



DPC: 25/30470379 DC

BSI Group Headquarters

389 Chiswick High Road London W4 4AL

Tel: +44 (0)20 8996 9000

Fax: +44 (0)20 8996 7400

www.bsigroup.com

Date: 08 January 2025

Origin: International

Latest date for receipt of comments: 03 March 2025

Project No. 2023/00420

Responsible committee: CPW/42 Photography

Interested committees:

Title: Draft BS ISO 12234-4 Digital imaging - Image storage. Part 4: Digital negative format.

Please notify admin.start@bsigroup.com if you are aware of any keywords that might assist in clarifying or identifying the standard or of the content of this standard:

- i) has any issues related to 3rd party IPR, patent or copyright
- ii) affects other national standard(s)
- iii) requires additional national guidance or information

WARNING: THIS IS A DRAFT AND MUST NOT BE REGARDED OR USED AS A BRITISH STANDARD. THIS DRAFT IS NOT CURRENT BEYOND 03 March 2025

This draft is issued to allow comments from interested parties; all comments will be given consideration prior to publication. No acknowledgement will normally be sent. **See overleaf for information on the submission of comments.**

No copying is allowed, in any form, without prior written permission from BSI except as permitted under the Copyright, Designs and Patent Act 1988 or for circulation within a nominating organization for briefing purposes. Electronic circulation is limited to dissemination by e-mail within such an organization by committee members.

Further copies of this draft may be purchased from BSI Shop <http://shop.bsigroup.com> or from BSI Customer Services, Tel: +44(0) 20 8996 9001 or email cservices@bsigroup.com. British, International and foreign standards are also available from BSI Customer Services.

Information on the co-operating organizations represented on the committees referenced above may be obtained from <https://standardsdevelopment.bsigroup.com/>.

E-mail: admin.start@bsigroup.com

Introduction

This draft standard is based on national and international discussions. Your comments on this draft are invited and will assist in the preparation of the consequent standard.

For international standards, comments will be reviewed by the relevant UK national committee before sending the consensus UK vote and comments to the international committee, which will then decide appropriate action. If the international standard is approved, it is usual for the text to be published as a British Standard.

For national standards, comments will be reviewed by the relevant UK national committee and the resulting standards published as a British Standard.

UK Vote

Please indicate whether you consider the UK should submit a negative (with supporting technical reasons) or positive vote on this draft. Please indicate if you are aware of any reason why this draft standard should not be published as a British Standard.

Submission of Comments

- **Annotated drafts are not acceptable and will be rejected.**
- All comments should be submitted online at <https://standardsdevelopment.bsigroup.com/>. You will need to register in order to comment.

Template for comments and secretariat observations

Date: xx/xx/20xx	Document: ISO/DIS xxxx
------------------	------------------------

1	2	(3)	4	5	(6)	7
MB	Clause No./ Subclause No./Annex (e.g. 3.1)	Paragraph/ Figure/Table/Note	Type of comment	Comment (justification for change) by the MB	Proposed change by the MB	Secretariat observations on each comment submitted
	3.1	Definition 1	ed	Definition is ambiguous and needs clarifying.	Amend to read '...so that the mains connector to which no connection...'	
	6.4	Paragraph 2	te	The use of the UV photometer as an alternative cannot be supported as serious problems have been encountered in its use in the UK.	Delete reference to UV photometer.	



DRAFT International Standard

ISO/DIS 12234-4

Digital imaging — Image storage — Part 4: Digital negative format

ICS: 37.040.99

ISO/TC 42

Secretariat: **ANSI**

Voting begins on:
2025-01-08

Voting terminates on:
2025-04-02

This document is circulated as received from the committee secretariat.

THIS DOCUMENT IS A DRAFT CIRCULATED FOR COMMENTS AND APPROVAL. IT IS THEREFORE SUBJECT TO CHANGE AND MAY NOT BE REFERRED TO AS AN INTERNATIONAL STANDARD UNTIL PUBLISHED AS SUCH.

IN ADDITION TO THEIR EVALUATION AS BEING ACCEPTABLE FOR INDUSTRIAL, TECHNOLOGICAL, COMMERCIAL AND USER PURPOSES, DRAFT INTERNATIONAL STANDARDS MAY ON OCCASION HAVE TO BE CONSIDERED IN THE LIGHT OF THEIR POTENTIAL TO BECOME STANDARDS TO WHICH REFERENCE MAY BE MADE IN NATIONAL REGULATIONS.

RECIPIENTS OF THIS DRAFT ARE INVITED TO SUBMIT, WITH THEIR COMMENTS, NOTIFICATION OF ANY RELEVANT PATENT RIGHTS OF WHICH THEY ARE AWARE AND TO PROVIDE SUPPORTING DOCUMENTATION.

Reference number
ISO/DIS 12234-4:2025(en)

Full document details: 25/30470379 DC

© ISO 2025  **Documents**
Document Sharing Hub



COPYRIGHT PROTECTED DOCUMENT

© ISO 2025

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword	viii
Introduction	ix
1 Scope	1
2 Normative references	1
3 Terms and definitions	1
4 DNG Format Overview	3
4.1 File Extensions	3
4.2 SubIFD Trees	3
4.3 Byte Order	4
4.4 Masked Pixels	4
4.5 Defective Pixels	4
4.6 Metadata	4
4.7 Proprietary Data	4
4.8 Camera Profiles	4
4.9 Opcode Lists	6
4.10 Floating Point Image Data	6
4.11 Transparency	6
4.12 Proxy DNG Files	6
4.13 Depth Maps	6
4.14 Enhanced Image Data	7
4.15 Semantic Masks	7
4.16 64-bit Format	8
5 Restrictions and Extensions to Existing TIFF 6.0 Tags	8
5.1 NewSubFileType	8
5.1.1 Previews	8
5.1.2 Transparency	9
5.1.3 Depth Maps	9
5.1.4 Enhanced Image Data	9
5.1.5 Semantic Masks	9
5.2 BitsPerSample	9
5.3 SampleFormat	9
5.4 Compression	9
5.5 Predictor	10
5.6 PhotometricInterpretation	11
5.7 Orientation	11
6 DNG Image Processing Model	11
6.1 Mapping Raw Values to Linear Reference Values	11
6.1.1 Linearization	12
6.1.2 Black Subtraction	12
6.1.3 Rescaling (Normalization)	12
6.1.4 Clipping	12
6.2 Mapping Camera Colour Space to CIE XYZ Space	12
6.2.1 Camera Calibration Matrices	12
6.2.2 One, Two, or Three Colour Calibrations	12
6.2.3 Definitions used in the following sections	13
6.2.4 Translating White Balance xy Coordinates to Camera Neutral Coordinates	13
6.2.5 Translating Camera Neutral Coordinates to White Balance xy Coordinates	14
6.2.6 Camera to XYZ (D50) Transform	14
6.3 Applying the Hue/Saturation/Value Mapping Table	15
6.3.1 Special compatibility note with DNG 1.2	15
6.4 Opcodes	16
7 DNG Tags	17

