columns), allowing minification or placement of the image pixel data within a larger Specified Displayed Area. There is no provision to position a zoomed selected sub-area of the image pixel data within a larger Specified Displayed Area.

Item: Add to PS 3.3 new Section C.10.5

C.10.5 Graphic Annotation Module

This Module defines Attributes of vector graphics and text annotation that shall be made available by a display device to be applied to an image. The graphics and text are defined in position and size relative to the image pixel coordinates or the Specified Displayed Area space (defined in C.10.4 Displayed Area Module). A Graphic Annotation shall be related to an Image.

Note: This Module uses a Sequence of Items rather than a Repeating Group (such as the Curve Repeating Group) to avoid limiting the maximum number of annotation items that may be present. The use of a Repeating Group would limit the number of items to 16. The use of Repeating Groups is also noted in PS 3.5 to be deprecated.

Table C.10-5
GRAPHIC ANNOTATION MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Graphic Annotation Sequence	(0070,0001)	1	A sequence of Items each of which represents a group of annotations composed of graphics or text or both. One or more Items shall be present.
>Referenced Image Sequence	(0008,1140)	1C	Sequence of Repeating Items where each Item provides reference to a selected set of Image SOP Class/SOP Instance pairs that are defined in the Presentation State Module. Required if a sequence item is present, and if graphic annotations in this Item do not apply to all the images listed in the Presentation State Module.
>>Referenced SOP Class UID	(0008,1150)	1C	Uniquely identifies the referenced SOP Class. Required if sequence item is present.
>>Referenced SOP Instance UID	(0008,1155)	1C	Uniquely identifies the referenced SOP Instance. Required if sequence item is present.
>>Referenced Frame Number	(0008,1160)	1C	Identifies the frame numbers within the referenced SOP Instance to which this group of annotations applies. Required if sequence item is present and the referenced image is a multi-frame image and the annotations do not apply to all frames.
>Graphic Layer	(0070,0002)	1	The layer defined in the Graphic Layer Module C.10.7 in which the graphics or text is to be rendered.

>Text Object Sequence	(0070,0008)	1C	Sequence that describes a text annotation. One or more Items may be present. Either one or both of Text Object Sequence (0070,0008) or Graphic Object Sequence (0070,0009) are required if the Sequence Item is present.
>>Bounding Box Annotation Units	(0070,0003)	1C	Units of measure for the axes of the text bounding box. Defines whether or not the annotation is Image or Displayed Area relative. Both dimensions shall have the same units. Enumerated Values: PIXEL = Image relative position specified with sub-pixel resolution such that the origin at the Top Left Hand Corner (TLHC) of the TLHC pixel is 0.0\0.0, the Bottom Right Hand Corner (BRHC) of the TLHC pixel is 1.0\1.0, and the BRHC of the BRHC pixel is Columns\Rows (see figure C.10.5-1). The values must be within the range 0\0 to Columns\Rows. DISPLAY = Fraction of Specified Displayed Area where 0.0\0.0 is the TLHC and 1.0\1.0 is the BRHC. The values must be within the range 0.0 to 1.0. Required if Bounding Box Top Left Hand Corner (0070,0010) or Bounding Box Bottom Right Hand Corner (0070,0011) is present.
>>Anchor Point Annotation Units	(0070,0004)	1C	Units of measure for the axes of the text anchor point annotation. Enumerated Values for Anchor Point Annotation Units (0070,0004) are the same as for Bounding Box Annotation Units (0070,0003). Required if Anchor Point (0070,0014) is present.

>>Unformatted Text Value	(0070,0006)	1	Text data which is unformatted and whose manner of display within the defined bounding box or relative to the specified anchor point is implementation dependent. See C.10.5.1.1. The text value may contain spaces, as well as multiple lines separated by either LF, CR, CR LF or LF CR, but otherwise no format control characters (such as horizontal or vertical tab and form feed) shall be present, even if permitted by the Value Representation of ST. The text shall be interpreted as specified by Specific Character Set (0008,0005) if present in the SOP Common Module. Note: The text may contain single or multi-byte characters and use code extension techniques as described in PS 3.5 if permitted by the values of Specific Character Set (0008,0005).
>>Bounding Box Top Left Hand Corner	(0070,0010)	1C	Location of the Top Left Hand Corner (TLHC) of the bounding box in which Unformatted Text Value (0070,0006) is to be displayed, in Bounding Box Annotation Units (0070,0003), given as column\row . Column is the horizontal offset and row is the vertical offset. Required if Anchor Point (0070,0014) is not present.
>>Bounding Box Bottom Right Hand Corner	(0070,0011)	1C	Location of the Bottom Right Hand Corner (BRHC) of the bounding box in which Unformatted Text Value (0070,0006) is to be displayed, in Bounding Box Annotation Units (0070,0003), given as column\row . Column is the horizontal offset and row is the vertical offset. Required if Anchor Point (0070,0014) is not present.
>>Bounding Box Text Horizontal Justification	(0070,0012)	1C	Location of the text relative to the vertical edges of the bounding box. Enumerated Values: LEFT = closest to left edge RIGHT = closest to right edge CENTER = centered Required if Bounding Box Top Left Hand Corner (0070,0010) is present.

>>Anchor Point	(0070,0014)	1C	Location of a point in the image or Specified Displayed Area to which the Unformatted Text Value (0070,0006) is related, in Anchor Point Annotation Units (0070,0004), given as column\row . Column is the horizontal offset and row is the vertical offset. Required if Bounding Box Top Left Hand Corner (0070,0010) and Bounding Box Bottom Right Hand Corner (0070,0011) are not present. May be present even if a bounding box is specified (i.e. Bounding Box Top Left Hand Corner (0070,0010) and Bounding Box Bottom Right Hand Corner (0070,0011) are present).
>>Anchor Point Visibility	(0070,0015)	1C	Flag to indicate whether or not a visible indication (such as a line or arrow) of the relationship between the text and the anchor point is to be displayed. Enumerated Values: Y = yes N = no Required if Anchor Point (0070,0014) is present.
>Graphic Object Sequence	(0070,0009)	1C	Sequence that describes a graphic annotation. One or more Items may be present. Either one or both of Text Object Sequence (0070,0008) or Graphic Object Sequence (0070,0009) are required if the Sequence Item is present.
>>Graphic Annotation Units	(0070,0005)	1	Units of measure for the axes of the graphic annotation. Enumerated Values for Graphic Annotation Units (0070,0005) are the same as for Bounding Box Annotation Units (0070,0003).
>>Graphic Dimensions	(0070,0020)	1	Enumerated Value: 2
>>Number of Graphic Points	(0070,0021)	1	Number of data points in this graphic.
>> Graphic Data	(0070,0022)	1	Coordinates that specify this graphic annotation . See C.10.5.1.2 for further explanation.

>>Graphic Type	(0070,0023)	1	The shape of graphic that is to be drawn. See C.10.5.1.2. Enumerated Values: POINT POLYLINE INTERPOLATED CIRCLE ELLIPSE
>>Graphic Filled	(0070,0024)	1C	Whether or not the closed graphics element is displayed as filled (in some unspecified manner that shall be distinguishable from an outline) or as an outline. See C.10.5.1.2. Enumerated Values: Y = yes N = no Required if Graphic Data (0070,0022) is "closed", that is Graphic Type (0070,0023) is CIRCLE or ELLIPSE, or Graphic Type (0070,0023) is POLYLINE or INTERPOLATED and the first data point is the same as the last data point.

C.10.5.1 GRAPHIC ANNOTATION ATTRIBUTE DESCRIPTIONS

C.10.5.1.1 Unformatted Text Value

The text shall be displayed if any part of the bounding box or anchor point is within the Specified Display Area.

The text need not be confined to within the bounding box, but shall be rendered in a direction from the Top Left Hand Corner (TLHC) of the bounding box to the Bottom Right Hand Corner (BRHC) of the bounding box, even if these coordinates have been specified in an image relative space and then transformed (rotated, flipped or scaled).

Notes: 1. An implementation may render text outside the confines of the bounding box if necessary to display all the specified text.

2. Alternatively, an implementation may choose to render the text in a scrolling box, or a link to another fixed or popup window as appropriate.

Whether the contents of the bounding box completely opacify the underlying image or whether the box is "transparent" is undefined.

Notes: 1. For example, an implementation may choose an "exclusive or" style opacification to be sure that the text is discernible over light and dark portions of the image.

2. Commonly, the region of the bounding box around the text will be rendered "transparently", i.e. the image will be visible, though some implementations may choose to opacify the bounding box behind the text to improve its readability.

An alternative to specifying a bounding box, is to specify an Anchor Point (0070,0014), i.e. some point in an image or Specified Displayed Area that is related to the text. The semantics of this relationship, and the manner of positioning or linking the text to this point, are unspecified.

Notes: 1. For example, a description of a feature may be linked to a point in the image, and when that image is displayed, if it is magnified and panned, the rendered text (and any arrow or line drawn in response to Anchor Point Visibility (0070,0015)) might be repositioned as appropriate so as not to be cropped out of the Specified Displayed Area.

2. As another example, the text could be rendered in a pop-up window when a hypertext link flagged on the displayed image at the location of the Anchor Point (0070,0014) is selected.

3. The bounding box and anchor point need not be defined with the same axis units, i.e. one can be image pixel relative, and the other displayed area relative.

The size, font and rotation of the individual rendered text characters are unspecified.

C.10.5.1.2 Graphic Data and Graphic Type

Graphic Data (0070,0022) contains the points in the graphic annotation, each dimension for the first point, followed by dimensions for second point, etc. For a two dimensional curve: X1, Y1, X2, Y2, etc. The first (X) dimension corresponds to the image or Specified Displayed Area column (horizontal offset), and the second (Y) dimension corresponds to the image or Specified Displayed Area row (vertical offset). The Value Representation of all components of the N-tuple shall be the same. The image or Specified Displayed Area relative drawing space is defined in Graphic Annotation Units (0070,0005).

If Graphic Type (0070,0023) is POINT, then two values (one point) shall be specified and the single point specified is to be drawn.

If Graphic Type (0070,0023) is POLYLINE, then the points are to be interpreted as an n-tuple list of end points between which straight lines are to be drawn.

If Graphic Type (0070,0023) is INTERPOLATED, then the points are to be interpreted as an n-tuple list of end points between which some form of implementation dependent curved lines are to be drawn. The rendered line shall pass through all the specified points.

If Graphic Type (0070,0023) is CIRCLE, then exactly two points shall be present; the first point is to be interpreted as the center and the second point as a point on the circumference of a circle, some form of implementation dependent representation of which is to be drawn.

If Graphic Type (0070,0023) is ELLIPSE, then exactly four points shall be present; the first two points are to be interpreted as the endpoints of the major axis and the second two points as the endpoints of the minor axis of an ellipse, some form of implementation dependent representation of which is to be drawn.

The notion of “open” or “closed” has no inherent meaning in the context of an arbitrary graphic, other than in the condition for the presence of Graphic Filled (0070,0024). The graphic has no semantic notion of an associated observation such as a region of interest, except that which the unformatted text in the same Item may describe.

The choice of pixel value used to represent the graphic on a display is defined in the Graphic Layer Module C.10.7.

