

Network Working Group
Request for Comments: 2301
Category: Standards Track

L. McIntyre
Xerox Corporation
S. Zilles
Adobe Systems, Inc.
R. Buckley
Xerox Corporation
D. Venable
Xerox Corporation
G. Parsons
Northern Telecom
J. Rafferty
Human Communications
March 1998

File Format for Internet Fax

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

This document describes the TIFF (Tag Image File Format) representation of image data specified by the ITU-T Recommendations for black-and-white and color facsimile. This file format specification is commonly known as TIFF-FX. It formally defines minimal, extended and lossless JBIG modes (Profiles S, F, J) for black-and-white fax, and base JPEG, lossless JBIG and Mixed Raster Content modes (Profiles C, L, M) for color and grayscale fax. These modes or profiles correspond to the content of the applicable ITU-T Recommendations. Files formatted according to this specification use the image/tiff MIME Content Type.

Table of Contents

- 1. INTRODUCTION.....4
 - 1.1. Scope.....5
 - 1.2. Approach.....5
 - 1.3. Overview of this draft.....5
- 2. TIFF and Fax.....7
 - 2.1. TIFF Overview.....7
 - 2.1.1. File Structure.....7
 - 2.1.2. Image Structure.....9
 - 2.1.3. TIFF File Structure for Fax Applications.....10
 - 2.2 TIFF Fields for All Fax Applications.....11
 - 2.2.1. TIFF Fields required for all fax modes.....12
 - 2.2.2. Additional TIFF Fields required for all fax modes.....13
 - 2.2.3. TIFF Fields recommended for all fax modes.....15
 - 2.2.4. New TIFF Fields recommended for fax modes.....16
- 3. Minimal Black-and-White Fax Mode.....18
 - 3.1. Overview.....18
 - 3.2. Required TIFF Fields.....18
 - 3.2.1 Baseline Fields.....18
 - 3.2.2 Extension Fields.....20
 - 3.2.3 New Fields.....20
 - 3.3. Recommended TIFF Fields.....20
 - 3.4. End of Line (EOL) and Return to Control (RTC).....20
 - 3.4.1 RTC Exclusion.....21
 - 3.5. File Structure.....22
 - 3.6. Minimal Black-and-White Mode Summary.....23
- 4. Extended Black-and-White Fax Mode.....24
 - 4.1. TIFF-F Overview.....25
 - 4.2. Required TIFF Fields.....26
 - 4.2.1. Baseline Fields.....26
 - 4.2.2. Extension Fields.....28
 - 4.2.3. New Fields.....29
 - 4.3. Recommended TIFF Fields.....29
 - 4.3.1. Baseline Fields.....29
 - 4.3.2. Extension Fields.....29
 - 4.3.3. New Fields.....29
 - 4.4. Technical Implementation Issues.....30
 - 4.4.1. Strips.....30
 - 4.4.2. Bit Order.....31
 - 4.4.3. Multi-Page.....31
 - 4.4.4. Compression.....31
 - 4.4.5. Example Use of Page-quality Fields.....32
 - 4.4.6. Practical Guidelines for Writing and Reading Multi-Page
TIFF-F Files.....33
 - 4.4.7. Use of TIFF-F for Streaming Applications.....34
 - 4.5. Implementation Warnings.....34
 - 4.5.1. Uncompressed Data.....34

- 4.5.2. Encoding and Resolution.....35
- 4.5.3. EOL byte-aligned.....35
- 4.5.4. EOL.....36
- 4.5.5. RTC Exclusion.....36
- 4.5.6. Use of EOFB for T.6 Compressed Images.....37
- 4.6. Example Use of TIFF-F.....37
- 4.7. Extended Black-and-white Fax Mode Summary.....37
- 5. Lossless JBIG Black-and-White Fax Mode.....39
 - 5.1. Overview.....40
 - 5.2. Required TIFF Fields.....40
 - 5.2.1. Baseline Fields.....40
 - 5.2.2. Extension Fields.....40
 - 5.2.3. New Fields.....41
 - 5.3. Recommended TIFF Fields.....41
 - 5.4. Lossless JBIG Black-and-White Mode Summary.....41
- 6. Base Color Fax Mode.....43
 - 6.1. Overview.....43
 - 6.2. Required TIFF Fields.....43
 - 6.2.1. Baseline Fields.....43
 - 6.2.2. Extension Fields.....45
 - 6.2.3. New Fields.....46
 - 6.3. Recommended TIFF Fields.....47
 - 6.4. Base Color Fax Mode Summary.....47
- 7. Lossless Color Mode.....49
 - 7.1. Overview.....50
 - 7.1.1. Color Encoding.....50
 - 7.1.2. JBIG Encoding.....50
 - 7.2. Required TIFF Fields.....51
 - 7.2.1. Baseline Fields.....51
 - 7.2.2. Extension Fields.....52
 - 7.2.3. New Fields.....53
 - 7.3. Recommended TIFF Fields.....53
 - 7.4. Lossless Color Fax Mode Summary.....53
- 8. Mixed Raster Content Mode.....55
 - 8.1 Overview.....55
 - 8.1.1. MRC 3-layer model.....55
 - 8.1.2. A TIFF Representation for the MRC 3-layer model.....56
 - 8.2. Required TIFF Fields.....58
 - 8.2.1. Baseline Fields.....58
 - 8.2.2. Extension Fields.....59
 - 8.2.3. New Fields.....60
 - 8.3. Recommended TIFF Fields.....62
 - 8.4. Rules and Requirements for Images.....62
 - 8.5. MRC Fax Mode Summary.....63
- 9. MIME content-type image/tiff.....66
 - 9.1 Refinement of MIME content-type image/tiff for Facsimile Applications.....66
- 10. Security Considerations.....67

11. References.....	67
12. Authors' Addresses.....	69
Annex A: Summary of TIFF Fields for Internet Fax	70
Annex B. IANA Registration for image/tiff Application Parameter Values used for facsimile.....	75
Full Copyright Statement.....	77

1. Introduction

This document describes the use of TIFF (Tag Image File Format) to represent the data content and structure generated by the current suite of ITU-T Recommendations for Group 3 facsimile. These Recommendations and the TIFF fields described here support the following facsimile modes or profiles:

- S: minimal black-and-white mode, using binary MH compression [T.4]
- F: extended black-and-white mode, using binary MH, MR and MMR compression [T.4, T.6]
- J: lossless JBIG black-and-white mode, with JBIG compression [T.85, T.82]
- C: lossy color and grayscale mode, using JPEG compression [T.42, T.81]
- L: lossless color and grayscale mode, using JBIG compression [T.43, T.82]
- M: mixed raster content mode [T.44], using a combination of existing compression methods

Each profile corresponds to the content of ITU-T Recommendations shown and is a subset of the full TIFF for facsimile specification.

Profile S describes a minimal interchange set of fields, which will guarantee that, at least, binary black-and-white images will be supported. Implementations are required to support this minimal interchange set of fields.

With the intent of specifying a file format for Internet Fax, this draft:

1. specifies the structure of TIFF files for facsimile data,
2. defines ITU fax-compatible values for existing TIFF fields,
3. defines new TIFF fields and values required for compatibility with ITU color fax.

This specification of TIFF for facsimile is known as TIFF-FX.

1.1 Scope

This document defines a TIFF-based file format specification for enabling standardized messaging-based fax over the Internet. It specifies the TIFF fields and field values required for compatibility with the existing ITU-T Recommendations for Group 3 black-and-white, grayscale and color facsimile. TIFF has historically been used for handling fax image files in applications such as store-and-forward messaging. Implementations that support this file format specification for import/export may elect to support it as a native format. This document recommends a TIFF file structure that is compatible with low-memory and page-level streaming implementations.

Unless otherwise noted, the current TIFF specification [TIFF] and selected TIFF Technical Notes [TTN1, TTN2] are the primary references for describing TIFF and defining TIFF fields. This document is the primary reference for defining TIFF field values for fax applications.

1.2 Approach

The basic approach to using TIFF for facsimile data is to insert the compressed fax image data in a TIFF file and use TIFF fields to encode the parameters that describe the image data. These fields will have values that comply with the ITU-T Recommendations. The MIME content type of the resulting file will be image/tiff, with an optional Application parameter [TIFF-REG]; see Section 9.

This approach takes advantage of TIFF features and structures that bridge the data formats and performance requirements of both legacy fax machines and host-based fax applications. TIFF constructs for pages, images, and strips allow a TIFF file to preserve the fax data stream structure and the performance advantages that come with it. A TIFF-based approach also builds on an established base of users and implementors and ensures backward compatibility with existing TIFF-based IETF proposals and work in progress for Internet fax.

1.3 Overview of this draft

Section 2 gives an overview of TIFF. Section 2.1 describes the structure of TIFF files, including general guidelines for structuring multi-page TIFF files. Section 2.2 lists the TIFF fields that are required or recommended for all fax modes. The TIFF fields used only by specific fax modes are described in Sections 3-8, which describe the individual fax modes. These sections also specify the ITU-compatible field values (image parameters) for each mode.


The full set of permitted fields of TIFF for facsimile are included in the current TIFF specification, Section 2 of this document and the sections on specific modes of facsimile operation. This document defines profiles of TIFF for facsimile, where a profile is a subset of the full set of permitted fields and field values of TIFF for facsimile.

Section 3 defines the minimal black-and-white facsimile mode (Profile S), which is required in all implementations. Section 4 defines the extended black-and-white fax mode (Profile F), which provides a standard definition of TIFF-F. Section 5 describes the lossless black-and-white mode using JBIG compression (Profile J). Section 6 defines the base color mode, required in all color implementations, for the lossy JPEG representation of color and grayscale facsimile data (Profile C). Section 7 defines the lossless JBIG color and grayscale facsimile mode (Profile L) and Section 8 defines the Mixed Raster Content facsimile mode (Profile M). Each of these sections concludes with a table summarizing the required and recommended fields for each mode and the values they can have.