7.1	ActiveArea	17
7.2	AnalogBalance	17
7.3	AntiAliasStrength	18
7.4	AsShotICCProfile	18
7.5	AsShotNeutral	19
7.6	AsShotPreProfileMatrix	19
7.7	AsShotProfileName	20
7.8	AsShotWhiteXY	20
7.9	BaselineExposure	20
7.10	BaselineExposureOffset	21
7.11	BaselineNoise	21
7.12	BaselineSharpness	22
7.13	BayerGreenSplit	22
7.14	BestQualityScale	23
7.15	BlackLevel	23
7.16	BlackLevelDeltaH	23
7.17	BlackLevelDeltaV	24
7.18	BlackLevelRepeatDim	24
7.19	CalibrationIlluminant1	24
7.20	CalibrationIlluminant2	25
7.21	CalibrationIlluminant3	25
7.22	CameraCalibration1	26
7.23	CameraCalibration2	26
7.24	CameraCalibration3	26
7.25	CameraCalibrationSignature	27
7.26	CameraSerialNumber	27
7.27	CFALayout	27
7.28	CFAPattern	28
7.29	CFAPlaneColour	28
7.30	CFARepetPatternDim	29
7.31	ChromaBlurRadius	29
7.32	ColourimetricReference	30
7.33	ColourMatrix1	30
7.34	ColourMatrix2	30
7.35	ColourMatrix3	31
7.36	ColumnInterleaveFactor	31
	7.36.1 Potential Usage	31
7.37	CurrentICCProfile	32
7.38	CurrentPreProfileMatrix	32
7.39	DefaultBlackRender	32
7.40	DefaultCropOrigin	33
7.41	DefaultCropSize	33
7.42	DefaultScale	33
7.43	DefaultUserCrop	34
7.44	DepthFar	34
7.45	DepthFormat	34
7.46	DepthMeasureType	35
7.47	DepthNear	35
7.48	DepthUnits	36
7.49	DNGBackwardVersion	36
7.50	DNGPrivateData	36
7.51	DNGVersion	37
7.52	EnhanceParams	37
7.53	ExtraCameraProfiles	37
7.54	ForwardMatrix1	38
7.55	ForwardMatrix2	38
7.56	ForwardMatrix3	38
7.57	IlluminantData1	39
7.58	IlluminantData2	40

7.59	IlluminantData3	40
7.60	ImageSequenceInfo	40
7.61	ImageStats	41
7.62	JXLDistance	44
7.63	JXLEffort	44
7.64	JXLDecodeSpeed	44
7.65	LensInfo	45
7.66	LinearizationTable	45
7.67	LinearResponseLimit	45
7.68	LocalizedCameraModel	46
7.69	MakerNoteSafety	46
7.70	MaskedAreas	47
7.71	MaskSubArea	47
	7.71.1 Example	48
7.72	NewRawImageDigest	48
7.73	NoiseProfile	49
7.74	NoiseReductionApplied	50
7.75	OpcodeList1	50
7.76	OpcodeList2	50
7.77	OpcodeList3	51
7.78	OriginalBestQualityFinalSize	51
7.79	OriginalDefaultCropSize	51
7.80	OriginalDefaultFinalSize	52
7.81	OriginalRawFileData	52
7.82	OriginalRawFileDigest	53
7.83	OriginalRawFileName	54
7.84	PreviewApplicationName	54
7.85	PreviewApplicationVersion	54
7.86	PreviewColourSpace	54
7.87	PreviewDateTime	55
7.88	PreviewSettingsDigest	55
7.89	PreviewSettingsName	56
7.90	ProfileCalibrationSignature	56
7.91	ProfileCopyright	56
7.92	ProfileDynamicRange	56
7.93	ProfileEmbedPolicy	58
7.94	ProfileGainTableMap	58
7.95	ProfileGainTableMap2	60
	7.95.1 DataType	61
	7.95.2 Gamma	62
	7.95.3 GainMin and GainMax	62
	7.95.4 Application	62
	7.95.5 Compatibility with ProfileGainTableMap	62
	7.95.6 IFD 0 vs Camera Profile IFD	62
	7.95.7 IFD 0 vs Raw IFD	63
7.96	ProfileGroupName	63
7.97	ProfileHueSatMapData1	63
7.98	ProfileHueSatMapData2	64
7.99	ProfileHueSatMapData3	64
7.100	ProfileHueSatMapDims	64
7.101	ProfileHueSatMapEncoding	65
7.102	ProfileLookTableEncoding	66
7.103	ProfileLookTableData	67
7.104	ProfileLookTableDims	67
7.105	ProfileName	68
7.106	ProfileToneCurve	68
7.107	RawDataUniqueID	69
7.108	RawImageDigest	69
7.109	RawToPreviewGain	69

7.110	ReductionMatrix1	70
7.111	ReductionMatrix2	70
7.112	ReductionMatrix3	70
7.113	RGBTables	71
	7.113.1 Notes on the tag fields	72
	7.113.2 Background Table	73
	7.113.3 Table transform pipeline	74
	7.113.4 Notes on applying the colour transforms	75
7.114	RowInterleaveFactor	76
7.115	SemanticInstanceID	76
7.116	SemanticName	77
7.117	SubTileBlockSize	77
7.118	UniqueCameraModel	78
7.119	WhiteLevel	78
8	Opcodes	78
8.1	DeltaPerColumn	78
	8.1.1 Syntax	78
	8.1.2 Description	79
8.2	DeltaPerRow	79
	8.2.1 Syntax	79
	8.2.2 Description	80
8.3	FixBadPixelsConstant	80
	8.3.1 Syntax	80
	8.3.2 Description	80
8.4	FixBadPixelsList	81
	8.4.1 Syntax	81
	8.4.2 Description	81
8.5	FixVignetteRadial	81
	8.5.1 Syntax	82
	8.5.2 Description	82
8.6	GainMap	83
	8.6.1 Syntax	83
	8.6.2 Description	83
8.7	MapPolynomial	84
	8.7.1 Syntax	84
	8.7.2 Description	85
8.8	MapTable	85
	8.8.1 Syntax	85
	8.8.2 Description	86
8.9	ScalePerColumn	86
	8.9.1 Syntax	86
	8.9.2 Description	87
8.10	ScalePerRow	87
	8.10.1 Syntax	87
	8.10.2 Description	88
8.11	TrimBounds	88
	8.11.1 Syntax	88
	8.11.2 Description	88
8.12	WarpFisheye	88
	8.12.1 Syntax	89
	8.12.2 Description	89
	8.12.3 Notes and Restrictions	90
8.13	WarpRectilinear	90
	8.13.1 Syntax	91
	8.13.2 Description	91
	8.13.3 Notes and Restrictions	92
8.14	WarpRectilinear2	93
	8.14.1 Syntax	93

ISO/DIS 12234-4:2025(en)

8.14.2	Description	94
8.14.3	Special Compatibility Note and Relationship to Previous Warp Opcodes	94
Annex A	(informative) Compatibility with Legacy Versions of DNG	96
Bibliography		100

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html.

This document was prepared by Technical Committee ISO/TC 42 *Photography*.

A list of all parts in the ISO 12234 series can be found on the ISO website.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

Introduction

Advanced and professional photographers often choose to capture and edit photos in raw format, for additional flexibility and artistic control. Unlike display-referred formats (such as JPEG, PNG, and TIFF) which store images that have already been processed by the camera, raw formats store unprocessed or minimally processed data directly from the camera sensor. This enables photographers to adjust many parameters to taste, including white balance, tone mapping, noise reduction, sharpening, and so on. Photos in raw format are analogous to film negatives in a photographer's workflow, so they are often referred to as "digital negatives."

The purpose of this document is to define a standardized file format for storing raw photographs. This file format can be used in a wide range of hardware and software applications for generating, processing, managing, transcoding, or archiving raw photographs. Since raw photographs are closely tied to the camera sensor and other characteristics, camera manufacturers may use and support their own Vendor Raw formats to meet various design requirements.

The name of this file format is Digital Negative (DNG). DNG is an extension of TIFF 6.0 and is compatible with TIFF-EP standard. It is possible (but not required) for a DNG file to simultaneously comply with both this specification and the TIFF-EP standard.