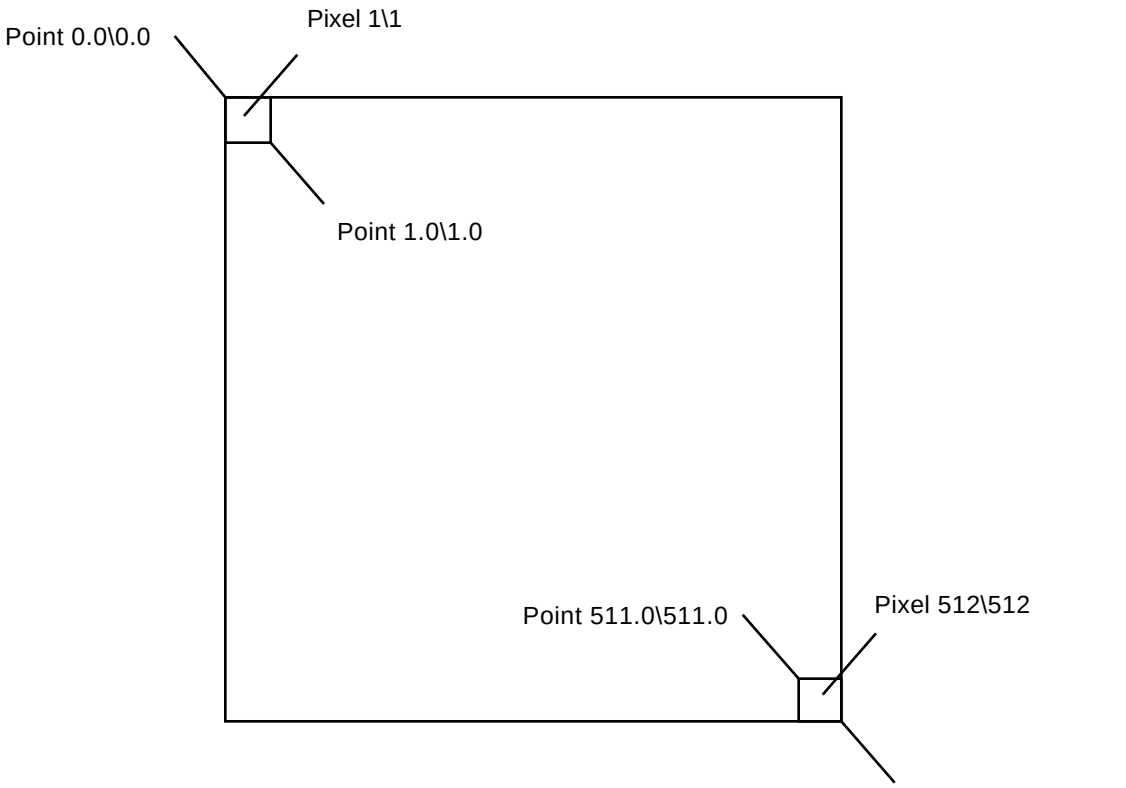

Figure C.10.5-1
Sub-pixel Addressing Units in PIXEL Space

Item: Add to PS 3.3 new Section C.10.6

C.10.6 Spatial Transformation Module

This Module defines a manner of rotating an image by increments of ninety degrees and flipping an image.

Table C.10-6
SPATIAL TRANSFORMATION MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Image Rotation	(0070,0042)	1	How far to rotate the image clockwise in degrees, before any Image Horizontal Flip (0070,0041) is applied. Enumerated Values: 0, 90,180,270 Notes: Negative values are not permitted since the Value Representation is unsigned.
Image Horizontal Flip	(0070,0041)	1	Whether or not to flip the image horizontally after any Image Rotation has been applied such that the left side of the image becomes the right side. Enumerated Values: Y = yes, N = no Note: No vertical flip is specified since the same result can be achieved by a combination of a 180 degree rotation and a horizontal flip.

Note: Given the definition of the Grayscale Transformation Sequence in PS 3.4, it is apparent that the rotation, flipping and magnification will be applied AFTER the application of any bit-mapped overlays or graphic annotations that are specified in the image pixel spaces, but BEFORE the application of graphic annotations that apply in the Specified Displayed Area relative space.

Item: Add to PS 3.3 new Section C.10.7

C.10.7 Graphic Layer Module

This Module defines the characteristics of the layers in which curves, overlays, graphic and text may be rendered.

Layers group together graphics which are related. It is recommended that a layer be displayed such that it may be distinguished from other layers that have a different value for Graphic Layer Order (0070,0062).

Note: The transparency, opacity, and any other interaction (such as exclusive or) with underlying layers or image data are not specified and are at the discretion of the implementation.

Table C.10-7
GRAPHIC LAYER MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Graphic Layer Sequence	(0070,0060)	1	A sequence of Items each of which represents a single layer in which overlays, curves, graphics or text may be rendered. An Item is required for each layer referenced from the Graphic Annotation Module or the Overlay/Curve Activation Module.
>Graphic Layer	(0070,0002)	1	A string which identifies the layer. Note: This identifier may be used by other Attributes within the same presentation state instance to reference this layer. There is no requirement for the same identifiers to be used in different presentation states, and there is no mechanism for referencing layers in other presentation states. That is, a UID is not required.
>Graphic Layer Order	(0070,0062)	1	An integer indicating the order in which it is recommended that the layer be rendered, if the display is capable of distinguishing. Lower numbered layers are to be rendered first.
>Graphic Layer Recommended Display Grayscale Value	(0070,0066)	3	A default single gray unsigned value in which it is recommended that the layer be rendered on a monochrome display. The units are specified in P-Values from a minimum of 0000H (black) up to a maximum of FFFFH (white). Note: The maximum P-Value for this Attribute may be different from the maximum P-Value from the output of the Presentation LUT, which may be less than 16 bits in depth.
>Graphic Layer Recommended Display RGB Value	(0070,0067)	3	A triplet of unsigned RGB values in which it is recommended that the layer be rendered on a color display. The minimum intensity displayable is specified as 0000H\0000H\0000H (black) and the maximum intensity displayable as FFFFH\FFFFH\FFFFH (white).
>Graphic Layer Description	(0070,0068)	3	A free text description of the contents of this layer.

Item: Amend PS 3.3 C.11.1 Modality LUT

C.11.1 Modality LUT module

Table C.11-1 specifies the Attributes that describe the Modality LUT.

Either a Modality LUT Sequence containing a single Item or Rescale Slope and Intercept values shall be present but not both.

Note: This requirement for only a single transformation makes it possible to unambiguously define the input of succeeding stages of the grayscale pipeline such as the VOI LUT.

Table C.11-1
MODALITY LUT MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Modality LUT Sequence	(0028,3000)	~~3~~ **1C**	Defines a sequence of Modality LUTs. **Only one Item may be present. Shall not be present if Rescale Intercept (0028,1052) is present.**
>LUT Descriptor	(0028,3002)	1C	Specifies the format of the LUT Data in this Sequence. See C.11.1.1 for further explanation. Required if the Modality LUT Sequence (0028,3000) is sent.
>LUT Explanation	(0028,3003)	3	Free form text explanation of the meaning of the LUT.
>Modality LUT Type	(0028,3004)	1C	Specifies the output values of this Modality LUT. See C.11.1.1**.2** for further explanation. Required if the Modality LUT Sequence (0028,3000) is sent.
>LUT Data	(0028,3006)	1C	LUT Data in this Sequence. Required if the Modality LUT Sequence (0028,3000) is sent.
Rescale Intercept	(0028,1052)	1C	The value b in relationship between stored values (SV) and the output units specified in Rescale Type (0028,1054). Output units = m*SV + b. Required if Modality LUT Sequence (0028,3000) is not present. **Shall not be present otherwise.**
Rescale Slope	(0028,1053)	1C	m in the equation specified by Rescale Intercept (0028,1052). Required if Rescale Intercept is present.
Rescale Type	(0028,1054)	1C	Specifies the output units of Rescale Slope (0028,1053) and Rescale Intercept (0028,1052). **See C.11.1.1.2 for further explanation.** Required if Rescale Intercept is present.

C.11.1.1 LUT Attribute Descriptions

C.11.1.1.1 LUT descriptor

The three values of the LUT Descriptor (0028,3002) describe the format of the LUT Data in the corresponding Data Element (0028,3006).

The first value is the number of entries in the lookup table. **When the number of table entries is equal to 2^{16} then this value shall be 0.**

The second value is the first stored pixel value mapped. **The Value Representation of the second value (US or SS) is specified by Pixel Representation (0028,0103).** This **stored** pixel value is mapped to the first entry in the LUT. All ~~**image**~~ **stored** pixel values less than the first ~~**entry**~~ **value mapped** are also mapped to the first entry in the LUT Data. A~~**n image**~~ **stored** pixel value one greater than the first ~~**entry**~~ **value mapped** is mapped to the second entry in the LUT Data. Subsequent ~~**image**~~ **stored** pixel values are mapped to the subsequent entries in the LUT Data up to a~~**n image**~~ **stored** pixel value equal to number of entries + first ~~**entry**~~ **value mapped** - 1 which is mapped to the last entry in the LUT Data. ~~**Image**~~ **Stored** pixel values greater than **or equal to** number of entries + first ~~**entry**~~ **value mapped** are also mapped to the last entry in the LUT Data.

The third value specifies the number of bits for each entry in the LUT Data. It shall take the value 8 or 16. The LUT Data shall be stored in a format equivalent to 8 or 16 bits allocated ~~**and**~~ **where the** high bit **is** equal **to** ~~**1**~~ bits allocated **- 1**.

The third value **also** conveys the range of LUT entry values. It shall take the value 8 or 16, corresponding with the LUT entry value range of 256 or 65536.

~~**Note: The third value is not required for describing the LUT data and is only included for informational usage and for maintaining compatibility with ACR-NEMA 2.0.**~~

Note: Since the LUT Descriptor (0028,3002) Attribute is multi-valued, in an Explicit VR Transfer Syntax, only one value representation (US or SS) may be specified, even though the first and third values are always by definition interpreted as unsigned. The explicit VR actually used is dictated by the VR needed to represent the second value, which will be consistent with Pixel Representation (0028,0103).

The LUT Data contains the LUT entry values.

The output range of the Modality LUT Module depends on whether or not Rescale Slope and Rescale Intercept or the Modality LUT Sequence are used.

In the case where Rescale Slope and Rescale Intercept are used, the output ranges from (minimum pixel value*Rescale Slope+Rescale Intercept) to (maximum pixel value*Rescale Slope+Rescale Intercept), where the minimum and maximum pixel values are determined by Bits Stored and Pixel Representation.

Note: This range may be signed even if Pixel Representation is unsigned.

In the case where the Modality LUT Sequence is used, the output range is from 0 to 2^n-1 where n is the third value of LUT Descriptor. This range is always unsigned.

C.11.1.1.2 Modality LUT and Rescale Type

Specifies the units of the output of the Modality LUT **or rescale operation**.

Defined Terms:

OD = The number in the LUT represents thousands of optical density. That is, a value of 2140 represents an optical density of 2.140.

HU = Hounsfield Units (CT)

US = Unspecified

Other values are permitted, but are not defined by the DICOM Standard.

Item: Amend PS 3.3 C.11.2 VOI LUT

C.11.2 VOI LUT module

Table C.11-2 specifies the Attributes that describe the VOI LUT.

Table C.11-2
VOI LUT MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
VOI LUT Sequence	(0028,3010)	3	Defines a sequence of VOI LUTs.
>LUT Descriptor	(0028,3002)	1C	Specifies the format of the LUT Data in this Sequence. See ~~**C.11.1.1**~~ **C.11.2.1.1** for further explanation. Required if the VOI LUT Sequence (0028,3010) is sent.
>LUT Explanation	(0028,3003)	3	Free form text explanation of the meaning of the LUT.
>LUT Data	(0028,3006)	1C	LUT Data in this Sequence. Required if the VOI LUT Sequence (0028,3010) is sent.
Window Center	(0028,1050)	3	Window Center for display. See C.11.2.1.~~**1**~~**2** for further explanation.
Window Width	(0028,1051)	1C	Window Width for display. See C.11.2.1.~~**1**~~**2** for further explanation. Required if Window Center (0028,1050) is sent.
Window Center & Width Explanation	(0028,1055)	3	Free form explanation of the meaning of the Window Center and Width. Multiple values correspond to multiple Window Center and Width values.

C.11.2.1 LUT Attribute Descriptions

C.11.2.1.1 LUT Descriptor

The three values of the LUT Descriptor (0028,3002) describe the format of the LUT Data in the corresponding Data Element (0028,3006).

The first value is the number of entries in the lookup table. When the number of table entries is equal to 2^{16} then this value shall be 0.