Section 9 describes the MIME content type image/tiff and the use of the optional Application parameter in connection with TIFF for facsimile. Sections 10, 11, 12 and 13 give Security Considerations, the ISOC Copyright Notice, References and Authors' Addresses. Annex A gives a summary of the TIFF fields used or defined in this document and provides a convenient reference for implementors.

To implement only the minimal interchange black-and-white set of fields and values (Profile S), one need read only Sections 1, 2, 3, 9 and 10.

The following tree diagram shows the relationship among profiles and between profiles and coding methods.



A profile is based on a collection of ITU-T facsimile coding methods.

For example, Profile S, the minimal mode, is based on Modified Huffman (MH) compression, which are defined in ITU-T Rec. T.4. Profile F specifies Modified Read (MR) and Modified Modified Read (MMR) compressions, which are defined in ITU-T Rec. T.4 and T.6.

All implementations of TIFF for facsimile MUST implement Profile S, which is the root node of the tree. All color implementations of TIFF for facsimile MUST implement Profile C. The implementation of a particular profile MUST also implement those profiles on the path that connect it to the root node, and MAY optionally implement profiles not on the path connecting it to the root node. For example, an implementation of Profile M must also implement Profiles C and S, and may optionally implement Profile F, J or L. For another example, an implementation of Profile C must also implement Profile S, and may optionally implement Profile F or J.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [REQ].

2. TIFF and Fax

2.1. TIFF Overview

TIFF provides a means for describing, storing and interchanging raster image data. A primary goal of TIFF is to provide a rich environment within which applications can exchange image data. The current TIFF specification [TIFF] defines a commonly used, core set of TIFF fields known as Baseline TIFF. The current specification and TIFF Technical Notes 1 and 2 [TTN1, TTN2] define several TIFF extensions. The TIFF-based specification for fax applications uses a subset of Baseline TIFF fields, with selected extensions, as described in this document. In a few cases, this document defines new TIFF fields specifically for fax applications.

2.1.1. File Structure

TIFF is designed for raster images, which makes it a good match for facsimile documents, which are multi-page raster images. Each raster image consists of a number of rows or scanlines, each of which has the same number of pixels, the unit of sampling. Each pixel has at least one sample or component (exactly one for black-and-white images).

A TIFF file begins with an 8-byte image file header. The first two bytes describe the byte order used within the file. Legal values are "II" (0x4949) when bytes are ordered from least to most significant (little-endian), and "MM" (0x4D4D), when bytes are ordered from most

to least significant (big-endian) within a 16- or 32-bit integer. Either byte order can be used, except in the case of the minimal black-and-white mode, which SHALL use value "II". The next two bytes contain the value 42 that identifies the file as a TIFF file and is ordered according to the value in the first two bytes of the header. The last four bytes give the offset that points to the first image file directory (IFD). This and all other offsets in a TIFF file are with respect to the beginning of the TIFF file. An IFD can be at any location in the file after the header but must begin on a word boundary.

An IFD is a sequence of tagged fields, sorted in ascending order by tag value. An IFD consists of a 2-byte count of the number of fields, a sequence of field entries and a 4-byte offset to the next IFD. The fields contain information about the image and pointers to the image data. Each separate raster image in the file is represented by an IFD.

Each field entry in an IFD has 12 bytes and consists of a 2-byte Tag, 2 bytes identifying the field type (e.g. short, long, rational, ASCII), 4 bytes giving the count (number of values or offsets), and 4 bytes that either contain the offset to a field value stored outside the IFD, or, based on the type and count, the field value itself. Resolution and metadata such as dates, names and descriptions are examples of "long" field values that do not fit in 4 bytes and therefore use offsets in the field entry. Details are given in the TIFF specification [TIFF].


A TIFF file can contain more than one IFD, where each IFD is a subfile whose type is given in the NewSubfileType field. Multiple IFDs can be organized either as a linked list, with the last entry in each IFD pointing to the next IFD (the pointer in the last IFD is 0), or as a tree, using the SubIFDs field in the primary IFD [TTN1]. The SubIFDs field contains an array of pointers to child IFDs of the primary IFD.

Child IFDs describe related images, such as reduced resolution versions of the primary IFD image. The same IFD can point both to a next IFD and to child IFDs, and child IFDs can themselves point to other IFDs.

All fax modes represent a multi-page fax image as a linked list of IFDs, with a NewSubfileType field containing a bit that identifies the IFD as one page of a multi-page document. Each IFD has a PageNumber field, identifying the page number in ascending order, starting at 0 for the first page. While a Baseline TIFF reader is not

required to read any IFDs beyond the first, an implementation that reads the files that comply with this specification SHALL read multiple IFDs. Only the Mixed Raster Content fax mode, described in Section 8, requires the use of child IFDs.

The following figure illustrates the structure of a multi-page TIFF file.



2.1.2 Image Structure

An IFD stores an image as one or more strips, as shown in the preceding figure. A strip consists of 1 or more scanlines (rows) of raster image data in compressed form. An image may be stored in a single strip or may be divided into several strips, which would require less memory to buffer. (Baseline TIFF recommends about 8k bytes per strip, but existing fax usage is typically one strip per image.)

Each IFD requires three strip-related fields: StripOffsets, RowsPerStrip and StripByteCounts. The StripOffsets field is an array of pointers to the strip or strips that contain the actual image data. The StripByteCounts field gives the number of bytes in each strip after compression. TIFF requires that each strip, except the last, contain the same number of scanlines, which is given in the RowsPerStrip field. This document introduces the new StripRowCounts field that allows a variable number of scanlines per strip, which is required by the Mixed Raster Content fax mode (Section 8).

Image data is stored as uninterpreted, compressed image data streams within a strip. The formats of these streams follow the ITU-T Recommendations. The Compression field in the IFD indicates the type of compression, and other TIFF fields in the IFD describe image attributes, such as color encoding and spatial resolution. Compression parameters are stored in the compressed data stream, rather than in TIFF fields. This makes the TIFF representation and compressed data format specification independent of each another. This approach, modeled on [TTN2], allows TIFF to gracefully add new compression schemes as they become available.

Some attributes can be specified both in the compressed data stream and within a TIFF field. It is possible that the two values will differ. When this happens for values required to interpret the data stream, then the values in the data stream take precedence. For informational values that are not required to interpret the data stream, such as author name, then the TIFF field value takes precedence.

2.1.3 TIFF File Structure for Fax Applications

The TIFF specification has a very flexible file structure, which does not specify the ordering of IFDs, field values and image data in a file. Individual applications may require or recommend an ordering.

This specification recommends that when using a TIFF file for facsimile, A multi-page fax document SHOULD be represented as a linked list of IFDs. It also recommends that a TIFF file for facsimile SHOULD order pages in a TIFF file in the same way that they are ordered in a fax data stream. In a TIFF file, a page consists of several elements: one or more IFDs (including subIFDs), long field values that are stored outside the IFDs, and image data (in one or more strips).

The minimal black-and-white mode (Profile S) specifies a required ordering of pages and elements within a page (Section 3.5). The extended black-and-white mode (Profile F) provides guidelines for ordering pages and page elements (Section 4.4.6). Other profiles

SHOULD follow these guidelines. This recommendation is intended to simplify the implementation of TIFF writers and readers in fax applications and the conversion between TIFF file and fax data stream representations. However, for interchange robustness, readers SHOULD be prepared to read TIFF files whose structure is consistent with [TIFF], which supports a more flexible file structure than is recommended here.

This specification introduces an optional new GlobalParametersIFD field, defined in Section 2.2.4. This field has type IFD and indicates parameters describing the fax session. While it is often possible to obtain these parameters by scanning the file, it is convenient to make them available together in one place for fast and easy access. If the GlobalParametersIFD occurs in a TIFF file, it SHOULD be located in the first IFD, immediately following the 8-byte image file header.

2.2 TIFF Fields for All Fax Applications

The TIFF specification [TIFF] is organized as a baseline set and several extensions, including technical notes [TTN1, TTN2] that will be incorporated in the next release of TIFF. The baseline and extensions have required and optional fields.

Facsimile applications require (and recommend) a mixture of baseline and extensions fields, as well as some new fields that are not part of the TIFF specification and that are defined in this document. This sub-section lists the fields that are required or recommended for all modes. In particular, Section 2.2.1 lists the fields that are required by all modes and that have values that do not depend on the mode. Section 2.2.2 lists the fields that are required by all modes and that have values which do depend on the mode. Section 2.2.3 lists the fields that are recommended for all modes. Fields that are required or recommended by some but not all modes are given in the section (Section 3-8) that describes that mode. The sections for each fax mode have sub-sections for required and recommended fields; each sub-section organizes the fields according to whether they are baseline, extension or new.

The fields required for facsimile have only a few legal values, specified in the ITU-T Recommendations. Of these legal values, some are required and some are optional, just as they are required (mandatory) or optional in fax implementations that conform to the ITU-T Recommendations. The required and optional values are noted in the sections on the different fax modes.

This section describes the fields required or recommended by all fax modes. The pattern for the description of TIFF fields in this draft is:

```

FieldName(TagValueInDecimal) = allowable values.                TYPE
  WhetherRequiredByTIFForTIFFforFAX
  Count = (omitted if =1) = (if not in current spec but available)
  Explanation of the field, how it's used, and the values it can have.
  Default value, if any, as specified in [TIFF]

```

When a field's default value is the desired value, that field may be omitted from the relevant IFD unless specifically required by the text of this specification.