Digital imaging — Image storage —

Part 4: Digital negative format

1 Scope

This part of ISO 12234 specifies the Digital Negative (DNG) image file format, published by Adobe Systems as the "Digital Negative (DNG) Specification." This includes specifying DNG by capturing and reconciling major adoptions and implementations of Adobe DNG.

A DNG file shall meet the requirements provided in this specification.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO 12234-1, *Electronic still-picture imaging — Removable memory — Part 1: Basic removable-memory model*

ISO 12234-2, *Electronic still-picture imaging — Removable memory — Part 2: TIFF/EP image data format*

ISO 12234-3, *Electronic still picture imaging — Removable memory — Part 3: XMP for digital photography*

ISO/TS 22028-3, *Photography and graphic technology — Extended colour encodings for digital image storage, manipulation and interchange — Part 3: Reference input medium metric RGB colour image encoding (RIMM RGB)*

IEC 61966-2-1, *Multimedia systems and equipment — Colour measurement and management — Part 2-1: Colour management — Default RGB colour space — sRGB*

IEC 61966-2-2, *Multimedia systems and equipment — Colour measurement and management — Part 2-2: Colour management — Extended RGB colour space — scRGB*

ISO/IEC 10918-1, *Information technology — Digital compression and coding of continuous-tone still images: Requirements and guidelines*

ISO/IEC 14495-1, *Information technology — Lossless and near-lossless compression of continuous-tone still images: Baseline — Part 1:*

CIPA DC-010-2024, *Exif metadata for XMP*

IPTC Photo Metadata Standard 2023.2

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <https://www.electropedia.org/>

3.1 **opcode**

this specifies an image processing instruction to be executed on reading

3.2 **"raw"**

Unprocessed or minimally processed image data which is directly related to the color information captured by the sensor and suitable for further processing. Note that in this document we distinguish between "Vendor Raw" and "Standard Raw", defined below, and use 'raw' to refer to 'Standardized Raw'

3.3 **reader**

device that consumes files compliant with this standard

3.4 **Standard Raw**

standard Raw refers to the raw files conformant with this specification. This is in contrast to vendor Raw files which are typically vendor specific. In the context of this document, Raw==Standard Raw and Vendor Raw will be referred to as Vendor Raw

3.5 **structural metadata**

metadata required for the expected display of the image

3.6 **tag**

Each tag is a 12-byte record with the following format:

Bytes 0–1: The Tag ID Field, which is a number that identifies the particular tag.

Bytes 2–3: The Data Type Field, which is one of the following values:

- 1 = BYTE, an 8-bit unsigned integer.
- 2 = ASCII, a string of 8-bit bytes containing 7-bit ASCII character codes terminated with a NULL code.
- 3 = SHORT, a 16-bit (2-byte) unsigned integer.
- 4 = LONG, a 32-bit (4-byte) unsigned integer.
- 5 = RATIONAL, is two LONGs. The first LONG is the numerator and the second LONG is the denominator.
- 6 = SBYTE, is a signed 8-bit integer
- 7 = UNDEFINED, indicates an 8-bit byte that can take any value depending on the field definition.
- 8 = SSHORT, is a signed 16-bit integer
- 9 = SLONG, is a signed 32-bit integer
- 10 = SRATIONAL, is two SLONGs: the first represents the numerator of the fraction, the second represents the denominator, which shall be non-zero.
- 11 = FLOAT, is 32-bit (4-byte) float in IEEE-754 format
- 12 = DOUBLE, is 64-bit (8-byte) float in IEEE-754 format

Bytes 4–7: The Count Field, which indicates the number of Data Type values (not the number of bytes) stored by the tag.

Bytes 8–11: The Value or Offset Field, which indicates the tag value(s) if they can be stored in these 4 bytes, or the offset to the data if more room is required. The offset value is the relative offset from the Byte Order field in the file header.

3.7

TIFF field

logical entity consisting of a TIFF tag and its value, where this logical concept is implemented as an 12 byte IFD entry. The tag value(s) if they can be stored in the last 4 bytes, or the offset to the data if more room is required. The offset value is the relative offset from the Byte Order field, first two bytes of the file in the file header

3.8

DNG file

file compliant with this standard

3.9

Vendor Raw

in common usage, raw files can refer to a variety of non-standardized raw file formats independently specified by any number of vendors. For the purposes of this document, these will be referred to as 'Vendor Raw' and 'Raw' will refer to this 'Standardized Raw' specification

3.10

working colour space

colour space encoding in which manual edits or image enhancement processing is performed

Note 1 to entry: Manual edits in a working colour space are performed viewing the effect of the edits on some medium which becomes the reference medium for the editing process. Typically, the black and white points of this medium are encoded at the minimum and maximum working colour space values. In this case, in order to obtain accurate colourimetric values, the black and white scaling corresponding to the reference medium used must be undone. XYZ scaling is typically recommended (see ISO 22028-2).

Note 2 to entry: When a colour space encoding used for a standard colour image encoding is also used as a working space, it may still be necessary to perform a colour re-rendering of the values associated with the (potentially variable) working colour space reference medium to make them appropriate for the colour image encoding reference medium, and for open exchange. For example, When the sRGB colour space encoding specified in IEC 61966-2-1 is used as a working colour space for the editing of motion picture content, viewing it in a dark theatrical environment, the resulting images will be encoded in the sRGB colour space encoding but not in the sRGB colour image encoding. Consequently, they will not produce good results if exchanged and interpreted according to the sRGB colour image encoding standard.

Note 3 to entry: Image enhancement processing may involve the application of algorithms in different working spaces as different operations and algorithms may be optimized for particular working spaces.

4 DNG Format Overview

This section describes the DNG file format. A DNG file shall comply with all of the format requirements of TIFF 6.0, unless otherwise noted in this specification.

4.1 File Extensions

It is recommended to use the file extension DNG or dng. For example, IMG_1234.DNG or IMG_1234.dng.

4.2 SubIFD Trees

It is recommended that DNG files use SubIFD trees, as described in TIFF-EP. However, DNG files shall not use SubIFD chains.

The highest-resolution and quality IFD shall use a NewSubFileType value of 0.

Reduced-resolution or reduced-quality thumbnails or previews, if any, should use NewSubFileType equal to 1 (for a primary preview) or 10001.H (for an alternate preview).

It is recommended that the first IFD contain a low-resolution thumbnail, as described in TIFF-EP.

4.3 Byte Order

DNG files may use either byte order (big-endian or little-endian).

4.4 Masked Pixels

Some camera sensors measure the black encoding level using masked pixels at the edges of the sensor. These pixels can either be trimmed before storing the image, or they can be included in the stored image. If the masked pixels are not trimmed, the area of the non-masked pixels shall be specified using the ActiveArea tag.

The black encoding level information extracted from these masked pixels should be used to either pre-compensate the raw data stored in the file or they should be included in the file using the tags for specifying the black level.

This black encoding level information shall be included, even if the masked pixels are not trimmed, to allow DNG readers to process the image without requiring knowledge of the best way to compute the black levels for any given camera model.

4.5 Defective Pixels

There are two methods to treat defective pixels in a DNG file.

The first method is to remove (for instance, interpolate over) the defective pixels before storing the raw image into the DNG file.

The second is to include an appropriate opcode (either FixBadPixelsConstant or FixBadPixelsList) in the OpcodeList1 tag.

4.6 Metadata

Metadata can be stored in DNG in the following ways:

- Using TIFF-EP (ISO 12234-2) metadata tags
- Using EXIF metadata tags
- Using the IPTC metadata tag (tag code 33723)
- Using the XMP metadata tag (tag code 700)

NOTE TIFF-EP and EXIF use nearly the same metadata tag set, but TIFF-EP stores the tags in IFD 0, while EXIF store the tags in a separate IFD. It is recommended to use a separate EXIF IFD.