The second value is the first input value mapped. The Value Representation of the second value (US or SS) depends on the source of the input to the VOI LUT, and shall be:

- **the same as specified by Pixel Representation (0028,0103), if there is no Modality LUT or Rescale Slope and Intercept specified;**
- **SS if the possible output range after application of the Rescale Slope and Intercept may be signed;**

  Note: This is always the case for the CT Image IOD in which the Rescale Type is specified to be Hounsfield Units, which are always signed.

- **US otherwise.**

This input value is mapped to the first entry in the LUT. All input values less than the first value mapped are also mapped to the first entry in the LUT Data. An input value one greater than the first value mapped is mapped to the second entry in the LUT Data. Subsequent input values are mapped to the subsequent entries in the LUT Data up to an input value equal to number of entries + first value mapped - 1 which is mapped to the last entry in the LUT Data. Input values greater than or equal to number of entries + first value mapped are also mapped to the last entry in the LUT Data.

The third value specifies the number of bits for each entry in the LUT Data. If the VOI LUT is included in an Image IOD, the third value of LUT Descriptor (0028,3002) shall be 8 or 16 bits, unless otherwise specialized. If the VOI LUT is included in a Presentation State IOD, the third value of LUT Descriptor (0028,3002) shall be between 8 and 16 inclusive. The LUT Data shall be stored in a format equivalent to 8 or 16 bits allocated where the high bit is equal to bits stored - 1, where bits stored is the third value.

Note: Since the LUT Descriptor (0028,3002) Attribute is multi-valued, in an Explicit VR Transfer Syntax, only one value representation (US or SS) may be specified, even though the first and third values are always by definition interpreted as unsigned. The explicit VR actually used is dictated by the VR needed to represent the second value.

The LUT Data contains the LUT entry values.

The output range is from 0 to 2^n-1 where n is the third value of LUT Descriptor. This range is always unsigned.

C.11.2.1.~~1~~2 Window center and window width

Window Center (0028,1050) and Window Width (0028,1051) specif~~ies~~**y** a linear conversion from stored pixel values **(after any Modality LUT or Rescale Slope and Intercept specified in the IOD have been applied)** to values to be displayed. Window Center contains the ~~**pixel**~~ **input** value that is the center of the window. Window Width contains the width of the window.

Note: The terms “window center” and “window width” are not consistently used in practice, nor were they defined in previous versions of the standard. The definitions here are presented for the purpose of defining consistent meanings for identity and threshold transformations while preserving the common practice of using integral values for center and width.

Window Width (0028,1051) shall always be greater than or equal to 1.

When Window Width (0028,1051) is greater than 1, these Attributes select the range of input values that are to be mapped to the full range of the displayed output.

When Window Width (0028,1051) is equal to 1, they specify a threshold below which input values will be displayed as the minimum output value.

Note: Whether the minimum output value is rendered as black or white may depend on the value of Photometric Interpretation (0028,0004) or the presence of a Presentation LUT Module.

These Attributes are applied according to the following pseudo-code, where x is the input value, y is an output value with a range from y_{min} to y_{max}, c is Window Center (0028,1050) and w is Window Width (0028,1051):

if (x <= c - 0.5 - (w-1)/2), then y = y_{min}

else if (x > c - 0.5 + (w-1)/2), then y = y_{max},

else y = ((x - (c - 0.5)) / (w-1) + 0.5) * (y_{max} - y_{min})+ y_{min}

Notes: 1. For the purpose of this definition, a floating point calculation without integer truncation is assumed, though the manner of implementation may vary as long as the result is the same.

2. The pseudo-code function computes a continuous value over the output range without any discontinuity at the boundaries. The value of 0 for w is expressly forbidden, and the value of 1 for w does not cause division by zero, since the continuous segment of the function will never be reached for that case.

3. For example, for an output range 0 to 255:

c=2048, w=4096 becomes:

```
if (x <= 0) then y = 0
else if (x > 4095) then y = 255
else y = ((x - 2047.5) / 4095 + 0.5) * (255-0) + 0
```

c=2048, w=1 becomes:

```
if (x <= 2047.5) then y = 0
else if (x > 2047.5) then y = 255
else /* not reached */
```

c=0, w=100 becomes:

```
if (x <= -50) then y = 0
else if (x > 49) then y = 255
else y = ((x + 0.5) / 99 + 0.5) * (255-0) + 0
```

c=0, w=1 becomes:

if (x <= -0.5) then y = 0

else if (x > -0.5) then y = 255

else /* not reached */

4. A Window Center of 2^{n-1} and a Window Width of 2^n selects the range of input values from 0 to 2^n-1. This represents an identity VOI LUT transformation in the case where no Modality LUT is specified and the stored pixel data are n bit unsigned integers.

5. A Window Width of 1 is typically used to represent a "threshold" operation in which those integer input values less than the Window Center are represented as the minimum displayed value and those greater than or equal to the Window Center are represented as the maximum displayed value. A Window Width of 2 will have the same result for integral input values.

6. The application of Window Center (0028,1050) and Window Width (0028,1051) may select a signed input range. There is no implication that this signed input range is clipped to zero.

7. The selected input range may exceed the actual range of the input values, thus effectively "compressing" the contrast range of the displayed data into a narrower band of the available contrast range, and "flattening" the appearance. There are no limits to the maximum value of the window width, or to the minimum or maximum value of window level, both of which may exceed the actual or possible range of input values.

8. Input values "below" the window are displayed as the minimum output value and input values "above" the window are displayed as the maximum output value. This is the common usage of the window operation in medical imaging. There is no provision for an alternative approach in which all values "outside" the window are displayed as the minimum output value.

9. The output of the Window Center/Width or VOI LUT transformation is either implicitly scaled to the full range of the display device if there is no succeeding transformation defined, or implicitly scaled to the full input range of the succeeding transformation step (such as the Presentation LUT), if present. See C.11.6.1.

10. Fractional values of Window Center and Window Width are permitted (since the VR of these Attributes is Decimal String), and though they are not often encountered, applications should be prepared to accept them.

These Attributes shall be used only for Images with Photometric Interpretation (0028,0004) values of MONOCHROME1 and MONOCHROME2. They have no meaning for other Images.

If multiple values are present, both Attributes shall have the same number of values and shall be considered as pairs. Multiple values indicate that multiple alternative views ~~**shall**~~ **may** be presented.

~~**This transformation of pixel values shall be applied after any Modality LUT.**~~

If any VOI LUT Table is included or referenced by an Image, a Window Width and Window Center or the VOI LUT Table, but not both, ~~**it shall**~~ **may** be applied to the Image for display ~~**to a pixel**~~

~~**in the Image for display**~~. Inclusion of both indicates that multiple ~~**viewing**~~ **alternative views** ~~**should**~~ **may** be presented.

If multiple Items are present in VOI LUT Sequence (0028,3010), only one **may** ~~**shall**~~ be applied to the Image for display. Multiple Items indicate that multiple alternative views **may** ~~**should**~~ be presented.

Item: Add PS 3.3 C.11.6 Softcopy Presentation LUT

C.11.6 Softcopy Presentation LUT Module

Table C.11-4 specifies the Attributes that describe the Softcopy Presentation LUT.

Table C.11.6-1
SOFTCOPY PRESENTATION LUT MODULE ATTRIBUTES

Attribute name	Tag	Type	Description
Presentation LUT Sequence	(2050,0010)	1C	Defines a sequence of Presentation LUTs. Only a single item shall be included in this sequence. Required if Presentation LUT Shape (2050,0020) is absent.
>LUT Descriptor	(0028,3002)	1C	Specifies the format of the LUT Data in this Sequence. See C.11.6.1.1 for further explanation. Required if a Sequence Item is present.
>LUT Explanation	(0028,3003)	3	Free form text explanation of the meaning of the LUT.
>LUT Data	(0028,3006)	1C	LUT Data in this Sequence. Required if a Sequence Item is present.
Presentation LUT Shape	(2050,0020)	1C	Specifies predefined Presentation LUT transformation. Required if Presentation LUT Sequence (2050,0010) is absent. Enumerated Values : IDENTITY - no further translation necessary, input values are P-Values INVERSE - output values after inversion are P-Values See C.11.6.1.2.

Note: This Module differs from the Presentation LUT Module used in the hardcopy (print) related SOP Classes in that Optical Density is not supported for Presentation LUT Shape (since Optical Density has no meaning for softcopy display devices).

C.11.6.1 Softcopy Presentation LUT Attributes

When the Presentation LUT is specified as a Presentation LUT Sequence, then the input range of values is specified by the LUT Descriptor as the first value mapped and the number of entries (values mapped). However, there is an implicit linear scaling of the output range of the preceding

transformation (such as the VOI LUT transformation) so that it is always mapped to the specified input range of the Presentation LUT.

When the Presentation LUT is specified as Presentation LUT Shape, then the input range is implicitly specified to be the output range of the preceding transformation (VOI LUT, or if the VOI LUT is identity or absent, the Modality LUT, or if the Modality LUT and VOI LUT are identity or absent, the stored pixel values). In this case, the full range of the output of the preceding transformation will be mapped to the full input range of the display device that receives the output of the Presentation LUT.

Note: The output of the preceding transformation may be signed. This does not mean that signed P-Values actually need to be generated, only that the output of the preceding transformation is to be interpreted by the display device as perceptually linear over the range from the minimum to the maximum values output by the preceding step, and that the minimum value be mapped to the lowest JND Index (and hence luminance) that the display can generate, and the maximum value be mapped to the highest JND Index (and hence luminance) that the display can generate.

In other words, in both cases, the Presentation LUT Module is always implicitly specified to apply over the full range of output of the preceding transformation, and it never selects a subset or superset of the that range (unlike the VOI LUT).

The output bit precision of the VOI LUT Sequence is not required to match the input range of the Presentation LUT Sequence.

Notes: 1. For example, if the VOI LUT is specified as a Window Center of 0 and a Window Width of 100, then the range from -50 to +49 is selected to be mapped to the full range of the display or print device (the full range of P-Values) if the Presentation LUT Shape is specified as IDENTITY or INVERSE. This example demonstrates the conventional understanding of the meaning of Window Center and Width to select “values of interest” that are to be displayed across the full range of the output device, without explicitly having to map each choice to P-Values.

2. For example, if the VOI LUT is specified as a Window Center of 0 and a Window Width of 100, and the Presentation LUT Sequence is sent with a LUT Descriptor first value of 256 and second value of 0, then the range from -50 to +49 is implicitly linearly scaled from 0 to 255 before selecting values from the LUT Data in the Presentation LUT Sequence. This example demonstrates that it is not necessary to send a different Presentation LUT for different Window Center and Width values.

3. For example, if the VOI LUT is specified as VOI LUT Sequence with a LUT Descriptor with a 3rd Value of 16, then the range from 0 to 2^{16}-1 is selected to be mapped to the full range of the display or print device (the full range of P-Values) if the Presentation LUT Shape is specified as IDENTITY or INVERSE. This example demonstrates that a VOI LUT may be specified with the desired precision, without having to explicitly send a Presentation LUT to rescale that precision to whatever range of P-Values is preferred by the display application.

4. For example, if the VOI LUT is specified as VOI LUT Sequence with a LUT Descriptor with a 3rd Value of 16, and the Presentation LUT Sequence is sent with a LUT Descriptor first value of 4096 and second value of 0, then the range from 0 to 2^{16}-1 is implicitly linearly scaled to the range 0 to 4095 before selecting values from the LUT Data in the Presentation LUT Sequence. This example demonstrates the case where, to save space, the Presentation LUT is sent in a compact form that a display application may choose to interpolate more precisely, yet the VOI LUT output may be sent with 16 bit precision.

C.11.6.1.1 LUT Descriptor

The three values of the LUT Descriptor (0028,3002) describe the format of the LUT Data in the corresponding Data Element (0028,3006).

The first value is the number of entries in the lookup table. When the number of table entries is equal to 2^{16} then this value shall be 0.

The second value is the first implicitly scaled input value mapped, and shall always be 0. The Value Representation of the second value is always US. This implicitly scaled input value is mapped to the first entry in the LUT. There are no implicitly scaled input values less than the first value mapped. An implicitly scaled input value one greater than the first value mapped is mapped to the second entry in the LUT Data. Subsequent implicitly scaled input values are mapped to the subsequent entries in the LUT Data up to an implicitly scaled input value equal to number of entries + first value mapped - 1 which is mapped to the last entry in the LUT Data. There are no implicitly scaled input values greater than number of entries + first value mapped.