2.2.1. TIFF fields required for all fax modes

The TIFF fields listed in this section SHALL be used by all fax modes, but have field values that are not specified by the ITU standards, i.e. the fields do not depend on the mode. The next subsection lists the fields that SHALL be used by all fax modes, but which do have values specified by the ITU-specified or mode-specific values. Fields that SHALL be used by some but not all modes are given in the sections (3-8) which describe the modes that uses them.

```

ImageLength(257)                                                SHORT or LONG
  RequiredByTIFFBaseline
  Total number of scanlines in image.
  No default, must be specified.

```

```

PageNumber(297)                                                SHORT
  RequiredByTIFFforFAX, TIFFExtension
  Count = 2
  The first number represents the page number (0 for the first page);
  the second number is the total number of pages in the document. If
  the second value is 0, then the total page count is not available.
  No default, must be specified

```

```

RowsPerStrip(278)                                             SHORT or LONG
  RequiredByTIFFBaseline
  The number of scanlines per TIFF strip, except for the last strip.
  For a single strip image, this is the same as the value of the
  ImageLength field.
  Default = 2*32 - 1 (meaning all scanlines in one strip)

```

```

StripByteCounts(279)                                         SHORT or LONG
  RequiredByTIFFBaseline
  Count = number of strips
  For each strip, the number of bytes in that strip after compression.

```

No default, must be specified.

StripOffsets(273) SHORT or LONG
 RequiredByTIFFBaseline
 Count = number of strips
 For each strip, the byte offset from the beginning of the file to the start of that strip.
 No default, must be specified.

2.2.2 Additional TIFF fields required for all fax modes

The TIFF fields listed in this section SHALL be used by all fax modes, but the values associated with them depend on the mode being described and the associated ITU Recommendations. Therefore, only the fields are defined here; the values applicable to a particular fax mode are described in Sections 3-8. Fields that SHALL be used by some but not all modes are given in the section (3-8) describing the mode that uses them.

BitsPerSample(258) SHORT
 RequiredByTIFFBaseline
 Number of bits per image sample
 Default = 1 (field may be omitted if this is the value)

Compression(259) SHORT
 RequiredByTIFFBaseline
 Compression method used for image data
 Default = 1 (no compression, so may not be omitted for FAX)

FillOrder(266) SHORT
 RequiredByTIFFforFax
 The default bit order in Baseline TIFF per [TIFF] is indicated by FillOrder=1, where bits are not reversed before being stored. However, TIFF for Fax typically utilizes the setting of FillOrder=2, where the bit order within bytes is reversed before storage (i.e., bits are stored with the Least Significant Bit first).
 Default = 1 (field may be omitted if this is the value)
 Facsimile data appears on the phone line in bit-reversed order relative to its description in the relevant ITU compression Recommendation. Therefore, a wide majority of facsimile implementations choose this natural order for storage. Nevertheless, all readers conforming to this specification must be able to read data in both bit orders.

ImageWidth(256) SHORT or LONG
 RequiredByTIFFBaseline
 The number of pixels (columns) per scanline (row) of the image
 No default, must be specified.

NewSubFileType(254) LONG
RequiredByTIFFforFAX
A general indication of the kind of data contained in this IFD
Bit 1 is 1 if the image is a single page of a multi-page document.
Default = 0 (no subfile bits on, so may not be omitted for FAX)

PhotometricInterpretation(262) SHORT
RequiredByTIFFBaseline
The color space of the image data
No default, must be specified

ResolutionUnit(296) SHORT
RequiredByTIFFBaseline
The unit of measure for resolution. 2 = inch, 3 = centimeter;
Default = 2 (field may be omitted if this is the value)

SamplesPerPixel(277) SHORT
RequiredByTIFFBaseline
The number of color components per pixel; SamplesPerPixel is 1 for a
black-and-white, grayscale or indexed (palette) image.
Default =1 (field may be omitted if this is the value)

XResolution(282) RATIONAL
RequiredByTIFFBaseline
The horizontal resolution of the image in pixels per resolution
unit. The ITU-T Recommendations for facsimile specify a small number
of horizontal resolutions: 100, 200, 300, 400 pixels per inch, and
80, 160 pixels per centimeter (or 204, 408 pixels per inch). The
allowed XResolution values for each mode are given in the section
defining that mode. Per [T.4], it is permissible for applications to
treat the following XResolution values as being equivalent: <204,
200> and <400,408> in pixels/inch. These equivalencies were allowed
by [T.4] to permit conversions between inch and metric based
facsimile terminals.
TIFF for Facsimile Writers SHOULD express XResolution in inch based
units, for consistency with historical practice and to maximize
interoperability. See the table below for information on how to
convert from an ITU-T metric value to its inch based equivalent
resolution.
No default, must be specified

YResolution(283) RATIONAL
RequiredByTIFFBaseline
The vertical resolution of the image in pixels per resolution unit.
The ITU-T Recommendations for facsimile specify a small number of
vertical resolutions: 100, 200, 300, 400 pixels per inch, and 38.5,
77, 154 pixels per centimeter (or 98, 196, 391 pixels per inch). The
allowed YResolution values for each mode are given in the section

defining that mode. Per [T.4], it is permissible for applications to treat the following YResolution values as being equivalent: <98, 100>, <196, 200>, and <391, 400> in pixels/inch. These equivalencies were allowed by [T.4] to permit conversions between inch and metric based facsimile terminals. TIFF for Facsimile Writers SHOULD express YResolution in inch based units, for consistency with historical practice and to maximize interoperability. See the table below for information on how to convert from an ITU-T metric value to its inch based equivalent resolution. No default, must be specified

XResolution		YResolution	
ResolutionUnit =2 (inch)	ResolutionUnit =3 (cm)	ResolutionUnit =2 (inch)	ResolutionUnit =3 (cm)
100		100	
204 200	80	98 100	38.5
204 200	80	196 200	77
204	80	391	154
300		300	
408 400	160	391 400	154

2.2.3 TIFF fields recommended for all fax modes

The TIFF fields listed in this section MAY be used by all fax modes. However, Profile S writers (the minimal fax mode described in Section 3) SHOULD NOT use these fields. Recommended fields that are mode-specific are described in Sections 3-8.

DateTime(306) ASCII
 OptionalInTIFFBaseline
 Date/time of image creation in 24-hour format "YYYY:MM:DD HH:MM:SS".
 No default.

DocumentName(269) ASCII
 OptionalInTIFFExtension(DocumentStorageAndRetrieval)
 The name of the scanned document. This is a TIFF extension field,
 not a Baseline TIFF field.

No default.

ImageDescription(270) ASCII
 OptionalInTIFFBaseline
 A string describing the contents of the image.
 No default.

Orientation(274) = 1-8. SHORT
 OptionalInTIFFBaseline
 1: 0th row represents the visual top of the image; the 0th column represents the visual left side of the image. See the current TIFF spec [TIFF] for further values; Baseline TIFF only requires value=1. Default = 1.
 Note: It is recommended that a writer that is aware of the orientation will include this field to give a positive indication of the orientation, even if the value is the default. If the Orientation field is omitted, the reader SHALL assume a value of 1.

Software(305) ASCII
 OptionalInTIFFBaseline
 The optional name and release number of the software package that created the image.
 No default.

2.2.4 New TIFF fields recommended for fax modes

The new TIFF fields listed in this section MAY be used by all fax modes, but their support is not expected for the minimal fax mode described in Section 3. In addition, support for these new TIFF fields has not been included in historical TIFF-F readers described in Section 4 and [TIFF- F]. These fields describe "global" parameters of the fax session that created the image data. They are optional, not part of the current TIFF specification, and are defined in this document.

The first new field, GlobalParametersIFD, is an IFD that contains global parameters and is located in a Primary IFD.

GlobalParametersIFD (400) IFD
 An IFD containing global parameters. It is recommended that a TIFF writer place this field in the first IFD, where a TIFF reader would find it quickly.

Each field in the GlobalParametersIFD is a TIFF field that is legal in any IFD. Required baseline fields should not be located in the GlobalParametersIFD, but should be in each image IFD. If a conflict exists between fields in the GlobalParametersIFD and in the image IFDs, then the data in the image IFD shall prevail.

Among the GlobalParametersIFD entries is a new ProfileType field which generally describes information in this IFD and in the TIFF file.

ProfileType(401) LONG
The type of image data stored in this IFD.
0 = Unspecified
1 = Group 3 fax
No default

The following new global fields are defined in this document as IFD entries for use with fax applications.

FaxProfile(402) = 0 - 6. BYTE
The profile that applies to this file; a profile is subset of the full set of permitted fields and field values of TIFF for facsimile. The currently defined values are:
0: does not conform to a profile defined for TIFF for facsimile
1: minimal black & white lossless, Profile S
2: extended black & white lossless, Profile F
3: lossless JBIG black & white, Profile J
4: lossy color and grayscale, Profile C
5: lossless color and grayscale, Profile L
6: Mixed Raster Content, Profile M

CodingMethods(403) LONG
This field indicates which coding methods are used in the file. A bit value of 1 indicates which of the following coding methods is used:
Bit 0: unspecified compression,
Bit 1: 1-dimensional coding, ITU-T Rec. T.4 (MH - Modified Huffman),
Bit 2: 2-dimensional coding, ITU-T Rec. T.4 (MR - Modified Read),
Bit 3: 2-dimensional coding, ITU-T Rec. T.6 (MMR - Modified MR),
Bit 4: ITU-T Rec. T.82 coding, using ITU-T Rec. T.85 (JBIG),
Bit 5: ITU-T Rec. T.81 (Baseline JPEG),
Bit 6: ITU-T Rec. T.82 coding, using ITU-T Rec. T.43 (JBIG color),
Bits 7-31: reserved for future use
Note: There is a limit of 32 compression types to identify standard compression methods.

VersionYear(404) BYTE
Count: 4
The year of the standard specified by the FaxProfile field, given as 4 characters, e.g. '1997'; used in lossy and lossless color modes.