4.7 Proprietary Data

Proprietary data may be stored using private tags, private IFDs, or a private MakerNote. It is recommended to use the DNGPrivateData and MakerNoteSafety tags to ensure that programs that edit DNG files preserve this proprietary data.

NOTE Camera manufacturers can include proprietary data in a raw file for use in their own software. Use of such proprietary data may result in different images.

4.8 Camera Profiles

A DNG file shall contain at least 1 camera profile. The primary camera profile shall be stored in IFD 0.

A DNG file may store additional camera profiles using the ExtraCameraProfiles tag.

A camera profile consists of a set of tags. Some tags are mandatory, while others are optional.

The tags belonging to a camera profile are:

- BaselineExposureOffset
- CalibrationIlluminant1
- CalibrationIlluminant2
- CalibrationIlluminant3
- ColourMatrix1
- ColourMatrix2
- ColourMatrix3
- DefaultBlackRender
- ForwardMatrix1
- ForwardMatrix2
- ForwardMatrix3
- IlluminantData1
- IlluminantData2
- IlluminantData3
- ProfileCalibrationSignature
- ProfileCopyright
- ProfileDynamicRange
- ProfileEmbedPolicy
- ProfileGroupName
- ProfileHueSatMapData1
- ProfileHueSatMapData2
- ProfileHueSatMapData3
- ProfileHueSatMapDims
- ProfileHueSatMapEncoding
- ProfileLookTableData
- ProfileLookTableDims
- ProfileLookTableEncoding
- ProfileName
- ProfileToneCurve
- ReductionMatrix1
- ReductionMatrix2
- ReductionMatrix3

4.9 Opcode Lists

The concept of “Opcode Lists” enables additional image processing steps to be specified in an extensible manner. Opcodes are specified in one of the following three tags:

- OpcodeList1
- OpcodeList2
- OpcodeList3

4.10 Floating Point Image Data

Images in DNG files may be stored using IEEE-754 floating point data type with bit depths of 16, 24, or 32 bits per sample. In addition, images in DNG files may be stored using integer data types as described in TIFF 6.0.

4.11 Transparency

This feature allows transparency (or alpha) to be specified for the image data. Transparency is useful when not all pixels in the image bounding rectangle are defined. For example, after combining several source images into a panorama, the edge of the combined image is often not perfectly rectangular.

If transparency is stored in a DNG file, it shall be stored as an IFD with a NewSubFileType value of 4 or 5.

Fully transparent pixels shall be stored with a pixel value of zero.

For floating-point transparency images, fully opaque pixels shall be stored with a value of 1.0. For integer transparency images, fully opaque pixels shall be stored with the maximum integer value ($2^N - 1$ for a bit depth of N).

NOTE 1 Intermediate values (between zero and the maximum encoding value) correspond to semi-transparent pixels.

NOTE 2 For semi-transparent pixels, the image values shall be stored as if the pixel is fully opaque.

4.12 Proxy DNG Files

A DNG file may be a “proxy DNG” – a lower resolution or lower quality version of an original DNG. Proxy DNG files may be useful placeholders for original DNGs in applications with limited bandwidth or available storage space.

In a proxy DNG, the tags OriginalDefaultFinalSize, OriginalBestQualitySize, and OriginalCropSize may be used to indicate the size of the corresponding original DNG. These tags provide information on the pixel size of the original DNG file to allow rendering parameters to be transferred back to the original DNGs with similar visual results.

It is recommended to use either JPEG and JPEG XL (lossy mode) for compressing images in proxy DNG files, because doing so will greatly reduce the file size.

4.13 Depth Maps

A DNG file may store a depth map, which allows estimated or measured depth information to be specified for the image.

If a depth map is stored in a DNG file, it shall be stored in an IFD with NewSubFileType value of 8 or 9.

The depth map IFD shall use a PhotometricInterpretation value of 51177.

The depth map IFD shall be stored as an 8 or 16-bit single channel unsigned integer image. Zero encodes the “nearest distance” (closest to the camera), and the maximum unsigned integer encodes the “farthest

distance". The mapping of the depth map integer values to distance values is specified by the DepthFormat, DepthNear, DepthFar, DepthUnits, and DepthMeasureType tags.

NOTE The depth map should match the field of view of the main image. It may have a different pixel resolution than the main image. If the depth map is resized to match the size of the main image, the pixels should correspond. The depth values should be filled in by an estimation algorithm for any areas of the depth map that cannot be measured directly (i.e., no holes are allowed in depth maps).

4.14 Enhanced Image Data

The concept of "enhanced image data" in DNG is a more processed (and always at least demosaiced) version of the raw image data.

Enhanced image data is optional. If included, it shall be stored in an IFD with a NewSubFileType value of 16.

The enhanced image data IFD shall use a PhotometricInterpretation value of 34892 (LinearRaw).

The enhanced image data IFD shall use the same colour space as the raw image data.

The enhanced image should use either 16-bit unsigned integer encoding, or 16 or 24 or 32-bit floating point encoding.

The enhanced image data shall be mapped to linear reference values, with the exception that it may have a non-zero BlackLevel tag if unsigned integer encoding is used. This black level shall be the same value for all colour planes of the enhanced image data.

The enhanced IFD may have its own instances of the BaselineSharpness, NoiseProfile, and NoiseReductionApplied tags, which may differ from the raw IFD versions of these tags due to sharpening or noise reduction having been applied to the enhanced image data.

The enhanced IFD should have a non-null EnhanceParams tag to document the operations used to enhance the image data.

4.15 Semantic Masks

The concept of "Semantic Masks" in DNG may be used for processing specific image regions, based on image segmentation. For example, applying a RGB lookup table to the sky.

Semantic masks are optional. If included, a semantic mask shall be stored in an IFD with a NewSubFileType value of 65540 (10004.H).

A semantic mask IFD shall use a PhotometricInterpretation of 52527 (CD2F.H).

A semantic mask shall be stored as a single channel image with one of the following pixel data types, with valid encoding range indicated in parentheses:

- 8-bit unsigned integer (0 to 255)
- 16-bit unsigned integer (0 to 65535)
- 32-bit floating point (0.0 to 1.0)

The minimum encoding value (zero) represents "not a part of the semantic attribute."

The maximum encoding value represents "is a part of the semantic attribute."

Encoding values in between the minimum and maximum values are allowed; they represent a partial semantic attribute. These fractional values may be used for matting or blending purposes.

A semantic mask should match the field of view of the main image, as specified by the ActiveArea tag. It may have a different pixel resolution than the main image.

Semantic masks shall be stored uncompressed, or compressed with one of the following methods:

Compression Method	Compression Tag Value	Data Types Allowed
Lossless JPEG	7 decimal (7.H)	8-bit and 16-bit unsigned integer
Baseline DCT (lossy) JPEG	34892 decimal (884C.H)	8-bit unsigned integer
Deflate / Zip	8 decimal (8.H)	8-bit and 16-bit unsigned integer 32-bit float
JPEG XL	52546 decimal (CD42.H)	8-bit and 16-bit unsigned integer 16-bit float

Multiple semantic masks may be stored in a DNG file. Each semantic mask IFD should include a SemanticName tag with a unique string that identifies the purpose of the mask.

A semantic mask shall have the same orientation as the main image.

4.16 64-bit Format

TIFF format was designed using 32-bit values to store file offsets, which limits the maximum size of a TIFF file to 4 gigabytes. There is a 64-bit extension of TIFF (known as BigTIFF) which does not have this file size limitation. The DNG format supports this same 64-bit extension. Note that 64-bit format support is independent of the DNG version.