The third value specifies the number of bits for each entry in the LUT Data. The third value of the LUT Descriptor (0028,3002) shall be between 8 and 16 inclusive. The LUT Data shall be stored in a format equivalent to 8 or 16 bits allocated where the high bit is equal to bits stored - 1, where bits stored is the third value.

Note: Since the LUT Descriptor (0028,3002) Attribute is multi-valued, in an Explicit VR Transfer Syntax, only one value representation (US or SS) may be specified. Since all three values are always by definition interpreted as unsigned, the explicit VR actually used will always be US.

The LUT Data contains the LUT entry values, which are P-Values.

The output range is from 0 to 2^n-1 where n is the third value of LUT Descriptor. This range is always unsigned.

This range specifies the output range of the P-Values.

C.11.6.1.2 Presentation LUT Shape

A value of INVERSE shall mean the same as a value of IDENTITY, except that the minimum output value shall convey the meaning of the maximum available luminance, and the maximum value shall convey the minimum available luminance. In other words:

P-Value = maximum value - output value

Item: Add PS 3.3 C.11.7 Overlay/Curve Activation Module

C.11.7 Overlay/Curve Activation Module

This Module defines a manner of controlling whether or not bit-mapped overlay and curve information are displayed.

In the case of Curves, these Curves are contained within the referenced image(s).

Note: Curves may not be present within the Presentation State, since the same function is served by the Graphic Annotation Module which provides additional features.

In the case of Overlays, if the corresponding Overlay Group activated is present within the Presentation State, then that Overlay shall be activated and any corresponding Overlay in the referenced image(s) ignored, otherwise the Overlay within the referenced image(s) shall be activated.

An Overlay Group referenced in the Bitmap Display Shutter Module described in C.7.6.15 shall not be activated using the Overlay/Curve Activation Module.

Table C.11.7-1 specifies the Attributes that describe the Overlay/Curve Activation Module.

Table C.11.7-1
OVERLAY/CURVE ACTIVATION MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Overlay Activation Layer	(60xx,1001)	2C	The layer (defined in Graphic Layer (0070,0002) of the Graphic Layer Module C.10.7) in which the Overlay described in group 60xx shall be displayed. If no layer is specified (zero length) then the overlay shall not be displayed. Required if Group 60xx is present in the referenced image(s) or the Presentation State instance containing this Module.
Curve Activation Layer	(50xx,1001)	2C	The layer (defined in Graphic Layer (0070,0002) of the Graphic Layer Module C.10.7) in which the Curve described in group 50xx shall be displayed. If no layer is specified (zero length) then the curve shall not be displayed. Required if Group 50xx is present in the referenced image(s) and Type of Data (50xx,0020) is POLY or ROI. Note: Curves with other types of data are not expected to be displayed.

Note: Those bits which are stored in Pixel data (7FE0,0010) above High Bit(0028,0102) may be used as overlay bit planes if they are referenced by an Overlay Bit Position (60xx,0102). If they are not so referenced, their contents are unspecified in DICOM and should not be displayed. Usually they will be zero, though if the pixel data is signed, i.e. Pixel Representation (0028,0103) is 0001H, then it is possible that the sign bit may be “extended” through these values. Alternatively, they may have been “masked off” even if the value is signed and negative.

Item: Add PS 3.3 C.11.8 Softcopy VOI LUT Module

C.11.8 Softcopy VOI LUT module

Table C.11-7 specifies the Attributes that describe the Softcopy VOI LUT. These Attributes have the same meaning and behavior as defined in the VOI LUT Module Section C.11.2.

Table C.11.8-1
SOFTCOPY VOI LUT MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Softcopy VOI LUT Sequence	(0028,3110)	1	Defines a sequence of VOI LUTs or Window Centers and Widths and to which images and frames they apply. No more than one VOI LUT Sequence containing a single Item or one pair of Window Center/Width values shall be specified for each image or frame. One or more Items shall be present.
>Referenced Image Sequence	(0008,1140)	1C	Sequence of Repeating Items where each Item provides reference to a selected set of Image SOP Class/SOP Instance pairs that are defined in the Presentation State Module. Required if a sequence item is present, and if the VOI LUT transformation in this Item does not apply to all the images listed in the Presentation State Module.
>>Referenced SOP Class UID	(0008,1150)	1C	Uniquely identifies the referenced SOP Class. Required if sequence item is present.
>>Referenced SOP Instance UID	(0008,1155)	1C	Uniquely identifies the referenced SOP Instance. Required if sequence item is present.
>>Referenced Frame Number	(0008,1160)	1C	Identifies the frame numbers within the referenced SOP Instance to which this VOI LUT or Window Center and Width applies. Required if sequence item is present and the referenced image is a multi-frame image and the VOI LUT or Window Center and Width does not apply to all frames.
>VOI LUT Sequence	(0028,3010)	1C	Defines a sequence of VOI LUTs. Only one Item may be present. Required if Window Center (0028,1050) is not present.
>>LUT Descriptor	(0028,3002)	1C	Specifies the format of the LUT Data in this Sequence. See C.11.2.1.1 for further explanation. Required if the VOI LUT Sequence (0028,3010) is sent.

>>LUT Explanation	(0028,3003)	3	Free form text explanation of the meaning of the LUT.
>>LUT Data	(0028,3006)	1C	LUT Data in this Sequence. Required if the VOI LUT Sequence (0028,3010) is sent.
>Window Center	(0028,1050)	1C	Window Center for display. See C.11.2.1.2 for further explanation. Required if VOI LUT Sequence (0028,3010) is not present.
>Window Width	(0028,1051)	1C	Window Width for display. See C.11.2.1.2 for further explanation. Required if Window Center (0028,1050) is sent.
>Window Center & Width Explanation	(0028,1055)	3	Free form explanation of the meaning of the Window Center and Width.

Item: Add to PS 3.3 new Section C.12.2

C.12.2 Presentation Series Module

Table C.12-5 contains Attributes that identify and describe a Presentation Series.

Table C.12-5
PRESENTATION SERIES MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Modality	(0008,0060)	1	Type of equipment that originally acquired the data. Enumerated Value: PR = Presentation State See C.7.3.1.1.1.

Note: This implies that presentation states will be in different series from the images to which they apply, which will have different values for Modality.

Item: Add to PS 3.3 new Section C.12.3

C.12.3 Presentation State Module

Table C.12-6 contains Attributes that identify and describe a Presentation State.

Table C.12-6
PRESENTATION STATE MODULE ATTRIBUTES

Attribute Name	Tag	Type	Attribute Description
Instance Number	(0020,0013)	1	A number that identifies this presentation (SOP Instance). Note: In previous versions of the Standard this Attribute was referred to as Image Number.
Presentation Label	(0070,0080)	1	A label that is used to identify this presentation. Note: This value may be used by an application as a Defined Term in order to imply some grouping of different presentation states, i.e. it may have the same value for different presentation state instances that share some common concept.
Presentation Description	(0070,0081)	2	A description of this presentation.
Presentation Creation Date	(0070,0082)	1	Date on which this presentation was created. Note: This date may be different from the date that the DICOM SOP Instance was created, since the presentation state information contained may have been recorded earlier.
Presentation Creation Time	(0070,0083)	1	Time at which this presentation was created. Note: This time may be different from the time that the DICOM SOP Instance was created, since the presentation state information contained may have been recorded earlier.
Presentation Creator's Name	(0070,0084)	2	Name of operator saving the presentation state (such as a technologist or physician).
Referenced Series Sequence	(0008,1115)	1	Sequence of Repeating Items where each Item includes the Attributes of one or more Series.
>Series Instance UID	(0020,000E)	1C	Unique identifier of a Series that is part of this Study. Required if sequence item is present.

>Retrieve AE Title	(0008,0054)	3	Title of the DICOM Application Entity where the Image(s) may be retrieved on the network.
>Storage Media File-Set ID	(0088,0130)	3	The user or implementation specific human readable identifier that identifies the Storage Media on which the Image(s) reside.
>Storage Media File-Set UID	(0088,0140)	3	Uniquely identifies the Storage Media on which the Image(s) reside.
>Referenced Image Sequence	(0008,1140)	1C	Sequence of Repeating Items where each Item provides reference to a selected set of Image SOP Class/SOP Instance pairs that are part of this Study and the Series defined by Series Instance UID (0020,000E). Required if a sequence item is present.
>>Referenced SOP Class UID	(0008,1150)	1C	Uniquely identifies the referenced SOP Class. Required if sequence item is present. Shall be the same for all Images referenced by this presentation state.
>>Referenced SOP Instance UID	(0008,1155)	1C	Uniquely identifies the referenced SOP Instance. Required if sequence item is present.
>>Referenced Frame Number	(0008,1160)	1C	Identifies the frame numbers within the referenced SOP Instance to which the presentation applies. Note: This Attribute may be multi-valued. Required if sequence item is present and the referenced SOP Instance is a multi-frame image and the presentation does not apply to all frames.
Shutter Presentation Value	(0018,1622)	1C	The value used to replace those parts of the image occluded by the shutter, in P-Values. Required if the Display Shutter Module or Bitmap Display Shutter Module is present. Note: The requirement in this module is type 1C which overrides the type 3 in the Display Shutter Module.

Mask Subtraction Sequence	(0028,6100)	1C	Required if Mask Module is present. Only one Item shall be present. Applicable Frame Range (0028,6102) shall not be included in the Sequence Item. See C.7.6.10 for a complete definition of the Attributes in the Items of this Sequence other than Mask Operation(0028,6101) and Applicable Frame Range (0028,6102). Notes: 1. This Sequence is replicated here in order to specify one Item, additional conditions on Mask Operation (0028,6101) and to forbid Applicable Frame Range (0028,6102). 2. The role of Applicable Frame Range (0028,6102) is replaced by Referenced Frame Number (0008,1160).
>Mask Operation	(0028,6101)	1	Type of mask operation to be performed Enumerated Values: AVG_SUB TID See C.7.6.10.1 for further explanation. Note: The requirement in this module is for Enumerated Values which override the requirements of the Mask Module.
>Contrast Frame Averaging	(0028,6112)	1C	Specified the number of contrast frames to average together before performing the mask operation. Required if Mask Frame Numbers (0028,6110) specifies more than one frame (i.e. is multi-valued). Note: The requirement in this module is conditional and overrides the optional requirements of the Mask Module.
Recommended Viewing Mode	(0028,1090)	1C	Specifies the recommended viewing protocol(s). Enumerated Value: SUB = for subtraction with mask images Required if Mask Subtraction Sequence (0028,6100) is present. Note: The requirement in this module is type 1C and an Enumerated Value is specified which override the requirements of the Mask Module.

Item: Modify Figure F.4.1 to include the Presentation State media storage feature by adding "Presentation DR" under the Series DR.

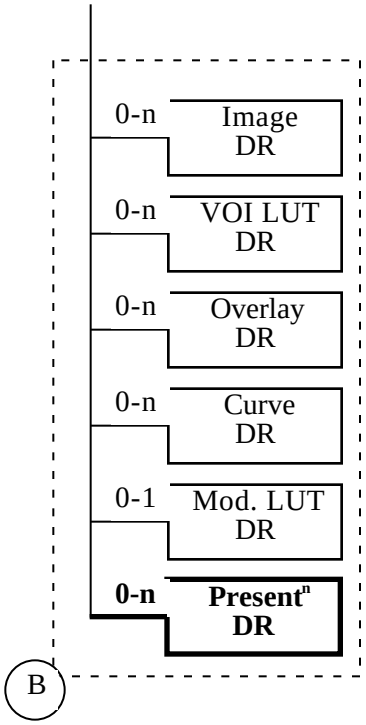

Item: Modify Table F.3-3 in section F.3.2.2 to include "PRESENTATION" in the Directory Record Type item of the Directory Record sequence.