ModeNumber (405) BYTE
The mode of the standard specified by the FaxProfile field. A value of 0 indicates Mode 1.0; used in Mixed Raster Content mode.

3. Minimal Black-and-White Fax Mode

This section defines the minimal black-and-white subset of TIFF for facsimile. This subset is designated Profile S. All implementations of TIFF for facsimile SHALL support the minimal subset.

Black-and-white mode is the binary fax application most users are familiar with today. This mode is appropriate for black-and-white text and line art. Black-and-white mode is divided into two levels of capability. This section describes the minimal interchange set of TIFF fields that must be supported by all implementations in order to assure that some form of image, albeit black-and-white, can be interchanged. This minimum interchange set is a strict subset of the fields and values defined for the extended black-and-white mode (TIFF-F or Profile F) in Section 4, which describes extensions to the minimal interchange set of fields that provide a richer set of black-and-white capabilities.

3.1. Overview

The minimal interchange portion of the black-and-white facsimile mode supports 1-dimensional Modified Huffman (MH) compression, with the original Group 3 fax resolutions, commonly called "standard" and "fine."

To assure interchange, this mode uses the minimal set of fields, with a minimal set of values. There are no recommended fields in this mode. Further, the TIFF file is required to be "little endian," which means that the byte order value in the TIFF header is "II". This mode defines a required ordering for the pages in a fax document and for the IFDs and image data of a page. It also requires that a single strip contain the image data for each page; see Section 3.5. The image data may contain RTC sequences, as specified in Section 3.4.

3.2. Required TIFF Fields

Besides the fields listed in Section 2.2.1, the minimal black-and-white fax mode requires the following fields. The fields listed in Section 2.2.1 and the fields and fax-specific values specified in this sub-section must be supported by all implementations.

3.2.1 Baseline fields

BitsPerSample(258) = 1. SHORT
RequiredByTIFFBaseline
Binary data only.
Default = 1 (field may be omitted if this is the value)

Compression(259) = 3. SHORT
RequiredByTIFFBaseline
3 = 1- or 2- dimensional coding.
The value 3 is a TIFF extension value [TIFF]. The T4Options field must be specified and its value specifies that the data is encoded using the Modified Huffman (MH) encoding of [T.4].

FillOrder(266) = 2. SHORT
RequiredByTIFFBaseline
2 = Least Significant Bit first

NOTE: Baseline TIFF readers are only required to support FillOrder = 1, where the lowest numbered pixel is stored in the MSB of the byte. However, because many devices, such as modems, transmit the LSB first when converting the data to serial form, it is common for black-and-white fax products to use the second FillOrder =2, where the lowest numbered pixel is stored in the LSB. Therefore, this value is specified in the minimal black-and-white mode.

ImageWidth(256) = 1728. SHORT or LONG
RequiredByTIFFBaseline
This mode only supports a page width of 1728 pixels. This width corresponds to North American Letter and Legal and to ISO A4 size pages.
No default, must be specified.

NewSubFileType(254) = (Bit 1=1). LONG
RequiredByTIFFforFAX
Bit 1 is 1 if the image is a single page of a multi-page document.
Default = 0 (no subfile bits on, so may not be omitted for fax)

PhotometricInterpretation(262) = 0. SHORT
RequiredByTIFFBaseline
0 = pixel value 1 means black
No default, must be specified

ResolutionUnit(296) = 2. SHORT
RequiredByTIFFBaseline
The unit of measure for resolution. 2 = inch.
Default = 2 (field may be omitted if this is the value)

SamplesPerPixel(277) = 1. SHORT
RequiredByTIFFBaseline
The number of components per pixel; 1 for black-and-white
Default =1 (field may be omitted if this is the value)

XResolution(282) = 200, 204. RATIONAL
RequiredByTIFFBaseline

The horizontal resolution of the image is expressed in pixels per resolution unit. In pixels/inch, the allowed values are 200 and 204, which may be treated as equivalent. See Section 2.2.2 for inch-metric equivalency.

No default, must be specified

YResolution(283) = 98, 100, 196, 200. RATIONAL

RequiredByTIFFBaseline

The vertical resolution of the image is expressed in pixels per resolution unit. In pixels/inch, the allowed values are 98, 100, 196 and 200; 98 and 100 may be treated as equivalent, and 196 and 200 may be treated as equivalent. See Section 2.2.2 for inch-metric equivalency.

No default, must be specified

3.2.2 Extension fields

T4Options(292) = (Bit 0 = 0, Bit 1 = 0, Bit 2 = 0, 1) LONG

RequiredTIFFExtension (when Compression = 3)

Bit 0 = 0 indicates MH encoding.

Bit 1 must be 0

Bit 2 = 1 indicates that EOLs are byte aligned, = 0 EOLs not byte aligned

Default is all bits are 0 (applies when EOLs are not byte aligned)

Note: The T4Options field is required when the Compression field has a value of 3. Bit 0 of this field specifies the encoding used (MH only in this mode) and Bit 2 indicates whether the EOL codes are byte-aligned or not. If they are byte aligned, then fill bits have been added as necessary so that the End of Line (EOL) codes always end on byte boundaries. See Section 3.4 for details.

3.2.3. New Fields

None.

3.3. Recommended TIFF Fields

None.

3.4. End of Line (EOL) and Return to Control (RTC)

The handling of End of Line (EOL) codes and Return to Control (RTC) sequences illustrate the differences between conventional fax, which is bit and stream oriented, and TIFF, which is byte and file oriented. Conventional fax, Baseline TIFF and TIFF extensions for fax all handle EOLs and RTCs differently.

In conventional fax, an MH-compressed fax data stream for a page consists of the following sequence:

EOL, compressed data (first line), EOL, compressed data, ... ,
EOL, compressed data (last line), RTC (6 consecutive EOL codes)

Baseline TIFF does not use EOL codes or Return to Control (RTC) sequences for MH-compressed data. However, the TIFF extension field T4Options used in this specification for MH compression (Compression = 3) requires EOLs.

Furthermore, Bit 2 in the T4Options field indicates whether or not the EOL codes are byte aligned. If Bit 2 = 1, indicating the EOL codes are byte aligned, then fill bits have been added as necessary before EOL codes so that an EOL code always ends on a byte boundary, and the first bit of data following an EOL begins on a byte boundary. Without fill bits, an EOL code may end in the middle of a byte. Byte alignment relieves application software of the burden of bit-shifting every byte while parsing scan lines for line-oriented image manipulation (such as writing a TIFF file). Not all TIFF readers historically used for fax are able to deal with non-byte aligned data.

While TIFF extension requires EOL codes, TIFF in fax applications has traditionally prohibited RTC sequences. Implementations that want common processing and interfaces for fax data streams and Internet fax files would prefer that the TIFF data include RTC sequences.

To reconcile these differences, RTCs are allowed in cases where EOL codes are not byte aligned and no fill bits have been added to the data. This corresponds to situations where the fax data is simply inserted in a strip without being processed or interpreted. RTCs should not occur in the data when EOLs have been byte aligned. This is formally specified in the next sub-section.

3.4.1. RTC Exclusion

Implementations which wish to maintain strict conformance with TIFF and compatibility with the historical use of TIFF for fax SHOULD NOT include the RTC sequence when writing TIFF files. However, implementations which need to support "transparency" of T.4-generated image data MAY include RTCs when writing TIFF files if the flag settings of the T4Options field are set for non-byte aligned data, i.e. Bit 2 is 0. Implementors of TIFF readers should be aware that there are some existing TIFF implementations for fax that include the RTC sequence in MH image data. Therefore, minimal set readers MUST be able to process files which do not include RTCs and SHOULD be able to process files which do include RTCs.

3.5. File Structure


The TIFF header, described in Section 2.1.1, contains two bytes which describe the byte order used within the file. For the minimal black-and-white mode, these bytes SHALL have the value "II" (0x4949), denoting that the bytes in the TIFF file are in LSByte-first order (little-endian). The first or 0th IFD immediately follows the header, so that offset to the first IFD is 8. The headers values are shown in the following table:

Offset	Description	Value
0	Byte Order	0x4949 (II)
2	Identifier	42 decimal
4	Offset of 0th IFD	0x 0000 0008

The minimal black-and-white mode SHALL order IFDs and image data within a file as follows: 1) there SHALL be an IFD for each page in a multi-page fax document; (2) the IFDs SHALL occur in the same order in the file as the pages occur in the document; (3) the IFD SHALL precede the image data to which it has offsets; (4) the image data SHALL occur in the same order in the file as the pages occur in the document; (5) the IFD, the value data and the image data it has offsets to SHALL precede the next image IFD; and (6) the image data for each page SHALL be contained within a single strip.

As a result of (6), the StripOffsets field will contain the pointer to the image data. With two exceptions, the field entries in the IFD contain the field values instead of offsets to field values located outside the IFD. The two exceptions are the values for the XResolution and YResolution fields, both of which are type RATIONAL and require 2 4-byte numbers. These "long" field values SHALL be placed immediately after the IFD which contains the offsets to them, and before the image data pointed to by that IFD.

The effect of these requirements is that the IFD for the first page SHALL come first in the file after the TIFF header, followed by the long field values for XResolution and YResolution, followed by the image data for the first page, then the IFD for second page, etc. This is shown in the following figure. Each IFD is required to have a PageNumber field, which has value 0 for the first page, 1 for the second page, and so on.



Using this file structure may reduce the memory requirements in implementations. It also provides some support for streaming, in which a file can be processed as it is received and before the entire file is received.

3.6 Minimal Black-and-white Mode Summary

The table below summarizes the TIFF fields that comprise the minimal interchange set for black-and-white facsimile. The Baseline and Extension fields and field values MUST be supported by all implementations. For convenience in the table, certain fields which have a value that is a sequence of flag bits are shown taking integer values that correspond to the flags that are set. An implementation should test the setting of the relevant flag bits individually, however, to allow extensions to the sequence of flag bits to be appropriately ignored. (See, for example, T4Options below.)