The differences from 32-bit DNG to 64-bit DNG are:

1. The 16-bit value at file offset 2 is the constant 43 (instead of 42).
2. The 16-bit value at file offset 4 is the constant 8.
3. The 16-bit value at file offset 6 is the constant 0.
4. The 64-bit value at file offset 8 is the offset to the first image file directory.
5. The entry count for each image file directory is a 64-bit value (instead of 16).
6. The element count for an image file directory entry is a 64-bit value (instead of 32).
7. The data entry for each image file directory entry is a 64-bit value (instead of 32).
8. The offset to the next image file directory is a 64-bit value (instead of 32).

Also, three new data types have been defined for IFD entries:

- 16, which is an unsigned 64-bit integer.
- 17, which is a signed 64-bit integer.
- 18, which is a 64-bit file offset for an IFD.

To maximize backwards compatibility, it is recommended that DNG writers only use the 64-bit extension of DNG when the resulting file would be larger than 4 gigabytes.

5 Restrictions and Extensions to Existing TIFF 6.0 Tags

5.1 NewSubFileType

5.1.1 Previews

A NewSubFileType equal to 10001.H may be used for alternative or non-primary rendered previews. Doing so enables multiple renderings (not just multiple sizes of a single rendering) to be stored in a DNG file. DNG reading software that displays a preview for a DNG file should, by default, display a preview from an IFD with NewSubFileType equal to 1. Alternative renderings should only be displayed if requested by the user.

5.1.2 Transparency

Transparency information for the full resolution raw image, if included, shall use NewSubFileType equal to 4. Transparency information for reduced resolution versions of raw images, if included, shall use NewSubFileType equal to 5.

5.1.3 Depth Maps

Depth map information of the highest resolution, if included, shall use NewSubFileType equal to 8. Depth map information of lower resolutions, if included, shall use a NewSubFileType of 9.

5.1.4 Enhanced Image Data

Enhanced image data, if included, shall use a NewSubFileType equal to 16.

5.1.5 Semantic Masks

Semantic mask information, if included, shall use NewSubFileType equal to 65540 (10004.H). Multiple semantic masks can be stored in a DNG file.

5.2 BitsPerSample

For integer images, BitsPerSample shall be an integer in the range 8 to 32.

For floating-point images, BitsPerSample shall be 16, 24, or 32.

If SamplesPerPixel is not equal to 1, the bit depth shall be the same value for each sample.

If BitsPerSample is not equal to 8 or 16 or 32, then the bits shall be packed into bytes using the TIFF default FillOrder of 1 (big-endian), even if the DNG file itself uses little-endian byte order.

5.3 SampleFormat

SampleFormat value shall be either 1 (unsigned integer) or 3 (IEEE floating point).

5.4 Compression

The compression tag codes supported by DNG are:

- 1: Uncompressed data.
- 7: JPEG compressed data, either baseline DCT JPEG, or lossless JPEG compression.
- 8: Deflate (ZIP)
- 34892: Lossy JPEG
- 52546: JPEG XL

If PhotometricInterpretation = 6 (YCbCr) and BitsPerSample = 8/8/8, or if PhotometricInterpretation = 1 (BlackIsZero) and BitsPerSample = 8, then the JPEG variant shall be baseline DCT JPEG.

Otherwise, the JPEG variant shall be lossless Huffman JPEG. For lossless JPEG, the internal width/length/components in the JPEG stream are not required to match the strip or tile's width/length/components. Only the total sample counts shall match.

NOTE 1 It is common for Colour Filter Array images to be encoded with a different width, length or component count to allow the JPEG compression predictors to work across like colours.

Deflate (8) compression may be used for floating-point image data, 32-bit integer image data, transparency mask data, and depth map data.

Lossy JPEG (34892) may be used in IFDs that use 8-bit integer data and one of the following PhotometricInterpretation values:

- 34892 (LinearRaw)
- 52527 (PhotometricMask)

NOTE 2 This compression code informs the DNG reader to use a lossy JPEG decoder rather than a lossless JPEG decoder for this combination of PhotometricInterpretation and BitsPerSample.

JPEG XL compression (see ISO/IEC 18181-1:2022) may be used for N-bit integer image data, where $8 \leq N \leq 16$, and 16-bit floating-point image data. The number of images planes shall be either 1 or 3. The PhotometricInterpretation shall be one of the following values:

- 0 (piWhiteIsZero)
- 1 (piBlackIsZero)
- 2 (piRGB)
- 4 (piTransparencyMask)
- 32803 (Colour Filter Array)
- 34892 (LinearRaw)
- 51177 (Depth)
- 52527 (PhotometricMask)

JPEG XL supports two types of bitstreams:

- A “bare” format with no metadata
- A “container” format that uses the ISO Base Media File Format box structure with optional metadata

Either type of bitstream may be used in DNG.

NOTE 3 When using JPEG XL with multiple tiles or strips, it is recommended to use the bare format, because each compressed tile will be slightly smaller, and per-tile metadata is not required.

NOTE 4 When using JPEG XL to represent a preview IFD, such as a thumbnail or full-resolution rendered preview, it is recommended to use a single tile with the container format, so that utilities can extract a readable, standalone JPEG XL file by simply copying the bytes from the IFD.

5.5 Predictor

The following predictor codes are allowed for IFDs using compression code 8 (deflate):

- 1 = Null Predictor. No predictor is used as part of the compression.
- 2 = Horizontal Difference.
- 3 = Floating Point.
- 34892 = Horizontal Difference X2. Same as Horizontal Difference except the pixel two to the left is used rather than the pixel one to the left.
- 34893 = Horizontal Difference X4. Same as Horizontal Difference except the pixel four to the left is used rather than the pixel one to the left.
- 34894 = Floating Point X2. Same as Floating Point except the pixel two to the left is used rather than the pixel one to the left.

- 34895 = Floating Point X4. Same as Floating Point except the pixel four to the left is used rather than the pixel one to the left.

5.6 PhotometricInterpretation

The following PhotometricInterpretation values may be used only in thumbnail and preview IFDs:

- 1 = BlackIsZero. The image is assumed to be in a gamma 2.2 colour space, unless otherwise specified using PreviewColourSpace tag.
- 2 = RGB. The image is assumed to be in the sRGB colour space, unless otherwise specified using the PreviewColourSpace tag.
- 6 = YCbCr. Used for JPEG-encoded preview images.

The following values may be used in the raw IFD. They are assumed to be in the camera's native colour space:

- 32803 = CFA (Colour Filter Array).
- 34892 = LinearRaw.

If PhotometricInterpretation is set to CFA (32803), CFARepetPatternDim and CFAPattern shall be included in the same IFD. The origin of the repeating CFA pattern is the top-left corner of the ActiveArea rectangle.

The LinearRaw PhotometricInterpretation value is intended for use by cameras that do not use colour filter arrays, but instead capture all colour components at each pixel. It can also be used for CFA data that has already been demosaiced.

The LinearRaw value may be used in reduced resolution IFDs and enhanced IFDs, even if the raw IFD uses the CFA PhotometricInterpretation value.

The following PhotometricInterpretation value shall be used for depth map IFDs:

- 51177 = Depth

The following PhotometricInterpretation value shall be used for semantic mask IFDs:

- 52527 = PhotometricMask

5.7 Orientation

The TIFF (274==112.H) Orientation tag shall be included in a DNG file.

NOTE With the Orientation tag present, file browsers can perform lossless rotation of DNG files by modifying a single byte of the file. DNG readers should support all possible orientations, including mirrored orientations. Note that the mirrored orientations are not allowed by the TIFF-EP specification, so writers should not use them if they want their files be compatible with TIFF-EP.