>Directory Record Type	(0004,1430)	1C	Defines a specialized type of Directory Record by reference to its position in the Media Storage Directory Information Model (see Section F.4). Required if the Directory Record Sequence (0004,1220) is not zero length. Enumerated Values (see Section F.5): PATIENT STUDY SERIES IMAGE OVERLAY MODALITY LUT VOI LUT CURVE TOPIC VISIT RESULTS INTERPRETATION STUDY COMPONENT STORED PRINT **PRESENTATION** PRIVATE = Privately defined record hierarchy position. Type shall be defined by Private Record UID (0004,1432). MRDR = Special Directory Record which allows indirect reference to a File by multiple Directory Records. Instead of directly referencing a File by its Referenced File ID (0004,1500), a Directory Record of any of the Types define above (except MRDR) may reference a Multi-Referenced File Directory Record which in turn will reference the File by its File ID.

Modify Table F.4-1 to include PRESENTATION.

Directory Record Type	Section	Directory Record Types which may be included in the next lower-level directory Entity
(Root Directory Entity)	¾	PATIENT, TOPIC, PRIVATE
PATIENT	F.5.1	STUDY, PRIVATE
STUDY	F.5.2	SERIES, VISIT, RESULTS, STUDY COMPONENT PRIVATE
SERIES	F.5.3	IMAGE, OVERLAY, MODALITY LUT, VOI LUT, CURVE, STORED PRINT, RT DOSE, RT STRUCTURE SET, RT PLAN, RT TREAT RECORD, **PRESENTATION,** PRIVATE
IMAGE	F.5.4	PRIVATE
OVERLAY	F.5.5	PRIVATE
MODALITY LUT	F.5.6	PRIVATE
VOI LUT	F.5.7	PRIVATE
CURVE	F.5.8	PRIVATE
STORED PRINT	F.5.18	PRIVATE

RT DOSE	F.5.19	PRIVATE
RT STRUCTURE SET	F.5.20	PRIVATE
RT PLAN	F.5.21	PRIVATE
RT TREAT RECORD	F.5.22	PRIVATE
PRESENTATION	**F.5.23**	**PRIVATE**
TOPIC	F.5.9	STUDY, SERIES, IMAGE, OVERLAY, MODALITY LUT, VOI LUT, CURVE, STORED PRINT, RT DOSE, RT STRUCTURE SET, RT PLAN, RT TREAT RECORD, **PRESENTATION,** PRIVATE
VISIT	F.5.10	PRIVATE
RESULTS	F.5.11	INTERPRETATION, PRIVATE
INTERPRETATION	F.5.12	PRIVATE
STUDY COMPONENT	F.5.13	PRIVATE
PRIVATE	F.6.1	PRIVATE, (any of the above as privately defined)
MRDR	F.6.2	(Not applicable)

Add section F.5.23: Presentation State Directory Record Definition

F.5.23 Presentation State Directory Record Definition

The Directory Record is based on the specification of Section F.3. It is identified by a Directory Record Type of Value "PRESENTATION". Table F.5-23 lists the set of keys with their associated Types for such a Directory Record Type. The description of these keys may be found in the Modules related to Grayscale Softcopy Presentation State Storage IODs. This Directory Record shall be used to reference a Grayscale Softcopy Presentation State Storage SOP Instance. This Type of Directory Record may reference a Lower-Level Directory Entity which includes one or more Directory Records as defined in Table F.4-2.

Table F.5-23
PRESENTATION KEYS

Key	Tag	Type	Attribute Description
Specific Character Set	(0008,0005)	1C	Required if an extended or replacement character set is used in one of the keys.
Instance Number	(0020,0013)	1	A number that identifies this presentation state.
Presentation Label	(0070,0080)	1	A label that is used to identify this presentation.
Presentation Description	(0070,0081)	2	A description of this presentation.
Presentation Creation Date	(0070,0082)	1	Date on which this presentation was created. Note: This date may be different from the date that the DICOM SOP Instance was created, since the presentation state information contained may have been recorded earlier.

Presentation Creation Time	(0070,0083)	1	Time at which this presentation was created. Note: This time may be different from the time that the DICOM SOP Instance was created, since the presentation state information contained may have been recorded earlier.
Presentation Creator's Name	(0070,0084)	2	Name of operator saving the presentation state (such as a technologist or physician).
Referenced Series Sequence	(0008,1115)	1	Sequence of Repeating Items where each Item includes the Attributes of one or more Series.
>Series Instance UID	(0020,000E)	1C	Unique identifier of a Series that is part of this Study. Required if sequence item is present.
>Referenced Image Sequence	(0008,1140)	1C	Sequence of Repeating Items where each Item provides reference to a selected set of Image SOP Class/SOP Instance pairs that are part of this Study and the Series defined by Series Instance UID (0020,000E). Required if a sequence item is present.
>>Referenced SOP Class UID	(0008,1150)	1C	Uniquely identifies the referenced SOP Class. Required if sequence item is present. Shall be the same for all Images referenced by this presentation state.
>>Referenced SOP Instance UID	(0008,1155)	1C	Uniquely identifies the referenced SOP Instance. Required if sequence item is present.
Any other Attribute of the Presentation IE Modules		3	

Note: Because (0004,1511) Referenced SOP Instance UID in File may be used as a "pseudo" Directory Record Key (See Table F.3-3), it is not duplicated in this list of keys.

Changes to:

NEMA Standards Publication PS 3.4-1996

Digital Imaging and Communications in Medicine (DICOM)
Part 4: Service Class Specifications

Annex B STORAGE SERVICE CLASS (Normative)

Item: Add to PS 3.4 B.5

B.5 STANDARD SOP CLASSES

The SOP Classes in the Storage Service Class identify the Composite IODs to be stored. Table B.5-1 identifies Standard SOP Classes.

Table B.5-1
STANDARD SOP CLASSES

SOP Class Name	SOP Class UID	IOD Specification
...		
Grayscale Softcopy Presentation State Storage	1.2.840.10008.5.1.4.1.1.11.1	Grayscale Softcopy Presentation State Storage

Item: Add to PS 3.4 I.5

I.4 MEDIA STORAGE STANDARD SOP CLASSES

The SOP Classes in the Media Storage Service Class identify the Composite and Normalized IODs to be stored. Table I.4-1 identifies Standard SOP Classes.

Table I.4--1
MEDIA STORAGE STANDARD SOP CLASSES

SOP Class Name	SOP Class UID	IOD Specification
...		
Grayscale Softcopy Presentation State Storage	1.2.840.10008.5.1.4.1.1.11.1	Grayscale Softcopy Presentation State Storage

Item: Add Annex N to PS 3.4

Annex N GRAYSCALE SOFTCOPY PRESENTATION STATE STORAGE SOP CLASS (Normative)

N.1. OVERVIEW

N.1.1 SCOPE

The Grayscale Softcopy Presentation State Storage SOP Class extends the functionality of the Storage Service class (defined in Annex B) to add the ability to convey an intended presentation state or record an existing presentation state. The SOP Class specifies information and behavior that may be used to present (display) images that are referenced from within the SOP Class.

It includes capabilities for specifying:

a. the output grayscale space in P-Values
b. grayscale contrast transformations including modality and VOI LUT
c. mask subtraction for multi-frame images
d. selection of the area of the image to display and whether to rotate or flip it
e. image and display relative annotations, including graphics, text and overlays

The softcopy presentation state refers to the grayscale image transformations that are to be applied in an explicitly defined manner to convert the stored image pixel data values in a Composite Image Storage Instance to presentation values (P-Values) when an image is displayed on a softcopy device. The P-Values are in a device independent perceptually linear space that is formally defined in PS 3.14 Grayscale Standard Display Function.

The Grayscale Softcopy Presentation State Storage SOP Class may be used to store a single state per image, or a common state to be shared by multiple selected images. All images to which the state applies must be a part of the same study that the stored state is a part of, and be of the same Composite Image Storage SOP Class.

How an SCU of this SOP Class records or generates this state is beyond the scope of the standard.

Note: For example, an acquisition device may acquire, reconstruct and store to a workstation or archive images that are later examined by an operator for the purpose of quality assurance or printing. At that time a selected grayscale transformation (such as a window level/width operation) may be applied by the operator, and that activity captured and saved as a Grayscale Softcopy Presentation State Storage SOP Instance to the same workstation or archive, from which it is subsequently available for use by another user. Another workstation may retrieve the state for later use. Alternatively, an automated algorithm may derive a state from analysis of image statistics, body part examined, or other characteristics.

How an SCP of this SOP Class chooses between multiple states that may apply to an image is beyond the scope of this standard, other than to state that a claim of conformance as an SCP of this SOP Class implies that the SCP shall make the presentation state available to the user of the device, and if selected by the

user, shall apply all the transformations stored in the state in the manner in which they are defined in the standard.

Notes: 1. For example, an acquisition device may automatically store appropriate presentation states for series of images as they are reconstructed that represent adequate defaults. A user or algorithm may subsequently determine a more appropriate presentation state that more effectively displays the contents of an image, or record some annotation related directly to the image, and record that as another presentation state for an image. An application subsequently may display the image by automatically choosing to use the more recently saved or more specific presentation state, or may use the more general default presentation state for all images but notify the user that alternative presentation states are available.

2. Choice of the same presentation state to display an image on two devices claiming conformance to these SOP Classes implies through the definition of the P-Value space that the displayed image on both devices will be perceptually similar within the limits defined in PS 3.14 Grayscale Standard Display Function, regardless of the actual capabilities of the display systems.

N.2 GRAYSCALE TRANSFORMATION SEQUENCE

The Grayscale Softcopy Presentation State Storage SOP Class supports a sequence of grayscale transformations that completely define the conversion of a stored image into a displayed image.

The sequence of grayscale transformations from stored pixel values into the Grayscale Standard Display Function P-Values is explicitly defined in a conceptual model. The actual sequence implemented may differ but must result in the same appearance. Figure N.2-1 describes this sequence of grayscale transformations.

Notes: 1. Even though a Composite Image Storage SOP Class may not include some modules that are part of the described grayscale transformations, the Grayscale Softcopy Presentation State Storage SOP Class does include them. For example, the CT Image Storage SOP Class includes Rescale Slope and Intercept in the CT Image Module, but does not include the Modality LUT Module, and hence is restricted to the description of linear transformations. A saved presentation state that refers to a CT Image Storage SOP Instance may include a Modality LUT, and hence may apply a non-linear transformation. This is a feature of the extended functionality of the Grayscale Softcopy Presentation State Storage SOP Class.

2. For the shutter, annotation and spatial transformations, the order in which they are applied relative to the other transformations should not result in a different appearance. The one exception is when a spatial transformation is applied that involves magnification implemented with interpolation. In this case, whether the interpolation is performed before or after the contrast transformations (such as VOI LUT) may result in a slightly different appearance. It is not considered necessary to constrain this sequence more precisely.

The grayscale transformations defined in the Grayscale Softcopy Presentation State Storage SOP Class replace those that may be defined in the Referenced Image SOP Instance. If a particular transformation is absent in the Grayscale Softcopy Presentation State Storage SOP Class, then it shall be assumed to be an identity transformation, and any equivalent transformation, if present, in the Referenced Image SOP Instance shall NOT be used instead.

Photometric Interpretation (0028,0004) in the Referenced Image SOP Instance shall be ignored, since its effect is defined by the application of the grayscale transformations.

Note: These requirements are in order to achieve complete definition of the entire grayscale transformation in the Grayscale Softcopy Presentation State Storage SOP Class, and not to depend on the content of the Referenced Image SOP Instance, which may change.

The Referenced Image Storage SOP Instance may also contain bit-mapped overlays and curves. The Grayscale Softcopy Presentation State Storage SOP Class specifies a mechanism for turning these on or off (i.e. displaying them or not).

The presentation related Attributes of the Grayscale Softcopy Presentation State Storage SOP Class are immutable. They shall never be modified or updated; only a derived SOP Instance with a new SOP Instance UID may be created to represent a different presentation.