Baseline Fields	Values
BitsPerSample	1
Compression	3: 1D Modified Huffman coding set T4Options = 0 or 4

FillOrder	2: least significant bit first
ImageWidth	1728
ImageLength	n: total number of scanlines in image
NewSubFileType	2: Bit 1 identifies single page of a multi-page document
PageNumber	n,m: page number n followed by total page count m
PhotometricInterpretation	0: pixel value 1 means black
ResolutionUnit	2: inch
RowsPerStrip	number of scanlines per strip = ImageLength, with one strip
SamplesPerPixel	1
StripByteCounts	number of bytes in TIFF strip
StripOffsets	offset from beginning of file to single TIFF strip
XResolution	204, 200 (pixels/inch)
YResolution	98, 196, 100, 200 (pixels/inch)
Extension Fields	
T4Options	0: MH coding, EOLs not byte aligned 4: MH coding, EOLs byte aligned

4. Extended Black-and-White fax mode

This section defines the extended black-and-white mode or Profile F of TIFF for facsimile. It provides a standard definition of what has historically been known as TIFF Class F and now TIFF-F. In doing so, it aligns this mode with current ITU-T Recommendations for black-and-white fax and with existing industry practice. Implementations of this profile include implementations of Profile S.

This section describes extensions to the minimal interchange set of fields (Profile S) that provide a richer set of black-and-white capabilities. The fields and values described in this section are a superset of the fields and values defined for the minimal interchange set in Section 3. In addition to the MH encoding, Modified READ (MR) and Modified Modified READ (MMR) encoding as described in [T.4] and [T.6] are supported.

Section 4.1 gives an overview of TIFF-F. Section 4.2 describes the TIFF fields that SHALL be used in this mode. Section 4.3 describes the fields that MAY be used in this mode. In the spirit of the original TIFF-F specification, Sections 4.4 and 4.5 discuss technical implementation issues and warnings. Section 4.6 gives an example use of TIFF-F. Section 4.7 gives a summary of the required and recommended fields and their values.

4.1 TIFF-F Overview

Though it has been in common usage for many years, TIFF-F has previously never been documented in the form of a standard. An informal TIFF-F document was originally created by a small group of fax experts led by Joe Campbell. The existence of TIFF-F is noted in [TIFF] but it is not defined. This document serves as the formal definition of the F application of [TIFF] for Internet applications. For ease of reference, the term TIFF-F will be used throughout this document as a shorthand for the extended black-and-white mode or profile of TIFF for facsimile.

Up until the TIFF 6.0 specification, TIFF supported various "Classes" which defined the use of TIFF for various applications. Classes were used to support specific applications. In this spirit, TIFF-F has been known historically as "TIFF Class F". Previous informal TIFF-F documents [TIFF-F0] used the "Class F" terminology. As of TIFF 6.0 [TIFF], the TIFF Class concept has been eliminated in favor of the concept of Baseline TIFF. Therefore, this document updates the definition of TIFF-F as the F profile of TIFF for facsimile, by using Baseline TIFF as defined in [TIFF] as the starting point and then adding the TIFF extensions to Baseline TIFF which apply for TIFF-F. In almost all cases, the resulting definition of TIFF-F fields and values remains consistent with those used historically in earlier definitions of TIFF Class F. Where some of the values for fields have been updated to provide more precise conformance with the ITU-T [T.4] and [T.30] fax recommendations, these differences are noted.

4.2. Required TIFF Fields

This section lists the required fields and the values they must have to be ITU-compatible. Besides the fields listed in Section 2.2.1, the extended black-and-white fax mode SHALL use the following fields.

4.2.1. Baseline fields

BitsPerSample(258) = 1. SHORT
 RequiredByTIFFBaseline
 Binary data only.
 Default = 1 (field may be omitted if this is the value)

Compression(259) = 3, 4. SHORT
 RequiredByTIFFBaseline
 3 = 1- or 2- dimensional coding, must have T4Options field This is a TIFF Extension value [TIFF].
 4 = 2-dimensional coding, ITU-T Rec. T.6 (MMR - Modified Modified Read, must have T6Options field) This is a TIFF Extension value.
 Default = 1 (and is not applicable; field must be specified)

NOTE: Baseline TIFF permits use of value 2 for Modified Huffman encoding, but data is presented in a form which does not use EOLs, and so TIFF for facsimile uses Compression=3 instead. See Sections 4.4.4, 4.5.1 and 4.5.2 for more information on compression and encoding.

FillOrder(266) = 1 , 2. SHORT
 RequiredByTIFFBaseline
 Profile F readers must be able to read data in both bit orders, but the vast majority of facsimile products store data LSB first, exactly as it appears on the telephone line.
 1 = Most Significant Bit first.
 2 = Least Significant Bit first

ImageWidth(256) SHORT or LONG
 RequiredByTIFFBaseline
 This mode supports the following fixed page widths: 1728, 2592, 3456 (corresponding to North American Letter and Legal, ISO A4 paper sizes), 2048, 3072, 4096 (corresponding to ISO B4 paper size), and 2432, 3648, 4864 (corresponding to ISO A3 paper size).
 No default; must be specified

NOTE: Historical TIFF-F did not include support for the following widths related to higher resolutions: 2592, 3072, 3648, 3456, 4096 and 4864. Historical TIFF-F documents also included the following values related to A5 and A6 widths: 816 and 1216. Per the most recent

version of [T.4], A5 and A6 documents are no longer supported in Group 3 facsimile, so the related width values are now obsolete. See section 4.5.2 for more information on inch/metric equivalencies and other implementation details.

NewSubFileType(254) = (Bit 1=1). LONG

RequiredByTIFFforFAX

Bit 1 is 1 if the image is a single page of a multi-page document.

Default = 0 (no subfile bits on, so may not be omitted for fax)

NOTE: Bit 1 is always set to 1 for TIFF-F, indicating a single page of a multi-page image. The same bit settings are used when TIFF-F is used for a one page fax image. See Section 4.4.3 for details on multi-page files.

PhotometricInterpretation(262) = 0, 1. SHORT

RequiredByTIFFBaseline

0 = pixel value 1 means black, 1 = pixel value 1 means white.

This field allows notation of an inverted or negative image.

No default, must be specified

ResolutionUnit(296) = 2, 3. SHORT

RequiredByTIFFBaseline

The unit of measure for resolution. 2 = inch, 3 = centimeter; TIFF-F has traditionally used inch-based measures.

Default = 2 (field may be omitted if this is the value)

SamplesPerPixel(277) = 1. SHORT

RequiredByTIFFBaseline

1 = monochrome, bilevel in this case (see BitsPerSample)

Default =1 (field may be omitted if this is the value)

XResolution(282) = 200, 204, 300, 400, 408 RATIONAL

RequiredByTIFFBaseline

The horizontal resolution of the image is expressed in pixels per resolution unit. In pixels/inch, the allowed values are: 200, 204, 300, 400, and 408. See Section 2.2.2 for inch-metric equivalency.

No default, must be specified

NOTE: The values of 200 and 408 have been added to the historical TIFF-F values, for consistency with [T.30]. Some existing TIFF-F implementations may also support values of 80 pixels/cm, which is equivalent to 204 pixels per inch. See section 4.5.2 for information on implementation details.

YResolution(283) = 98, 100, 196, 200, 300, 391, and 400 RATIONAL

RequiredByTIFFBaseline

The vertical resolution of the image is expressed in pixels per

resolution unit. In pixels/inch, the allowed values are: 98, 100, 196, 200, 300, 391, and 400 pixels/inch. See Section 2.2.2 for inch-metric equivalency. No default, must be specified

NOTE: The values of 100, 200, and 391 have been added to the historical TIFF-F values, for consistency with [T.30]. Some existing TIFF-F implementations may also support values of 77 and 38.5 (cm), which are equivalent to 196 and 98 pixels per inch respectively. See section 4.5.2 for more information on implementation details.

NOTE: Not all combinations of XResolution, YResolution and ImageWidth are legal. The following table gives the legal combinations and corresponding paper size [T.30].

XResolution x YResolution	ImageWidth		
200x100, 204x98 200x200, 204x196 204x391	1728	2048	2432
300 x 300	2592	3072	3648
408 x 391, 400 x 400	3456	4096	4864
	Letter, A4 Legal	B4	A3
	Paper Size		

4.2.2. Extension fields

T4Options(292) = (Bit 0 = 0 or 1, Bit 1 = 0, Bit 2 = 0 or 1) LONG
 RequiredTIFFExtension (when Compression = 3)

T4Options was also known as Group3Options in a prior version of [TIFF].

Bit 0 = 1 indicates MR encoding, = 0 indicates MH encoding.

Bit 1 must be 0

Bit 2 = 1 indicates that EOLs are byte aligned, = 0 EOLs not byte aligned

Default is all bits are 0 (applies when MH encoding is used and EOLs are not byte aligned EOLs) (See Section 3.2.2.)

The T4Options field is required when the Compression field has a value of 3. This field specifies the encoding used (MH or MR) and whether the EOL codes are byte-aligned or not. If they are byte

aligned, then fill bits have been added as necessary so that the End of Line (EOL) codes always end on byte boundaries See Sections 3.4, 4.5.3 and 4.5.4 for details.

T6Options(293) = (Bit 0 = 0, Bit 1 = 0). LONG
RequiredTIFFExtension (when Compression = 4)
Used to indicate parameterization of 2D Modified Modified Read compression. T6Options was also known as Group4Options in a prior version of [TIFF].
Bit 0 must be 0.
Bit 1 = 0 indicates uncompressed data mode is not allowed; = 1 indicates uncompressed data is allowed (see [TIFF]).
Default is all bits 0. For FAX, the field must be present and have the value 0. The use of uncompressed data where compression would expand the data size is not allowed for FAX.