6 DNG Image Processing Model

6.1 Mapping Raw Values to Linear Reference Values

The section describes the processing model for mapping stored raw sensor values into linear reference values.

Linear reference values encode zero light as 0.0, and the maximum useful value (limited by either sensor saturation or analog to digital converter clipping) as 1.0. If SamplesPerPixel is greater than one, each sample plane shall be processed independently.

The processing model follows these steps:

- Linearization

- Black Subtraction
- Rescaling
- Clipping

6.1.1 Linearization

The first step is to process the raw values through the look-up table specified by the `LinearizationTable` tag, if any. If the raw value is greater than the size of the table, it shall be mapped to the last entry of the table.

6.1.2 Black Subtraction

The black level for each pixel is then computed and subtracted. The black level for each pixel is the sum of the black levels specified by the `BlackLevel`, `BlackLevelDeltaH` and `BlackLevelDeltaV` tags.

6.1.3 Rescaling (Normalization)

The black subtracted values are then rescaled to map them to the logical 0.0 to 1.0 range. The scale factor is the inverse of the difference between the value specified in the `WhiteLevel` tag and the maximum computed black level for the sample plane.

6.1.4 Clipping

Rescaled values above 1.0 shall be clipped to 1.0. Rescaled values below 0.0 may be clipped to 0.0, but it is recommended to preserve negative values for at least the early stages of an image rendering pipeline (e.g., for better noise reduction in the shadows).

6.2 Mapping Camera Colour Space to CIE XYZ Space

This section describes the processing model for mapping between the camera colour space coordinates (linear reference values) and CIE XYZ (with a D50 white point).

6.2.1 Camera Calibration Matrices

Camera calibration tags using different reference cameras are supported.

When rendering a file using a camera profile, it is important to know if the selected camera profile was designed using the same reference camera used to create the camera calibration tags. If so, then the camera calibration tags should be used. If not, then it is preferable to ignore the camera calibration tags and use identity matrices instead in order to minimize the worst case calibration mismatch error.

This matching is done by comparing the `CameraCalibrationSignature` tag and the `ProfileCalibrationSignature` tag for the selected camera profile. If they match, then use the camera calibration tags. If not, then use identity matrices.

6.2.2 One, Two, or Three Colour Calibrations

DNG provides for one, two, or three sets of colour calibration tags, each set optimized for a different illuminant. If two or three sets of colour calibration tags are included, then the raw converter should interpolate between the calibrations based on the white balance selected by the user.

6.2.2.1 Information for two calibrations

If exactly two calibrations are included, then it is recommended that one of the calibrations be for a low colour temperature illuminant (e.g., Standard-A) and the second calibration illuminant be for a higher colour temperature illuminant (e.g., D55 or D65). This combination has been found to work well for a wide range of real-world digital camera images.

Interpolation between two illuminants shall be performed by linear interpolation using inverse correlated colour temperature. To find the interpolation weighting factor between the two tag sets, find the correlated colour temperature for the user-selected white balance and the two calibration illuminants. If the white balance temperature is between two calibration illuminant temperatures, then invert all the temperatures and use linear interpolation. Otherwise, use the closest calibration tag set.

6.2.2.2 Information for three calibrations

If three calibrations are included, then it is recommended that the first two calibrations are as described in the preceding paragraph, and the third calibration be used to characterize the "as shot" illuminant: that is, the illuminant used to photograph the main image. This third calibration provides a colour transform optimized for the capture conditions, while providing additional calibrations for two other illuminants in case the user wishes to change the white balance significantly during raw conversion.

6.2.2.3 Requirements for three calibrations

If CalibrationIlluminant3 is included, then the following are required:

- CalibrationIlluminant1 and CalibrationIlluminant2 shall be included.
- ColourMatrix3 shall be included.
- Either all three ForwardMatrix tags shall be included, or none of them shall be included.
- Either all three ReductionMatrix tags shall be included, or none of them shall be included.
- Either all three ProfileHueSatMapData tags shall be included, or none of them shall be included.
- The white points (x-y chromaticity values) of all three calibration illuminants shall be distinct.

6.2.3 Definitions used in the following sections

- Let n be the dimensionality of the camera colour space (usually 3 or 4).
- Let CM be the n -by-3 matrix interpolated from the ColourMatrix1, ColourMatrix2, and ColourMatrix3 tags.
- Let CC be the n -by- n matrix interpolated from the CameraCalibration1, CameraCalibration2, and CameraCalibration3 tags (or identity matrices, if the signatures don't match).
- Let AB be the n -by- n matrix, which is zero except for the diagonal entries, which are defined by the AnalogBalance tag.
- Let RM be the 3-by- n matrix interpolated from the ReductionMatrix1, ReductionMatrix2, and ReductionMatrix3 tags.
- Let FM be the 3-by- n matrix interpolated from the ForwardMatrix1, ForwardMatrix2, and ForwardMatrix3 tags.

6.2.4 Translating White Balance xy Coordinates to Camera Neutral Coordinates

If the white balance is specified in terms of a CIE xy coordinate, then a camera neutral coordinate can be derived by first calculating the weighting factor between the two or three sets of colour calibration tags.

The XYZ to camera space matrix is:

$$\text{XYZtoCamera} = \text{AB} * \text{CC} * \text{CM}$$

The camera neutral can be found by expanding the xy value to a 3-by-1 XYZ matrix (assuming $Y = 1.0$) and multiplying it by the XYZtoCamera matrix:

$$\text{CameraNeutral} = \text{XYZtoCamera} * \text{XYZ}$$

6.2.5 Translating Camera Neutral Coordinates to White Balance xy Coordinates

This process is slightly more complex than the transform in the other direction because it requires an iterative solution.

1. Guess an xy value. Use that guess to find the interpolation weighting factor between the colour calibration tags. Find the XYZtoCamera matrix as above.
2. Find a new xy value by computing:

$$XYZ = \text{Inverse} (XYZtoCamera) * CameraNeutral$$
 (If the XYZtoCamera matrix is not square, then use the pseudoinverse.)
3. Convert the resulting XYZ to a new xy value.
4. Iterate until the xy values converge to a solution.

6.2.6 Camera to XYZ (D50) Transform

Two methods of specifying the camera to XYZ (D50) transform are supported, depending on whether or not the forward matrix tags are included in the camera profile.

The use of the forward matrix tags is recommended for two reasons. First, it allows the camera profile creator to control the chromatic adaptation algorithm used to convert between the calibration illuminant and D50. Second, it causes the white balance adjustment (if the user white balance does not match the calibration illuminant) to be done by scaling the camera coordinates rather than by adapting the resulting XYZ values, which has been found to work better in extreme cases.

6.2.6.1 Without ForwardMatrix

If the ForwardMatrix tags are not included in the camera profile, then:

1. Invert the XYZtoCamera matrix, where N is the number of rows in the matrix.

If $n = 3$, this is:

$$CameraToXYZ = \text{Inverse} (XYZtoCamera)$$

If $n > 3$, and the reduction matrix tags are included, then:

$$CameraToXYZ = \text{Inverse} (RM * XYZtoCamera) * RM$$

Otherwise:

$$CameraToXYZ = \text{PseudoInverse} (XYZtoCamera)$$

2. Compute the white balanced transform:

$$CameraToXYZ_D50 = CA * CameraToXYZ$$

CA, above, is a chromatic adaptation matrix that maps from the white balance xy value to the D50 white point. The recommended method for computing this chromatic adaptation matrix is to use the linear Bradford algorithm.