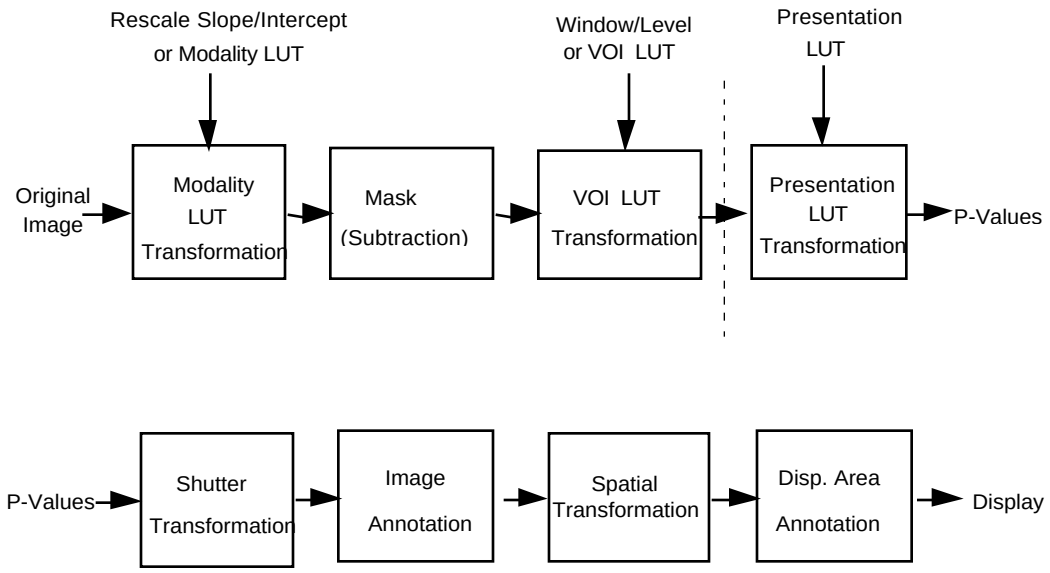

Figure N.2-1
Grayscale Image Transformation Model

N.2.1 Modality LUT

The Modality LUT transformation transforms the manufacturer dependent pixel values into pixel values which are meaningful for the modality and which are manufacturer independent (e.g., Hounsfield number for CT modalities, Optical Density for film digitizers). These may represent physical units or be dimensionless. The Modality LUT in the Presentation State is modality dependent and is analogous to the same module in an Image .

Note: In some cases, such as the CT Image Storage SOP Class, the same conceptual step as the Modality LUT is specified in another form, for example as Rescale Slope and Rescale Intercept Attributes in the CT Image Module, though the Modality LUT Module is not part of the CT Image IOD.

In the case of a linear transformation, the Modality LUT is described by the Rescale Slope (0028,1053) and Rescale Intercept (0028,1052). In the case of a non-linear transformation, the Modality LUT is described by the Modality LUT Sequence. The rules for application of the Modality LUT are defined in PS 3.3 Modality LUT Module.

If the Modality LUT or equivalent Attributes are part of both the Image and the Presentation State, then the Presentation State Modality LUT shall be used instead of the Image Modality LUT or equivalent Attributes

in the Image. If the Modality LUT is not present in the Presentation State it shall be assumed to be an identity transformation. Any Modality LUT or equivalent Attributes in the Image shall not be used.

N.2.2 Mask

The mask transformation may be applied in the case of multi-frame images for which other frames at a fixed frame position or time interval relative to the current frame may be subtracted from the current frame. Multiple mask frames may be averaged, and sub-pixel shifted before subtraction.

This transformation uses the Mask Module as used in the X-Ray Angiography Image Storage SOP Class, though it may be applied to any Image Storage SOP Instance that contains a multi-frame image.

In the case of X-Ray images, the subtraction is specified to take place in a space logarithmic to X-Ray intensity. If the stored pixel values are not already in such a space, an implementation defined transformation to such a space must be performed prior to subtraction. If a Modality LUT Module is present as well as a Mask Module, then the Modality LUT shall specify a transformation into such a logarithmic space, otherwise it shall not be present (even though a Modality LUT may be present in the referenced image(s) which shall be ignored).

Notes: 1. In the case of an XA or XRF image, if the Pixel Intensity Relationship (0028,1040) in the image is LOG, then even though a Modality LUT would be present in the image (to map pixel values back to linear to X-Ray intensity), no Modality LUT would be present in the presentation state (i.e. the Modality LUT would be an identity transformation) since log values are required for subtraction. See PS 3.3 C.8.7.1.1.2.

2. In the case of an XA or XRF image, if the Pixel Intensity Relationship (0028,1040) is LIN, then no Modality LUT would be present in the image, but a Modality LUT would need to be present in the presentation state since log values are required for subtraction.

3. In the case of an XA or XRF image, if the Pixel Intensity Relationship (0028,1040) in the image is DISP, then even though a Modality LUT may or may not be present in the image (to map pixel values back to linear to X-Ray intensity), a different Modality LUT would be present in the presentation state if the creator of the presentation state could create a transformation from DISP pixel values to a logarithmic space for subtraction, or the Modality LUT in the presentation state would be an identity transformation if the DISP pixel values were known to already be log values required for subtraction.

The result will be a signed value with a bit length one longer than the source frames.

When there is no difference between corresponding pixel values, the subtracted image pixel will have a value of 0.

If a pixel in the current frame has a greater value than in the mask frame, then the resulting frame shall have a positive value. If it has a lesser value, then the resulting frame shall have a negative value.

N.2.3 VOI LUT

The value of interest (VOI) LUT transformation transforms the modality pixel values into pixel values which are meaningful for the user or the application.

Note: Photometric Interpretation (0028,0004) is ignored, since its effect is defined by the application of the grayscale transformations.

The Softcopy VOI LUT Module in the Presentation State is analogous to the VOI LUT Module in an Image.

In the case of a linear transformation, the VOI LUT is described by the Window Center (0028,1050) and Window Width (0028,1051). In the case of a non-linear transformation, the VOI LUT is described by the VOI LUT Sequence. The rules for application of the VOI LUT are defined in PS 3.3 Softcopy VOI LUT Module.

The VOI LUT may have sections with negative slope.

Note: In the Basic Print Service Class a VOI LUT may not have negative slope.

If a VOI LUT is part of both the Image and the Presentation State then the Presentation State VOI LUT shall be used instead of the Image VOI LUT. If a VOI LUT (that applies to the Image) is not present in the Presentation State , it shall be assumed to be an identity transformation. Any VOI LUT or equivalent values in the Image shall not be used.

N.2.4 Presentation LUT

The Presentation LUT transformation transforms the pixel values into P-Values, a device independent perceptually linear space as defined in PS 3.14 Grayscale Display Function Standard. It may be an identity function if the output of the VOI LUT transformation is in P-Values.

Note: If the Presentation LUT and VOI LUT step are identity transformations, and the Mask Module is absent, then the output of the Modality LUT must be, by definition, P-Values.

No output space other than P-Values is defined for the Grayscale Softcopy Presentation State Storage SOP Classes.

In the case of a linear transformation, the Presentation LUT is described by the Presentation LUT Shape (2050,0020). In the case of a non-linear transformation, the Presentation LUT is described by the Presentation LUT Sequence. The rules for application of the Presentation LUT are defined in PS 3.3 Softcopy Presentation LUT Module.

Notes: 1. Since the grayscale transformation pipeline fully defines all transformations applied to the stored pixel values in the referenced image object, the value of Photometric Interpretation (0028,0004) in the referenced image object is ignored and overridden. This implies that either the creator of the presentation state chose a pipeline that reflects the Photometric Interpretation (0028,0004) , or chose to ignore or override the Photometric Interpretation, and invert the image relative to what is specified by Photometric Interpretation. If the Modality LUT and VOI LUT do not have a negative slope, one can achieve the effect of inversion of the polarity of an image by choosing Presentation LUT Shape of IDENTITY or INVERSE that displays the minimum pixel value as white rather than black in the case of a Photometric Interpretation of MONOCHROME2, or black rather than white in the case of a Photometric Interpretation of MONOCHROME1. If Presentation LUT Data is sent, then one can invert the order of the entries in the LUT table to achieve inversion of polarity.

2. The minimum P-Value (zero) always commands that the lowest intensity be displayed.

3. No separate Polarity transformation is defined.

A Softcopy Presentation LUT Module is always present in a Presentation State. If a Presentation LUT is present in the Image then the Presentation State Presentation LUT shall be used instead of the Image Presentation LUT.

N.2.5 Shutter

The Shutter transformation provides the ability to exclude the perimeter outside a region of an image. A gray level may be specified to replace the area under the shutter.

One form of this transformation uses the Display Shutter Module as used in the X-Ray Angiography Image Storage SOP Class, though it may be applied to any Image Storage SOP Instance, including single frame images.

Another form uses a bit-mapped overlay to indicate arbitrary areas of the image that should be excluded from display by replacement with a specified gray level, as described in the Bitmap Display Shutter Module.

Notes: 1. Since annotations follow the shutter operation in the pipeline, annotations in shuttered regions are not obscured and are visible.

2. Any shutter present in the referenced image object is ignored (i.e. not applied).

N.2.6 Pre-Spatial Transformation Annotation

The Pre-Spatial Transformation Annotation transformation includes the application of bit-mapped overlays as defined in the Overlay Plane Module, and free unformatted text or vector graphics as described in the Graphic Annotation Module that are defined in the image pixel space (as opposed to the displayed area space).

N.2.7 Spatial Transformation

Some modalities may not deliver the image in the desired rotation and need to specify a rotation into the desired position for presentation. This transformation, specified in the Spatial Transformation Module, includes a rotation of 90, 180, 270 degrees clockwise followed by a horizontal flip (L <--> R). Rotation by an arbitrary angle is not supported.

In addition, selection of a region of the image pixel space to be displayed is specified in the Displayed Area Module. This may have the effect of magnifying (or minifying) that region depending on what physical size the display is instructed to render the selected region. If so, the method of interpolation (or sub-sampling) is implementation dependent.

Note: In particular the number of displayed pixels may be different from the number of image pixels as a result of:

- minification (e.g. 1 display pixel for 4 image pixels),
- magnification (4 display pixels for each image pixel) ,
- interpolation (display pixels derived from values other than those in the image pixels), and
- sub-sampling.

N.2.8 Post-Spatial Transformation Annotation

The Post-Spatial Transformation Annotation transformation includes the application of free unformatted text or vector graphics as described in the Graphic Annotation Module that are defined in the displayed area space (as opposed to the image pixel space).

This implies that the displayed area space is defined as being the image after all Spatial Transformations have been applied.

These annotations are rendered in the displayed space, though they may be anchored to points in either the displayed area or image pixel space.

N.3 BEHAVIOR OF AN SCP

In addition to the behavior for the Storage Service Class specified in B.2.2 Behavior of an SCP, the following additional requirements are specified for the Grayscale Softcopy Presentation State Storage SOP Class:

— a display device acting as an SCP of this SOP Class shall make all mandatory presentation attributes available for application to the referenced images at the discretion of the display device user, for all Image Storage SOP Classes defined in the Conformance Statement for which the Grayscale Softcopy Presentation State Storage SOP Class is supported.

N.4 CONFORMANCE

In addition to the Conformance Statement requirements for the Storage Service Class specified in B.4.3, the following additional requirements are specified for the Grayscale Softcopy Presentation State Storage SOP Class:

N.4.1 Conformance Statement for An SCU

The following issues shall be documented in the Conformance Statement of any implementation claiming conformance to the Grayscale Softcopy Presentation State Storage SOP Class as an SCU:

— For an SCU of a Grayscale Softcopy Presentation State Storage SOP Class that is creating a SOP Instance of the Class, the manner in which presentation related attributes are derived from a displayed image, operator intervention or defaults, and how they are included in the IOD.

— For an SCU of a Grayscale Softcopy Presentation State Storage SOP Class, the Image Storage SOP Classes that are also supported by the SCU and which may be referenced by instances of the Grayscale Softcopy Presentation State Storage SOP Class.