NOTE: MMR compressed data is two-dimensional and does not use EOLs. Each MMR encoded image MUST include an "end-of-facsimile-block" (EOFB) code at the end of each coded strip; see Section 4.5.6.

4.2.3. New fields

None.

4.3. Recommended TIFF fields

4.3.1. Baseline fields

See Section 2.2.3.

4.3.2. Extension fields

See Section 2.2.3.

4.3.3. New fields

Three new, optional fields, used in the original TIFF-F description to describe page quality, are defined in this specification. The information contained in these fields is usually obtained from receiving facsimile hardware (if applicable). They SHOULD NOT be used in writing TIFF-F files for facsimile image data that is error corrected or otherwise guaranteed not to have coding errors. Some applications need to understand exactly the error content of the data. For example, a CAD program might wish to verify that a file has a low error level before importing it into a high-accuracy document. Because Group 3 facsimile devices do not necessarily perform error correction on the image data, the quality of a received page must be inferred from the pixel count of decoded scan lines. A

"good" scan line is defined as a line that, when decoded, contains the correct number of pixels. Conversely, a "bad" scan line is defined as a line that, when decoded, comprises an incorrect number of pixels.

BadFaxLines(326) SHORT or LONG
The number of "bad" scan lines encountered by the facsimile device during reception. A "bad" scanline is defined as a scanline that, when decoded, comprises an incorrect number of pixels. Note that $\text{PercentBad} = (\text{BadFaxLines}/\text{ImageLength}) * 100$
No default.

CleanFaxData(327) = 0, 1, 2. SHORT
Indicates if "bad" lines encountered during reception are stored in the data, or if "bad" lines have been replaced by the receiver.
0 = No "bad" lines
1 = "bad" lines exist, but were regenerated by the receiver,
2 = "bad" lines exist, but have not been regenerated.
No default.

NOTE: Many facsimile devices do not actually output bad lines. Instead, the previous good line is repeated in place of a bad line. Although this substitution, known as line regeneration, results in a visual improvement to the image, the data is nevertheless corrupted. The CleanFaxData field describes the error content of the data. That is, when the BadFaxLines and ImageLength fields indicate that the facsimile device encountered lines with an incorrect number of pixels during reception, the CleanFaxData field indicates whether these bad lines are actually still in the data or if the receiving facsimile device replaced them with regenerated lines.

ConsecutiveBadFaxLines(328) LONG or SHORT
Maximum number of consecutive "bad" scanlines received. The BadFaxLines field indicates only the quantity of bad lines.
No Default.

NOTE: The BadFaxLines and ImageLength data indicate only the quantity of bad lines. The ConsecutiveBadFaxLines field is an indicator of the distribution of bad lines and may therefore be a better general indicator of perceived image quality. See Section 4.4.5 for examples of the use of these fields.

4.4. Technical Implementation Issues

4.4.1 Strips

In general, TIFF files divide an image into "strips," also known as "bands." Each strip contains a few scanlines of the image. By using

strips, a TIFF reader need not load the entire image into memory, thus enabling it to fetch and decompress small random portions of the image as necessary.

The number of scanlines in a strip is described by the RowsPerStrip value and the number of bytes in the strip after compression by the StripByteCount value. The location in the TIFF file of each strip is given by the StripOffsets values.

Strip size is application dependent. The recommended approach for multi-page TIFF-F images is to represent each page as a single strip. Existing TIFF-F usage is typically one strip per page in multi-page TIFF-F files. See Sections 2.1.2 and 2.1.3.

4.4.2 Bit Order

The current TIFF specification [TIFF] does not require a Baseline TIFF reader to support FillOrder=2, i.e. lowest numbered 1-bit pixel in the least significant bit of a byte. It further recommends that FillOrder=2 be used only in special purpose applications.

Facsimile data appears on the phone line in bit-reversed order relative to its description in ITU-T Recommendation T.4. Therefore, a wide majority of facsimile applications choose this natural order for data in a file. Nevertheless, TIFF-F readers must be able to read data in both bit orders and support FillOrder values of 1 and 2.

4.4.3. Multi-Page

Many existing applications already read TIFF-F-like files, but do not support the multi-page field. Since a multi-page format greatly simplifies file management in fax application software, TIFF-F specifies multi-page documents (NewSubfileType = 2) as the standard case.

It is recommended that applications export multiple page TIFF-F files without manipulating fields and values. Historically, some TIFF-F writers have attempted to produce individual single-page TIFF-F files with modified NewSubFileType and PageNumber (page one-of-one) values for export purposes. However, there is no easy way to link such multiple single page files together into a logical multiple page document, so that this practice is not recommended.

4.4.4. Compression

In Group 3 facsimile, there are three compression methods which had been standardized as of 1994 and are in common use. The ITU-T T.4 Recommendation [T.4] defines a one-dimensional compression method

known as Modified Huffman (MH) and a two-dimensional method known as Modified READ (MR) (READ is short for Relative Element Address Designate). In 1984, a somewhat more efficient compression method known as Modified Modified READ (MMR) was defined in the ITU-T T.6 Recommendation [T.6]. MMR was originally defined for use with Group 4 facsimile, so that this compression method has been commonly called Group 4 compression. In 1991, the MMR method was approved for use in Group 3 facsimile and has since been widely utilized.

TIFF-F supports these three compression methods. The most common practice is the one-dimensional Modified Huffman (MH) compression method. This is specified by setting the Compression field value to 3 and then setting bit 0 of the T4Options field to 0. Alternatively, the two dimensional Modified READ (MR) method, which is much less frequently used in historical TIFF-F implementations, may be selected by setting bit 0 of the T4Options field to 1. The value of Bit 2 in this field is determined by the use of fill bits.

Depending upon the application, the more efficient two-dimensional Modified Modified Read (MMR) compression method from T.6 may be selected by setting the Compression field value to 4 and then setting the first two bits (and all unused bits) of the T6Options field to 0. More information to aid the implementor in making a compression selection is contained in Section 4.5.2.

Baseline TIFF also permits use of Compression=2 to specify Modified Huffman compression, but the data does not use EOLs. As a result, TIFF-F uses Compression=3 instead of Compression=2 to specify Modified Huffman compression.

4.4.5. Example Use of Page-quality Fields

Here are examples for writing the CleanFaxData, BadFaxLines, and ConsecutiveBadFaxLines fields:

1. Facsimile hardware does not provide page quality information: MUST NOT write page-quality fields.
2. Facsimile hardware provides page quality information, but reports no bad lines. Write only BadFaxLines = 0.
3. Facsimile hardware provides page quality information, and reports bad lines. Write both BadFaxLines and ConsecutiveBadFaxLines. Also write CleanFaxData = 1 or 2 if the hardware's regeneration capability is known.
4. Source image data stream is error-corrected or otherwise guaranteed to be error-free such as for a computer generated file: SHOULD NOT write page-quality fields.

TIFF Writers SHOULD only generate these fields when the image has been generated from a fax image data stream where error correction, e.g. Group 3 Error Correction Mode, was not used.

4.4.6. Practical Guidelines for Writing and Reading Multi-Page TIFF-F Files

Traditionally, historical TIFF-F has required readers and writers to be able to handle multi-page TIFF-F files. Based on the experience of various TIFF-F implementors, it has been seen that the implementation of TIFF-F can be greatly simplified if certain practical guidelines are followed when writing multi-page TIFF-F files.

The structure for a multi-page TIFF-F file will include one IFD per page of the document. In this case, this IFD will define the attributes for a single page. A second simplifying guideline is that the writer of TIFF-F files SHOULD present IFDs in the same order as the actual sequence of pages. (The pages are numbered within TIFF-F beginning with page 0 as the first page and then ascending (i.e. 0, 1, 2, ...)). However, any field values over 4 bytes will be stored separately from the IFD. TIFF-F readers SHOULD expect IFDs to be presented in page order, but be able to handle exceptions.

Per [TIFF], the exact placement of image data is not specified. However, the strip offsets for each strip of image are defined from within each IFD. Where possible, another simplifying guideline for the writing of TIFF-F files is to specify that the image data for each page of a multi-page document SHOULD be contained within a single strip (i.e. one image strip per fax page). The use of a single image strip per page is very useful for applications such as store and forward messaging, where the file is usually prepared in advance of the transmission, but other assumptions may apply for the size of the image strip for applications which require the use of "streaming" techniques (see section 4.4.7). In the event a different image strip size guideline has been used (e.g. constant size for image strips that may be less than the page size), this will immediately be evident from the values/offsets of the fields that are related to strips.

A third simplifying guideline is that each IFD SHOULD be placed in the TIFF-F file structure at a point which precedes the image which the IFD describes.

In addition, a fourth simplifying guideline for TIFF-F writers and readers is to place the actual image data in a physical order within the TIFF file structure which is consistent with the logical page order. In practice, TIFF-F readers will need to use the strip

offsets to find the exact physical location of the image data, whether or not it is presented in logical page order.

If the image data is stored in multiple strips, then the strips SHOULD occur in the file in the same order that the data they contain occurs in the facsimile transmission, starting at the top of the page.

TIFF-F writers MAY make a fifth simplifying guideline, in which the IFD, the value data and the image data to which the IFD has offsets precede the next image IFD. However, this guideline has been relaxed (writers MAY rather than SHOULD use it) compared to the other guidelines given here to reflect past practices for TIFF-F.

In the case of the minimal mode, which is also the minimal subset of Profile S, the SHOULD's and MAY's of these guidelines become SHALL's (see Section 3.5).

So, a TIFF-F file which is structured using the guidelines of this section will essentially be composed of a linked list of IFDs, presented in ascending page order, which in turn each point to a single page of image data (one strip per page), where the pages of image data are also placed in a logical page order within the TIFF-F file structure. (The pages of image data may themselves be stored in a contiguous manner, at the option of the implementor).