6.2.6.2 With ForwardMatrix

If the ForwardMatrix tags are included in the camera profile:

$$CameraToXYZ_D50 = FM * D * \text{Inverse} (AB * CC)$$

D, above, is a diagonal n-by-n matrix, computed so that the CameraToXYZ_D50 matrix maps the selected camera neutral to XYZ D50. The forward matrix is required to map a unit vector to XYZ D50 by definition, so D can be computed by finding the neutral for the reference camera:

$$\text{ReferenceNeutral} = \text{Inverse}(\text{AB} * \text{CC}) * \text{CameraNeutral}$$

And then:

$$D = \text{Invert}(\text{AsDiagonalMatrix}(\text{ReferenceNeutral}))$$

6.3 Applying the Hue/Saturation/Value Mapping Table

After the camera colours have been converted to XYZ (D50) values, the Hue/Saturation/Value mapping table, if any, is applied. If there are two Hue/Saturation/Value mapping tables, then they are interpolated in the same way that colour calibration tags are interpolated. If only one Hue/Saturation/Value table is included, then it is used regardless of the selected white balance.

1. First, the XYZ (D50) values are converted to linear RGB coordinates, using the ProPhoto RGB primaries. (This is also known as RIMM space).
2. If the camera profile is High Dynamic Range (see ProfileDynamicRange) and ValueDivisions is greater than 1, then apply the encoding function $f(x)$ to each colour component.
3. The linear RGB coordinates are converted to HSV coordinates (Hue-Saturation-Value).
4. The HSV coordinates are used to index the mapping table using tri-linear interpolation, resulting in three values: hue shift (in degrees); saturation scale factor; value scale factor. If the division count in a dimension is 1, then the table is constant for that dimension. If the camera profile is High Dynamic Range, the input value (V) coordinate is clamped to [0, 1] prior to indexing the mapping table.
5. Hue is indexed using “wrap-around” math. For example, if HueDivisions is equal to 3, then the table samples are at 0 degrees (red), 120 degrees (green), and 240 degrees (blue).
6. The hue coordinate is modified by adding the hue shift.
7. The saturation coordinate is modified by multiplying by the saturation scale factor, and then clipping to no more than 1.0.
8. The value coordinate is modified by multiplying by the value scale factor. If the camera profile is Standard Dynamic Range, clip the result to no more than 1.0.
9. The HSV coordinates are converted to linear RGB coordinates.
10. If the camera profile is High Dynamic Range (see ProfileDynamicRange) and ValueDivisions is greater than 1, then applying decoding function $f^{-1}(x)$ (inverse of the encoding function) to each colour component.
11. Convert from linear ProPhoto RGB back to XYZ (D50) values.

6.3.1 Special compatibility note with DNG 1.2

For maximum compatibility across legacy DNG readers created to be compliant with Adobe DNG version 1.7.1 and before, it is recommended to use the first two calibration illuminants for standard illuminant types (such as A and D65) and to reserve the third illuminant for custom and “as shot” illuminants. Legacy DNG readers older than version 1.6 will be able to read and interpret the color transforms (matrices and tables) from the first two illuminants. Newer DNG readers (DNG 1.6 and later) will be able to read and interpret the colour transforms matrices and tables from all three illuminants.

6.4 Opcodes

An opcode list is a list of opcodes, each of which specifies an image processing operation to be applied by the reader. There are three possible opcode list tags specifying operations at three distinct points in the processing pipeline.

Each opcode shall be applied in-order, following the sequence specified in the list.

Opcode lists shall be stored in big-endian byte order, regardless of the file's main byte order. This allows utility programs to copy opcode lists from file to file, without needing to understand their detailed internal structure.

At the start of the opcode list, there shall be a 32-bit unsigned integer count, which shall contain the number of opcodes in the list. This shall be followed by the data for each opcode.

Each opcode shall start with a 32-bit unsigned integer, which contains the opcode ID. The opcode ID identifies the specific opcode, which shall be one of the opcodes defined in this specification.

Next shall be a 32-bit unsigned integer, which contains the TIFF/DNG specification version in which the opcode ID was defined. The four-part version number shall be treated as a 4-digit number with the radix 256.

New opcode IDs may be defined in future specification versions. A reader shall not attempt to process an opcode with a version higher than the specification version that the reader was written to support.

Next shall be a 32-bit unsigned integer, which contain flag bits. There are two defined flag bits. When bit 0 (the least significant bit) is set to 1, the reader need not apply this opcode. When bit 1 (the second to least significant bit) is set to 1, the reader need not apply this opcode when doing “preview quality” processing, but the reader shall apply this opcode when doing “full quality” processing.

Next shall be a 32-bit unsigned integer, which contains the number of bytes in a variable size parameter area for the opcode. The format of this variable size parameter area is specified separately for each opcode.

After the application of each opcode, image values shall be clipped to the value range of the image being modified. For OpcodeList1, this range shall be 0 to 4294967295 ($2^{32} - 1$) for images with a bit depth greater than 16, otherwise 0 to 65535 ($2^{16} - 1$). For OpcodeList2 and OpcodeList3, the range shall be 0.0 to 1.0.

[Table 1](#) illustrates the opcode list structure.

Table 1 — Opcode List Structure

Type	Contents
LONG	Opcode Count (n)
LONG	Opcode ID 1
LONG	Version
LONG	Flag bits
LONG	Byte count (m_1)
BYTE[m_1]	Parameters
...	
LONG	Opcode ID n
LONG	Version
LONG	Flag bits
LONG	Byte count (m_n)
BYTE[m_n]	Parameters

As a special case, DNG readers should skip processing an opcode if both of the following conditions are met:

- it is a WarpRectilinear or WarpFisheye opcode, and

- the opcode immediately follows (in the same opcode list) is a WarpRectilinear2 opcode that is marked as Optional

This "skip rule" is intended to maximize compatibility across legacy DNG readers.

7 DNG Tags

This section describes the normative tags in a DNG file.

For each tag, the following information is given:

Tag The TIFF tag code as a decimal value, followed by the hexadecimal value in parentheses.

Type The TIFF type of the tag data. For example, "SHORT" for a 16-bit unsigned integer.

Count The TIFF tag count of the tag data (the number of typed elements).

Value The contents of the tag data.

Mandatory Indicates whether or not this tag is mandatory (must be present) in a DNG file.

Default The default value of the tag, if the tag is not present.

Usage Indicates the IFD(s) in which the tag may be included

7.1 ActiveArea

Tag 50829 (C68D.H)

Type SHORT or LONG

Count 4

Value See below

Mandatory No

Default 0, 0, ImageLength, ImageWidth

Usage Raw IFD

ActiveArea defines the rectangle of the active (non-masked) pixels of the sensor. The order of the rectangle coordinates specified in this tag shall be: top, left, bottom, right. The rectangle shall not extend beyond the image. The right value shall be greater than the left value (i.e., width of the rectangle is at least 1). The bottom value shall be greater than the top value (i.e., height of the rectangle is at least 1).

7.2 AnalogBalance

Tag 50727 (C627.H)

Type RATIONAL

Count ColourPlanes

Value See below

Mandatory No

Default All 1.0

Usage IFD 0

AnalogBalance defines the gain, either analog or digital, that has been already applied to the stored raw values. If a gain is applied, it is recommended that it be an analog gain.

NOTE Normally, the stored raw values are not white balanced, since any digital white balancing will reduce the dynamic range of the final image if the user decides to later adjust the white balance; however, if camera hardware is capable of white balancing the colour channels before the signal is digitized, it can improve the dynamic range of the final image.

7.3 AntiAliasStrength

Tag 50738 (C632.H)

Type RATIONAL

Count 1

Value Relative strength of the camera's anti-alias filter

Mandatory No

Default 1.0

Usage Raw IFD

The AntiAliasStrength tag shall specify the strength of the camera's anti-alias filter (also known as Optical LowPass Filter, OLPF).