N.4.2 Conformance Statement for An SCP

The following issues shall be documented in the Conformance Statement of any implementation claiming conformance to the Grayscale Softcopy Presentation State Storage SOP Class as an SCP:

— For an SCP of a Grayscale Softcopy Presentation State Storage SOP Class that is displaying an image referred to by a SOP Instance of the Class, the manner in which presentation related attributes are used to influence the display of an image.

— For an SCP of a Grayscale Softcopy Presentation State Storage SOP Class, the Image Storage SOP Classes that are also supported by the SCP and which may be referenced by instances of the Grayscale Softcopy Presentation State Storage SOP Class.

Digital Imaging and Communications in Medicine

PART 5 Addendum

Grayscale Softcopy Presentation State Encoding

Part 5 Addendum : Grayscale Softcopy Presentation State Encoding

Amend PS 3.5-1998 Section 7.3 to include note about packing bits within bytes for pixel & overlay data.

7.3 BIG ENDIAN VERSUS LITTLE ENDIAN BYTE ORDERING

Another component of the encoding of a Data Set that shall be agreed upon by communicating Application Entities is the Byte Ordering.

Little Endian byte ordering is defined as follows:

- ¾ In a binary number consisting of multiple bytes (e.g. a 32-bit unsigned integer value, the Group Number, the Element Number, etc.), the least significant byte shall be encoded first; with the remaining bytes encoded in increasing order of significance.
- ¾ In a character string consisting of multiple 8-bit single byte codes, the characters will be encoded in the order of occurrence in the string (left to right).

Big Endian byte ordering is defined as follows:

- ¾ In a binary number consisting of multiple bytes, the most significant byte shall be encoded first; with the remaining bytes encoded in decreasing order of significance.
- ¾ In a character string consisting of multiple 8-bit single byte codes, the characters will be encoded in the order of occurrence in the string (left to right).

Note: The packing of bits within values of OB or OW Value representation for Pixel Data and Overlay Data is described in Section 8.

Byte ordering is a component of an agreed upon Transfer Syntax (see Section 10). The default DICOM Transfer Syntax, which shall be supported by all AEs, uses Little Endian encoding and is specified in Annex A.1. Alternate Transfer Syntaxes, some of which use Big Endian encoding, are also specified in Annex A.

Note: The Command Set structure as specified in PS 3.7 is encoded using the Little Endian Implicit VR Transfer Syntax.

...

Amend PS 3.5-1998 Section 8 to include note about packing bits within bytes for pixel & overlay data.

Section 8 Encoding of Pixel *and Overlay* Data

The Pixel Data Element (7FE0,0010) **and Overlay Data Element (60xx,3000**) shall be used for the exchange of encoded graphical image data. ~~**This element**~~ **These elements** along with additional Data Elements, specified as Attributes of the Image Information Entities defined in PS 3.3, shall be used to describe the way in which the Pixel Data **and Overlay Data** ~~**is**~~ **are** encoded and shall be interpreted. Finally, depending on the negotiated Transfer Syntax (see Section 10 and Annex A), Pixel Data may be compressed.

The Pixel Data Element (7FE0,0010) **and Overlay Data Element (60xx,3000**) ~~**has**~~ **have** a VR of OW or OB, depending on the negotiated Transfer Syntax (see Annex A). The only difference between OW and OB being that OB, a string of bytes, shall be unaffected by Byte Ordering (see Section 7.3).

8.1 PIXEL AND OVERLAY DATA ENCODING OF RELATED DATA ELEMENTS

8.1.1 Pixel data encoding of related data elements

Encoded Pixel Data of various bit depths shall be accommodated. The following three Data Elements shall define the Pixel structure:

- Bits Allocated (0028,0100)
- Bits Stored (0028,0101)
- High Bit (0028,0102)

Each Pixel Cell shall contain a single Pixel Sample Value. The size of the Pixel Cell shall be specified by Bits Allocated (0028,0100). Bits Stored (0028,0101) defines the total number of these allocated bits that will be used to represent a Pixel Sample Value. Bits Stored (0028,0101) shall never be larger than Bits Allocated (0028,0100). High Bit (0028,0102) specifies where the high order bit of the Bits Stored (0028,0101) is to be placed with respect to the Bits Allocated (0028,0100) specification. Bits not used for Pixel Sample Values can be used for overlay planes described further in PS 3.3.

Note: For example, in Pixel Data with 16 bits (2 bytes) allocated, 12 bits stored, and bit 15 specified as the high bit, one pixel sample is encoded in each 16-bit word, with the 4 least significant bits of each word not containing Pixel Data. See Annex D for other examples of the basic encoding schemes.

Starting with DICOM Version 3.0, restrictions are placed on acceptable Values for Bits Allocated (0028,0100), Bits Stored (0028,0101), and High Bit (0028,0102) and are specified in the Information Object Definitions in PS 3.3. Also, the Value Field containing Pixel Data, like all other Value Fields in DICOM, shall be an even number of bytes in length. This means that the Value Field may need to be padded with data that is not part of the image and shall not be considered significant. If needed, the padding bits shall be appended to the end of the Value Field.

The field of bits representing the value of a Pixel Sample shall be a binary 2's complement integer or an unsigned integer, as specified by the Data Element Pixel Representation (0028,0103). The sign bit shall be the High Bit in a Pixel Sample Value that is a 2's complement integer. The minimum actual Pixel Sample Value encountered in the Pixel Data is specified by Smallest Image Pixel Value (0028,0106) while the maximum value is specified by Largest Image Pixel Value (0028,0107).

8.1.2 Overlay data encoding of related data elements

Encoded Overlay Planes always have a bit depth of 1, but may be encoded in bits not used for Pixel Sample Values in the Pixel Data (7FE0,0010), or separate from the Pixel Data in Overlay Data (60xx,3000). The following two Data Elements shall define the Overlay Plane structure:

- **Overlay Bits Allocated (60xx,0100)**
- **Overlay Bit Position (60xx,0102)**

Notes: 1. There is no Data Element analogous to Bits Stored (0028,0101) since Overlay Planes always have a bit depth of 1.

2. Restrictions on the allowed values for these Data Elements are defined in PS 3.3.

3. For example, in Pixel Data with 16 bits (2 bytes) allocated, 12 bits stored, and bit 11 specified as the high bit, one pixel sample is encoded in each 16-bit word, with the 4 most significant bits of each word not containing Pixel Data. These 4 most significant bits can be used to store Overlay Planes. For example, a single plane can be stored in bit 15 by specifying 15 for Overlay Bit Position. Overlay Bits Allocated would be 16, since it is always equal to Bits Allocated for the case of overlays embedded in the Pixel Data, as defined in PS 3.3. See Annex D for other examples of the basic encoding schemes.

If Overlay Planes are sent in the Overlay Data Element (60xx,3000), the Value Representation OW is most often required. The Value Representation OB may also be used for Overlay Data in cases where the Value Representation is explicitly conveyed (see Annex A).

Note: The DICOM default Transfer Syntax (Implicit VR Little Endian) does not explicitly convey Value Representation and therefore the VR of OB may not be used for Pixel Data when using the default Transfer Syntax.

Overlay Data is encoded as the direct concatenation of the bits of a single Overlay Plane, where the first bit of an Overlay Plane is encoded in the least significant bit, immediately followed by the next bit of the Overlay Plane in the next most significant bit. When the Overlay Data crosses a word boundary in the OW case, or a byte boundary in the OB case, it shall continue to be encoded, least significant bit to most significant bit, in the next word, or byte, respectively (see Annex D). For Pixel Data encoded with the Value Representation OW, the byte ordering of the resulting 2-byte words is defined by the Little Endian or Big Endian Transfer Syntaxes negotiated at the Association Establishment (see Annex A).

Note: For Overlay Data encoded with the Value Representation OB, the Overlay Data encoding is unaffected by Little Endian or Big Endian byte ordering.

8.2 NATIVE OR ENCAPSULATED FORMAT ENCODING

Pixel data conveyed in the Pixel Data Element (7FE0,0010) may be sent either in a Native (uncompressed) Format or in an Encapsulated Format (e.g. compressed) defined outside the DICOM standard.

If Pixel Data is sent in a Native Format, the Value Representation OW is most often required. The Value Representation OB may also be used for Pixel Data in cases where Bits Allocated has a value less than or equal to 8, but only with Transfer Syntaxes where the Value Representation is explicitly conveyed (see Annex A).

Note: The DICOM default Transfer Syntax (Implicit VR Little Endian) does not explicitly convey Value Representation and therefore the VR of OB may not be used for Pixel Data when using the default Transfer Syntax.

Native format Pixel Cells are encoded as the direct concatenation of the bits of each Pixel Cell, where the most significant bit of a Pixel Cell is immediately followed by the least significant bit of the next Pixel Cell. The number of bits of each Pixel Cell is defined by the Bits Allocated (0028,0100) Data Element Value. When a Pixel Cell crosses a word boundary in the OW case, or a byte boundary in the OB case, it shall continue to be encoded, least significant bit to most significant bit, in the next word, or byte, respectively (see Annex D). For Pixel Data encoded with the Value Representation OW, the byte ordering of the resulting 2-byte words is defined by the Little Endian or Big Endian Transfer Syntaxes negotiated at the Association Establishment (see Annex A).

Notes: 1. For Pixel Data encoded with the Value Representation OB, the Pixel Data encoding is unaffected by Little Endian or Big Endian byte ordering.

2. If encoding Pixel Data with a Value for Bits Allocated (0028,0100) not equal to 16 be sure to read and understand Annex D.

If sent in an Encapsulated Format (i.e. other than the Native Format) the Value Representation OB is used. The Pixel Cells are encoded according to the encoding process defined by one of the negotiated Transfer Syntaxes (see Annex A). The encapsulated pixel stream of encoded pixel data is segmented in one or more Fragments which convey their explicit length. The sequence of Fragments of the encapsulated pixel stream is terminated by a delimiter, thus allowing the support of encoding processes where the resulting length of the entire pixel stream is not known until it is entirely encoded. This Encapsulated Format supports both Single-Frame and Multi-Frame images (as defined in PS 3.3).

8.2.1 JPEG Image Compression

...

8.2.2 Run Length Encoding Compression

...

Amend PS 3.5-1998 Section A to include notes to clarify LUT Data and Overlay Data VRs.

A.1 DICOM IMPLICIT VR LITTLE ENDIAN TRANSFER SYNTAX

— **Data Element (0028,3006) Lookup Table Data has the Value Representation US, SS or OW and shall be encoded in Little Endian.**

Note: Previous versions of the Standard did not specify the encoding of these Data Elements in this Part, but specified a VR of US or SS in PS 3.6 (1998). A VR of OW has been added to support explicit VR transfer syntaxes. The actual encoding of the values and their byte order would be identical in each case.

— **Data Element (0028,3002) Lookup Table Descriptor has the Value Representation SS or US (depending on rules specified in the IOD in PS 3.3), and shall be encoded in Little Endian. The first and third values are always interpreted as unsigned, regardless of the Value Representation.**

A.2 DICOM LITTLE ENDIAN TRANSFER SYNTAX (EXPLICIT VR)

¾ Data Element (60xx,3000) Overlay Data

— ~~**where Bits Allocated (60xx,0100) has a value greater than 8 shall have Value Representation OW and shall be encoded in Little Endian;**~~

¾ ~~**where Bits Allocated (60xx,0100) has a value less than or equal to 8**~~ shall have the Value Representation OB or OW and shall be encoded in Little Endian.

Note: Previous versions of the Standard specified that the choice of OB or OW VR was based on whether or not Overlay Bits Allocated (60xx,0100) was greater than, or less than or equal to, 8. However, since only one bit plane can be encoded in each Overlay Data Element (60xx,3000), no value of Overlay Bits Allocated other than 1 makes sense. Such a restriction is now present in PS 3.3.

...

— **Data Element (0028,3006) Lookup Table Data has the Value Representation US, SS or OW and shall be encoded in Little Endian.**

Note: Previous versions of the Standard did not specify the encoding of these Data Elements in this Part, but specified a VR of US or SS in PS 3.6 (1998). However, an explicit VR of US or SS cannot be used to encode a table of 2^{16} elements, since the Value Length is restricted to 16 bits. Hence a VR of OW has been added. The actual encoding of the values and their byte order would be identical in each case, though the explicitly encoded VR field would be different.