4.4.7. Use of TIFF-F for Streaming Applications

TIFF-F has historically been used for handling fax image files in applications such as store and forward messaging where the entire size of the file is known in advance. While TIFF-F may also possibly be used as a file format for cases such as streaming applications, assumptions may be required that differ from those provided in this section (e.g., the entire size and number of pages within the image are not known in advance). As a result, a definition for the streaming application of TIFF-F is beyond the scope of this document.

4.5. Implementation Warnings

4.5.1 Uncompressed data

TIFF-F requires the ability to read and write at least one-dimensional T.4 Huffman ("compressed") data. Uncompressed data is not allowed. This means that the "Uncompressed" bit in T4Options or T6Options must be set to 0.

4.5.2. Encoding and Resolution

Since two-dimensional encoding is not required for Group 3 compatibility, some historic TIFF-F readers have not been able to read such files. The minimum subset of TIFF-F REQUIRES support for one dimensional (Modified Huffman) files, so this choice maximizes portability. However, implementors seeking greater efficiency SHOULD use T.6 MMR compression when writing TIFF-F files. Some TIFF-F readers will also support two-dimensional Modified READ files. Implementors that wish to have the maximum flexibility in reading TIFF-F files should support all three of these compression methods (MH, MR and MMR).

For the case of resolution, almost all facsimile products support both standard (98 dpi) vertical resolution and "fine" (196 dpi) resolution. Therefore, fine-resolution files are quite portable in the real world.

In 1993, the ITU-T added support for higher resolutions in the T.30 recommendation including 200 x 200, 300 x 300, 400 x 400 in dots per inch based units. At the same time, support was added for metric dimensions which are equivalent to the following inch based resolutions: 391v x 204h and 391v x 408h. Therefore, the full set of inch-based equivalents of the new resolutions are supported in the TIFF-F writer, since they may appear in some image data streams received from Group 3 facsimile devices. However, many facsimile terminals and older versions of TIFF-F readers are likely to not support the use of these higher resolutions.

Per [T.4], it is permissible for applications to treat the following XResolution values as being equivalent: <204,200> and <400,408>. In a similar respect, the following YResolution values may also be treated as being equivalent: <98, 100>, <196, 200>, and <391, 400>. These equivalencies were allowed by [T.4] to permit conversions between inch and metric based facsimile terminals.

In a similar respect, the optional support of metric based resolutions in the TIFF-F reader (i.e. 77 x 38.5 cm) is included for completeness, since they are used in some legacy TIFF-F applications, but this use is not recommended for the creation of TIFF-F files by a writer.

4.5.3. EOL byte-aligned

The historical convention for TIFF-F has been that all EOLs in Modified Huffman or Modified READ data must be byte-aligned. However, Baseline TIFF has permitted use of non-byte-aligned EOLs by default, so that a large percentage of TIFF-F reader implementations support

both conventions. Therefore, the minimum subset of TIFF-F, or Profile S, as defined in Section 3 includes support for both byte-aligned and non-byte-aligned EOLs; see Section 3.2.2.

An EOL is said to be byte-aligned when Fill bits have been added as necessary before EOL codes such that EOL always ends on a byte boundary, thus ensuring an EOL-sequence of a one byte preceded by a zero nibble: xxxx0000 00000001.

Modified Huffman encoding encodes bits, not bytes. This means that the end-of-line token may end in the middle of a byte. In byte alignment, extra zero bits (Fill) are added so that the first bit of data following an EOL begins on a byte boundary. In effect, byte alignment relieves application software of the burden of bit-shifting every byte while parsing scan lines for line-oriented image manipulation (such as writing a TIFF file).

For Modified READ encoding, each line is terminated by an EOL and a one bit tag bit. Per [T.4], the value of the tag bit is 0 if the next line contains two dimensional data and 1 if the next line is a reference line. To maintain byte alignment, fill bits are added before the EOL/tag bit sequence, so that the first bit of data following an MR tag bit begins on a byte boundary.

4.5.4. EOL

As illustrated in FIGURE 1/T.4 in [T.4], facsimile documents encoded with Modified Huffman begin with an EOL, which in TIFF-F may be byte-aligned. The last line of the image is not terminated by an EOL. In a similar respect, images encoded with Modified READ two-dimensional encoding begin with an EOL, followed by a tag bit.

4.5.5. RTC Exclusion

Aside from EOLs, TIFF-F files have historically only contained image data. This means that applications which wish to maintain strict conformance with the rules in [TIFF] and compatibility with historical TIFF-F, SHOULD NOT include the Return To Control sequence (RTC) (consisting of 6 consecutive EOLs) when writing TIFF-F files. However, applications which need to support "transparency" of [T.4] image data MAY include RTCs if the flag settings of the T4Options field are set for non-byte aligned MH or MR image data. Implementors of TIFF readers should also be aware that there are some existing TIFF-F implementations which include the RTC sequence in MH/MR image data. Therefore, TIFF-F readers MUST be able to process files which do not include RTCs and SHOULD be able to process files which do include RTCs.

4.5.6 Use of EOFB for T.6 Compressed Images

TIFF-F pages which are encoded with the T.6 Modified Modified READ compression method MUST include an "end-of-facsimile-block" (EOFB) code at the end of each coded strip. Per [TIFF], the EOFB code is followed by pad bits as needed to align on a byte boundary. TIFF readers SHOULD ignore any bits other than pad bits beyond the EOFB.

4.6. Example Use of TIFF-F

The Profile F of TIFF (i.e. TIFF-F content) is a secondary component of the VPIM Message, as defined in [VPIM2]. Voice messaging systems can often handle fax store-and-forward capabilities in addition to traditional voice message store-and-forward functions. As a result, TIFF-F fax messages can optionally be sent between compliant VPIM systems, and may be rejected if the recipient system cannot deal with fax.

Refer to the VPIM Specification for proper usage of this content.

4.7. Extended Black-and-white Fax Mode Summary

Recommended fields are shown with an asterisk *.

Required fields or values are shown with a double asterisk **. If the double asterisk is on the field name, then all the listed values are required of implementations; if the double asterisks are in the Values column, then only the values suffixed with a double asterisk are required of implementations.

Baseline Fields	Values
BitsPerSample	1**
Compression	3**: 1D Modified Huffman and 2D Modified Read coding 4: 2D Modified Modified Read coding
DateTime*	{ASCII}: date/time in 24-hour format "YYYY:MM:DD HH:MM:SS"
FillOrder**	1: most significant bit first 2: least significant bit first

ImageDescription*	{ASCII}: A string describing the contents of the image.
ImageWidth	1728**, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864
ImageLength**	n: total number of scanlines in image
NewSubFileType	2**: Bit 1 identifies single page of a multi-page document
Orientation	1**-8, Default 1
PhotometricInterpretation**	0: pixel value 1 means black 1: pixel value 1 means white
ResolutionUnit**	2: inch 3: centimeter
RowsPerStrip**	n: number of scanlines per TIFF strip
SamplesPerPixel	1**
Software*	{ASCII}: name & release number of creator software
StripByteCounts**	<n>: number of bytes in TIFF strip
StripOffsets**	<n>: offset from beginning of file to each TIFF strip
XResolution	200, 204**, 300, 400, 408 (written in pixels/inch)
YResolution	98**, 196**, 100, 200, 300, 391, 400 (written in pixels/inch)
Extension Fields	

T4Options	0**: required if Compression is Modified Huffman, EOLs are not byte aligned 1: required if Compression is 2D Modified Read, EOLs are not byte aligned 4**: required if Compression is Modified Huffman, EOLs are byte aligned
T4Options (continued)	5: required if Compression is 2D Modified Read, EOLs are byte aligned
T6Options	0: required if Compression is 2D Modified Modified Read
DocumentName*	{ASCII}: name of scanned document
PageNumber**	n,m: page number followed by total page count
New Fields	
BadFaxLines*	number of "bad" scanlines encountered during reception
CleanFaxData*	0: no "bad" lines 1: "bad" lines exist, but were regenerated by receiver 2: "bad" lines exist, but have not been regenerated
ConsecutiveBadFaxLines*	Max number of consecutive "bad" lines received

5. Lossless JBIG Black-and-White Fax Mode

This section defines the lossless JBIG black-and-white mode or Profile J of TIFF for facsimile. Implementations of this profile are required to also implement Profile S.

The previous section described the extended interchange set of TIFF fields for black-and-white fax, which provided support for the MH, MR and MMR compression of black-and-white images. This section adds a mode with JBIG compression capability.

5.1. Overview

This section describes a black-and-white mode that uses JBIG compression. The ITU-T has approved the single-progression sequential mode of JBIG [T.82] for Group 3 facsimile. JBIG coding offers improved compression for halftoned originals. JBIG compression is used in accordance with the application rules given in ITU-T Rec. T.85 [T.85].

This mode is essentially the extended black-and-white mode with JBIG compression used instead of MH, MR or MMR.

5.2. Required TIFF Fields

This section lists the required fields and the values they must have to be ITU-compatible. Besides the fields listed in Section 2.2.1, the extended black-and-white fax mode requires the following fields.

5.2.1. Baseline fields

The TIFF fields that SHALL be used in this mode are the same as those described in Section 4.2.1 for the extended black-and-white mode, with two exceptions: the following text replaces the text in Section 4.2.1 for the Compression and FillOrder fields.

Compression(259) = 9. SHORT
RequiredByTIFFBaseline
9 = ITU-T Rec. T.82 coding, applying ITU-T Rec. T.85 (JBIG). This is a TIFF extension value.
Default = 1 (and is not applicable; field must be specified).