The value shall be in the range [0, 1].

A value of 0 shall mean no anti-alias filter (that is, the camera is prone to aliasing artefacts with some subjects).

A value of 1 shall mean a strong anti-alias filter (that is, the camera almost never has aliasing artefacts).

NOTE 1 This tag overlaps in functionality with the BaselineSharpness tag. The primary difference is the AntiAliasStrength tag is used as a hint to the de-mosaic algorithm, while the BaselineSharpness tag is used as a hint to a sharpening algorithm applied later in the processing pipeline.

NOTE 2 This tag is intended to help readers tune their demosaic algorithms based on the strength of the sensor's low-pass filter. For example, a value of 0 indicates that more false colour suppression may be needed.

7.4 AsShotICCProfile

Tag 50831 (C68F.H)

Type UNDEFINED

Count Length of ICC profile in bytes

Value See below

Mandatory No

Default None

Usage IFD 0

This tag contains an ICC profile that, in conjunction with the AsShotPreProfileMatrix tag, provides the camera manufacturer with a way to specify a default colour rendering from camera colour space coordinates (linear reference values) into the ICC profile connection space.

NOTE 1 The ICC profile connection space is an output referred colourimetric space, whereas the other colour calibration tags in DNG specify a conversion into a scene referred colourimetric space. This means that the rendering in this profile should include any desired tone and gamut mapping needed to convert between scene referred values and output referred values.

NOTE 2 DNG readers that have their own tone and gamut mapping controls should ignore AsShotICCProfile.

7.5 AsShotNeutral

Tag 50728 (C628.H)

Type SHORT or RATIONAL

Count ColourPlanes

Value See below

Mandatory No

Default None

Usage IFD 0

AsShotNeutral specifies the selected white balance at time of capture, encoded as the coordinates of a perfectly neutral colour in linear reference space values.

Each value shall be positive.

When the AsShotWhiteXY tag is present, the AsShotNeutral tag shall not be present.

7.6 AsShotPreProfileMatrix

Tag 50832 (C690.H)

Type SRATIONAL

Count 3 * ColourPlanes or

ColourPlanes * ColourPlanes

Value See below

Mandatory No

Default Identity matrix

Usage IFD 0

This tag is used in conjunction with the AsShotICCProfile tag. It specifies a matrix that should be applied to the camera colour space coordinates before processing the values through the ICC profile specified in the AsShotICCProfile tag.

The matrix shall be stored in the row scan order.

If ColourPlanes is greater than three, then this matrix may reduce the dimensionality of the colour data down to three components, in which case the AsShotICCProfile shall have three input components (rather than ColourPlanes input components).

7.7 AsShotProfileName

Tag 50934 (C6F6.H)

Type ASCII or BYTE

Count Length of string including null

Value Null-terminated UTF8-encoded string

Mandatory No

Default None

Usage IFD 0

This string may contain the name of the "as shot" camera profile, if any. That is, the name of the camera profile set in the camera at capture time.

7.8 AsShotWhiteXY

Tag 50729 (C629.H)

Type RATIONAL

Count 2

Value See below

Mandatory No

Default None

Usage IFD 0

AsShotWhiteXY specifies the selected white balance at time of capture, encoded as x-y chromaticity coordinates.

The values shall be in the open interval (0, 1).

When the AsShotNeutral tag is present, the AsShotWhiteXY tag shall not be present.

7.9 BaselineExposure

Tag 50730 (C62A.H)

Type SRATIONAL

Count 1

Value See below

Mandatory No

Default 0.0

Usage IFD 0

BaselineExposure specifies a default exposure adjustment in ExposureValue (APEX) units. Positive values result in brighter default results, while negative values result in darker default results.

NOTE Camera models vary in the tradeoff they make between highlight headroom and shadow noise. Some leave a significant amount of highlight headroom during a normal exposure. This allows significant negative exposure compensation to be applied during raw conversion, but also means normal exposures will contain more shadow noise. Other models leave less headroom during normal exposures. This allows for less negative exposure compensation but results in lower shadow noise for normal exposures. Because of these differences, a raw converter needs to vary the zero point of its exposure compensation control from model to model.

7.10 BaselineExposureOffset

Tag 51109 (C7A5.H)

Type RATIONAL

Count 1

Value See below

Mandatory No

Default 0.0

Usage IFD 0 or Camera Profile IFD

Provides a way for colour profiles to increase or decrease exposure during raw conversion.

BaselineExposureOffset specifies the amount (in EV units) to add to the BaselineExposure tag during image rendering. For example, if the BaselineExposure value for a given camera model is +0.3, and the BaselineExposureOffset value for a given camera profile used to render an image for that camera model is -0.7, then the actual default exposure value used during rendering will be $+0.3 - 0.7 = -0.4$.

7.11 BaselineNoise

Tag 50731 (C62B.H)

Type RATIONAL

Count 1

Value See below

Mandatory No

Default 1.0

Usage IFD 0

BaselineNoise specifies the relative noise level of the camera model at a baseline ISO value of 100, compared to a reference camera model.

A value of 1.0 shall mean that the camera has the same noise level as the reference camera model.

Since noise levels tend to vary approximately with the square root of the ISO value, a raw converter can use this value, combined with the current ISO, to estimate the relative noise level of the current image.

NOTE It is recommended that cameras use the newer NoiseProfile tag instead, because the NoiseProfile tag describes noise levels more precisely. Both tags are intended to help DNG readers and raw conversion software set an appropriate default amount of noise reduction when processing.

7.12 BaselineSharpness

Tag 50732 (C62C.H)

Type RATIONAL

Count 1

Value See below

Mandatory No

Default 1.0

Usage IFD 0 or Enhanced IFD

BaselineSharpness specifies the relative amount of sharpening required for this camera model, compared to a reference camera model.

A value of 0 shall mean no sharpening is to be applied.

A value of 1 shall mean reference-level sharpening.

NOTE Camera models vary in the strengths of their anti-aliasing filters. Cameras with weak or no filters require less sharpening than cameras with strong anti-aliasing filters. This tag is intended to help readers and raw conversion software set an appropriate default amount of sharpening when processing raw images.

7.13 BayerGreenSplit

Tag 50733 (C62D.H)

Type LONG

Count 1

Value See below

Mandatory No

Default 0

Usage Raw IFD

BayerGreenSplit tag specifies, in relative units, how closely the values of the green pixels in the blue/green rows track the values of the green pixels in the red/green rows.

A value of zero means the two kinds of green pixels track closely, while a non-zero value means they sometimes diverge. The useful range for this tag is from 0 (no divergence) to about 5 000 (quite large divergence).

If the image in an IFD is not a Bayer CFA image, BayerGreenSplit shall not be present in that IFD.

NOTE The useful range for this tag is from 0 (no divergence) to about 5 000 (quite large divergence).

This tag is intended to help DNG renders and raw conversion software set an appropriate amount of Green-red / Green-blue balance correction when demosaicing Bayer images.

It is recommended that writers store a tag value of 0 when the absolute ratio of the two greens ($R = \max(|Gr/Gb|, |Gb/Gr|)$) measured in a linear gamma encoding does not exceed 4 %. This indicates that the two kinds of green pixels track closely and that green-green balance correction can safely be skipped without risk of artifacts in the demosaiced image.

It is recommended that writers store a tag value between 1 and 1 000 when the absolute ratio of the two greens ($R = \max(|Gr/Gb|, |Gb/Gr|)$) measured in a linear gamma encoding is greater than 4 % but does not