— **Data Element (0028,3002) Lookup Table Descriptor has the Value Representation SS or US (depending on rules specified in the IOD in PS**

3.3), and shall be encoded in Little Endian. The first and third values are always interpreted as unsigned, regardless of the Value Representation.

A.3 DICOM BIG ENDIAN TRANSFER SYNTAX (EXPLICIT VR)

¾ Data Element (60xx,3000) Overlay Data

— ~~**where Bits Allocated (60xx,0100) has a value greater than 8 shall have Value Representation OW and shall be encoded in Little Endian;**~~

¾ ~~**where Bits Allocated (60xx,0100) has a value less than or equal to 8**~~ shall have the Value Representation OB or OW and shall be encoded in Big Endian.

Note: Previous versions of the Standard specified that the choice of OB or OW VR was based on whether or not Overlay Bits Allocated (60xx,0100) was greater than, or less than or equal to, 8. However, since only one bit plane can be encoded in each Overlay Data Element (60xx,3000), no value of Overlay Bits Allocated other than 1 makes sense. Such a restriction is now present in PS 3.3.

…

— **Data Element (0028,3006) Lookup Table Data has the Value Representation US, SS or OW and shall be encoded in Big Endian.**

Note: Previous versions of the Standard did not specify the encoding of these Data Elements in this Part, but specified a VR of US or SS in PS 3.6 (1998). However, an explicit VR of US or SS cannot be used to encode a table of 2^{16} elements, since the Value Length is restricted to 16 bits. Hence a VR of OW has been added. The actual encoding of the values and their byte order would be identical in each case, though the explicitly encoded VR field would be different.

— **Data Element (0028,3002) Lookup Table Descriptor has the Value Representation SS or US (depending on rules specified in the IOD in PS 3.3), and shall be encoded in Big Endian. The first and third values are always interpreted as unsigned, regardless of the Value Representation.**

A.4 TRANSFER SYNTAXES FOR ENCAPSULATION OF ENCODED PIXEL DATA

¾ Data Element (60xx,3000) Overlay Data

— ~~**where Bits Allocated (60xx,0100) has a value greater than 8 shall have Value Representation OW and shall be encoded in Little Endian;**~~

¾ ~~**where Bits Allocated (60xx,0100) has a value less than or equal to 8**~~ shall have the Value Representation OB or OW and shall be encoded in Little Endian.

Note: Previous versions of the Standard specified that the choice of OB or OW VR was based on whether or not Overlay Bits Allocated (60xx,0100) was greater than, or less than or equal to, 8. However, since only one bit plane can be encoded in each Overlay Data Element (60xx,3000), no value of Overlay Bits Allocated other than 1 makes sense. Such a restriction is now present in PS 3.3.

...

— **Data Element (0028,3006) Lookup Table Data has the Value Representation US, SS or OW and shall be encoded in Little Endian.**

Note: Previous versions of the Standard did not specify the encoding of these Data Elements in this Part, but specified a VR of US or SS in PS 3.6 (1998). However, an explicit VR of US or SS cannot be used to encode a table of 2^{16} elements, since the Value Length is restricted to 16 bits. Hence a VR of OW has been added. The actual encoding of the values and their byte order would be identical in each case, though the explicitly encoded VR field would be different.

— **Data Element (0028,3002) Lookup Table Descriptor has the Value Representation SS or US (depending on rules specified in the IOD in PS 3.3), and shall be encoded in Little Endian. The first and third values are always interpreted as unsigned, regardless of the Value Representation.**

Amend PS 3.5-1998 Section D to include notes to clarify Overlay Data encoding.

D.2 VARIOUS ADDITIONAL EXAMPLES OF PIXEL AND OVERLAY DATA CELLS

...

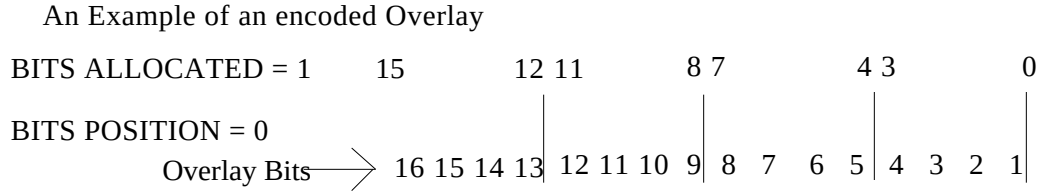

Figure D.2-4: Example 4 of ~~Pixel and~~ Overlay Data Cells

Note: In this example, the Overlay Bits are numbered in the same manner that Pixel Cells are numbered in the other examples in this Annex. That is Overlay Bit 1 is the first bit of the Overlay Plane, encoded from left to right and top to bottom, a row at a time.

Digital Imaging and Communications in Medicine

PART 6 Addendum

Grayscale Softcopy Presentation State Data Dictionary

Part 6 Addendum : Grayscale Softcopy Presentation State Data Dictionary

Section 6 : Amend the following entries in the table as indicated:

Tag	Name	VR	VM
(0028,3002)	LUT Descriptor	US~~\US~~ or SS~~\US~~	3
(0028,3006)	LUT Data	US or SS **or OW**	1-n **1**
(60xx,3000)	Overlay Data	**OB or** OW	1

Section 6 : Append the following entries to the table:

Tag	Name	VR	VM
(0018,1622)	Shutter Presentation Value	US	1
(0018,1623)	Shutter Overlay Group	US	1
(0028,3110)	Softcopy VOI LUT Sequence	SQ	1
(50xx,1001)	Curve Activation Layer	CS	1
(60xx,1001)	Overlay Activation Layer	CS	1
(0070,0001)	Graphic Annotation Sequence	SQ	1
(0070,0002)	Graphic Layer	CS	1
(0070,0003)	Bounding Box Annotation Units	CS	1
(0070,0004)	Anchor Point Annotation Units	CS	1
(0070,0005)	Graphic Annotation Units	CS	1
(0070,0006)	Unformatted Text Value	ST	1
(0070,0008)	Text Object Sequence	SQ	1

(0070,0009)	Graphic Object Sequence	SQ	1
(0070,0010)	Bounding Box Top Left Hand Corner	FL	2
(0070,0011)	Bounding Box Bottom Right Hand Corner	FL	2
(0070,0012)	Bounding Box Text Horizontal Justification	CS	1
(0070,0014)	Anchor Point	FL	2
(0070,0015)	Anchor Point Visibility	CS	1
(0070,0020)	Graphic Dimensions	US	1
(0070,0021)	Number of Graphic Points	US	1
(0070,0022)	Graphic Data	FL	2-n
(0070,0023)	Graphic Type	CS	1
(0070,0024)	Graphic Filled	CS	1
(0070,0041)	Image Horizontal Flip	CS	1
(0070,0042)	Image Rotation	US	1
(0070,0052)	Displayed Area Top Left Hand Corner	SL	2
(0070,0053)	Displayed Area Bottom Right Hand Corner	SL	2
(0070,005A)	Displayed Area Selection Sequence	SQ	1
(0070,0060)	Graphic Layer Sequence	SQ	1
(0070,0062)	Graphic Layer Order	IS	1
(0070,0066)	Graphic Layer Recommended Display Grayscale Value	US	1
(0070,0067)	Graphic Layer Recommended Display RGB Value	US	3
(0070,0068)	Graphic Layer Description	LO	1
(0070,0080)	Presentation Label	CS	1
(0070,0081)	Presentation Description	LO	1
(0070,0082)	Presentation Creation Date	DA	1
(0070,0083)	Presentation Creation Time	TM	1
(0070,0084)	Presentation Creator’s Name	PN	1
(0070,0100)	Presentation Size Mode	CS	1
(0070,0101)	Presentation Pixel Spacing	DS	2
(0070,0102)	Presentation Pixel Aspect Ratio	IS	2
(0070,0103)	Presentation Pixel Magnification Ratio	FL	1

Annex A : append the following entries to the table

UID Value	UID Name	UID Type	Part
1.2.840.10008.5.1.4.1.1.11.1	Grayscale Softcopy Presentation State Storage SOP Class	SOP Class	PS 3.4

Index of Attribute Tags and UIDs

(0004,1220) 52
(0004,1432) 52
(0004,1500) 52
(0004,1511) 54
(0008,0005) 27, 53
(0008,0054) 48
(0008,0060) 14, 46
(0008,1115) 47, 54
(0008,1140) 21, 25, 45, 48, 54
(0008,1150) 21, 25, 45, 48, 54
(0008,1155) 21, 25, 45, 48, 54
(0008,1160) 21, 25, 45, 48, 49
(0018,1164) 22, 23
(0018,1600) 17, 19
(0018,1602) 17
(0018,1604) 17
(0018,1606) 17
(0018,1608) 17
(0018,1610) 17
(0018,1612) 17
(0018,1620) 17
(0018,1622) 16, 18, 19, 48
(0018,1623) 19
(0020,000E) 21, 25, 45, 47, 48, 54
(0020,0013) 47, 53
(0028,0004) 39, 40, 58, 60, 61

(0028,0010)	18
(0028,0011)	18
(0028,0030)	22, 23
(0028,0100)	19, 67, 69
(0028,0101)	67, 68
(0028,0102)	44, 67
(0028,0103)	36, 38, 44, 67
(0028,0106)	67
(0028,0107)	67
(0028,1040)	60
(0028,1050)	8, 37, 38, 39, 40, 46, 61
(0028,1051)	8, 37, 38, 39, 40, 46, 61
(0028,1052)	8, 35, 59
(0028,1053)	35, 59
(0028,1054)	35
(0028,1055)	37, 46
(0028,1090)	15, 49
(0028,3000)	8, 35
(0028,3002)	35, 37, 38, 41, 42, 43, 45, 70, 71, 72
(0028,3003)	35, 37, 41, 46
(0028,3004)	35
(0028,3006)	35, 37, 41, 42, 46, 70, 71, 72
(0028,3010)	8, 37, 41, 45, 46
(0028,3110)	45
(0028,6100)	14, 49
(0028,6101)	14, 15, 16, 49
(0028,6102)	15, 16, 49
(0028,6110)	15, 16, 49
(0028,6112)	15, 16, 49

(0028,6114)	15, 16
(0028,6120)	15, 16
(0028,6190)	15
(0070,0001)	25
(0070,0002)	25, 33, 44
(0070,0003)	26, 27, 28
(0070,0004)	26, 28
(0070,0005)	28, 30
(0070,0006)	27, 28
(0070,0008)	26, 28
(0070,0009)	26, 28
(0070,0010)	26, 27, 28
(0070,0011)	26, 27, 28
(0070,0012)	27
(0070,0014)	26, 27, 28, 30
(0070,0015)	28, 30
(0070,0020)	28
(0070,0021)	28
(0070,0022)	28, 29, 30
(0070,0023)	29, 30
(0070,0024)	29, 30
(0070,0041)	32
(0070,0042)	32
(0070,0052)	21
(0070,0053)	21
(0070,005A)	21
(0070,0060)	33
(0070,0062)	32, 33
(0070,0066)	33

(0070,0068)	33
(0070,0080)	47, 53
(0070,0081)	47, 53
(0070,0082)	47, 53
(0070,0083)	47, 54
(0070,0084)	47, 54
(0070,0100)	20, 22, 23
(0070,0101)	20, 22
(0070,0102)	20, 23
(0070,0103)	23
(0070,0104)	20
(0088,0130)	48
(0088,0140)	48
(2050,0010)	41
(2050,0020)	41, 61
(50xx,1001)	44
(60xx,0010)	18, 19
(60xx,0011)	18
(60xx,0040)	18
(60xx,0050)	18
(60xx,0100)	18, 19, 68, 70, 71
(60xx,0102)	18, 19, 44, 68
(60xx,1001)	44
(60xx,1303)	20
(60xx,3000)	18, 19, 67, 68, 70, 71
(7FE0,0010)	19, 44, 67, 68
1.2.840.10008.5.1.4.1.1.11.1	56