FillOrder(266) = 2. SHORT
RequiredByTIFFBaseline
2 = Pixels are arranged within a byte such that pixels with lower column values are stored in the lower-order bits of the bytes, i.e., least significant bit first (LSB).

NOTE: The JBIG coding of black-and-white image data in Profile J follows ITU-T Rec. T.85 [T.85], which specifies LSB first ordering within a byte. Note that Baseline TIFF readers are only required to support MSB first ordering or FillOrder = 1.

5.2.2. Extension fields

Same fields as those in Section 2.2.1.

5.2.3. New fields

None.

5.3. Recommended TIFF Fields

See Section 2.2.3 and 2.2.4.

5.4. Lossless JBIG Black-and-white Fax Mode Summary

Recommended fields are shown with an asterisk *.

Required fields or values are shown with a double asterisk **. If the double asterisk is on the field name, then all the listed values are required of implementations; if the double asterisks are in the Values column, then only the values suffixed with a double asterisk are required of implementations.

Baseline Fields	Values
BitsPerSample	1**
Compression	9** : JBIG coding
DateTime*	{ASCII}: date/time in 24-hour format "YYYY:MM:DD HH:MM:SS"
FillOrder**	1: most significant bit first 2: least significant bit first
ImageDescription*	{ASCII}: A string describing the contents of the image.
ImageWidth	1728**, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864
ImageLength**	n: total number of scanlines in image
NewSubFileType**	2: Bit 1 identifies single page of a multi-page document
Orientation	1**-8, Default 1

PhotometricInterpretation**	0: pixel value 1 means black 1: pixel value 1 means white
ResolutionUnit**	2: inch 3: centimeter
RowsPerStrip**	n: number of scanlines per TIFF strip
SamplesPerPixel**	1
Software*	{ASCII}: name & release number of creator software
StripByteCounts**	<n>: number of bytes in TIFF strip
StripOffsets**	<n>: offset from beginning of file to each TIFF strip
XResolution	200, 204**, 300, 400, 408 (written in pixels/inch)
YResolution	98**, 196**, 100, 200, 300, 391, 400 (written in pixels/inch)
Extension Fields	
DocumentName*	{ASCII}: name of document scanned
PageNumber**	n,m: page number followed by total page count
New Fields	
GlobalParametersIFD*	IFD: global parameters IFD
ProfileType*	n: type of data stored in file
FaxProfile*	n: ITU-compatible fax mode
CodingMethods*	n: compression algorithms used in file

6. Base Color Fax Mode

6.1. Overview

This section defines the lossy color mode or Profile C of TIFF for facsimile. Implementations of this profile are required to also implement Profile S.

This is the base mode for color and grayscale facsimile, which means that all applications that support color fax must support this mode. The basic approach is the lossy JPEG compression [T.4, Annex E; T.81] of L*a*b* color data [T.42]. Grayscale applications use the L* lightness component; color applications use the L*, a* and b* components.

This mode uses a new PhotometricInterpretation field value to describe the L*a*b* encoding specified in [T.42]. This encoding differs in two ways from the other L*a*b* encodings used in TIFF [TIFF, TTN1]: it specifies a different default range for the a* and b* components, based on a comprehensive evaluation of existing hardcopy output, and it optionally allows selectable range for the L*, a* and b* components.

6.2. Required TIFF Fields

This section lists the required fields, in addition to those given in Section 2.2.1, and the values they must support to be compatible with ITU-T Rec. T.42 and Annex E in ITU-T Rec. T.4.

6.2.1. Baseline Fields

- ImageWidth(256). SHORT or LONG
 This mode supports the following fixed page widths: 864, 1024, 1216, 1728, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864.
- NewSubFileType(254) = (Bit 1=1). LONG
 RequiredByTIFFforFAX
 Bit 1 is 1 if the image is a single page of a multi-page document.
 Default = 0 (no subfile bits on, so may not be omitted for fax)
- BitsPerSample(258) = 8, 12. SHORT
 Count = SamplesPerPixel
 The base color fax mode requires 8 bits per sample, with 12 as an option. 12 bits per sample is not baseline TIFF.
- Compression(259) = 7. SHORT
 Base color fax mode uses Baseline JPEG compression. Value 7 represents JPEG compression as specified in [TTN2].

FillOrder(266) = 1 , 2. SHORT
 RequiredByTIFFBaseline
 Profile C readers must be able to read data in both bit orders,
 but the vast majority of facsimile products store data LSB
 first, exactly as it appears on the telephone line.
 1 = Most Significant Bit first.
 2 = Least Significant Bit first

PhotometricInterpretation(262) = 10. SHORT
 Base color fax mode requires pixel values to be stored using the CIE
 L*a*b* encoding defined in ITU-T Rec. T.42. This encoding is
 indicated by the PhotometricInterpretation value 10, referred to as
 ITULAB. With this encoding, the minimum sample value is mapped to 0
 and the maximum sample value is mapped to (2^n - 1), i.e. the
 maximum value, where n is the BitsPerSample value. The conversion
 from unsigned ITULAB-encoded samples values to signed CIE L*a*b*
 values is determined by the Decode field; see Sec. 6.2.3

NOTE: PhotometricInterpretation values 8 and 9 specify encodings for
 use with 8-bit-per-sample CIE L*a*b* [TIFF] and ICC L*a*b* [TTN1]
 data, but they are fixed encodings, which use different minimum and
 maximum samples than the T.42 default encoding. As currently defined,
 they are not able to represent fax-encoded L*a*b* data.

ResolutionUnit(296) = 2, 3. SHORT
 The unit of measure for resolution. 2 = inch, 3 = centimeter;
 Default = 2 (field may be omitted if this is the value)

SamplesPerPixel(277) = 1, 3. SHORT
 1: L* component only, required in base color mode
 3: L*, a*, b* components
 Encoded according to PhotometricInterpretation field

XResolution(282) = 100, 200, 300, 400. RATIONAL

YResolution(283) = 100, 200, 300, 400. RATIONAL

The resolution of the image is expressed in pixels per resolution
 unit. In pixels per inch, allowed XResolution values are: 100, 200,
 300, and 400. The base color fax mode requires the pixels to be
 square, hence YResolution must equal XResolution. Base resolution is
 200 pixels per inch and SHALL be supported by all implementations of
 this mode. See Section 2.2.2 for inch-metric equivalency.

NOTE: Not all combinations of XResolution, YResolution and ImageWidth
 are legal. The following table gives the legal combinations for inch-
 based resolutions and the corresponding paper sizes [T.30].

XResolution x YResolution	ImageWidth		
100 x 100	864	1024	1216
200 x 200	1728	2048	2432
300 x 300	2592	3072	3648
400 x 400	3456	4096	4864
	Letter, A4 Legal	B4	A3
	Paper Size		

6.2.2 Extension Fields

The JPEG compression standard allows for the a*b* chroma components of an image to be subsampled relative to the L* lightness component. The extension fields ChromaSubSampling and ChromaPositioning define the subsampling. They are the same as YCbCrSubSampling and YCbCrPositioning in [TIFF], but have been renamed to reflect their applicability to other color spaces.

ChromaSubSampling(530).

SHORT

Count = 2

Specifies the subsampling factors for the chroma components of a L*a*b* image. The two subfields of this field, ChromaSubsampleHoriz and ChromaSubsampleVert, specify the horizontal and vertical subsampling factors respectively.

SHORT 0: ChromaSubsampleHoriz = 1, 2.

1: equal numbers of lightness and chroma samples horizontally,
2: twice as many lightness samples as chroma samples horizontally,

SHORT 1: ChromaSubsampleVert = 1, 2.

1: equal numbers of lightness and chroma samples vertically,
2: twice as many lightness samples as chroma samples vertically,

The default value for ChromaSubSampling is (2,2), which is the default for chroma subsampling in color fax [T.4, Annex E]. No chroma subsampling, i.e. ChromaSubSampling = (1,1), is an option for color fax

ChromaPositioning(531) = 1.

SHORT

Specifies the spatial positioning of chroma components relative to

the lightness component.

1: centered,

A value of 1 means chrominance samples are spatially offset and centered with respect to luminance samples. See the current TIFF specification under YcbCr positioning for further information.

Default = 1, which is what ITU-T T.4, Annex E specifies.

6.2.3. New Fields

Decode(433).

SRATIONAL

Count = 2 * SamplesPerPixel

Describes how to map image sample values into the range of values appropriate for the current color space. In general, the values are taken in pairs and specify the minimum and maximum output value for each color component. For the base color fax mode, Decode has a count of 6 values and maps the unsigned ITULAB-encoded sample values (Lsample, asample, bsample) to signed L*a*b* values, as follows:.

$$L^* = \text{Decode}[0] + \text{Lsample} \times (\text{Decode}[1] - \text{Decode}[0]) / (2^n - 1)$$

$$a^* = \text{Decode}[2] + \text{asample} \times (\text{Decode}[3] - \text{Decode}[2]) / (2^n - 1)$$

$$b^* = \text{Decode}[4] + \text{bsample} \times (\text{Decode}[5] - \text{Decode}[4]) / (2^n - 1)$$

where Decode[0], Decode[2] and Decode[4] are the minimum values for L*, a* and b*; Decode[1], Decode[3] and Decode[5] are the maximum values for L*, a* and b*; and n is the BitsPerSample, either 8 or 12. For example, when n=8, L*=Decode[0] when Lsample=0 and L*=Decode[1] when Lsample=255.

ITU-T Rec. T.42 specifies the ITULAB encoding in terms of a range and offset for each component, which are related to the minimum and maximum values as follows:

$$\text{minimum} = - (\text{range} \times \text{offset}) / 2^n - 1$$

$$\text{maximum} = \text{minimum} + \text{range}$$

The Decode field default values depend on the color space. For the ITULAB color space encoding, the default values correspond to the base range and offset, as specified in ITU-T Rec. T.42 [T.42]. The following table gives the base range and offset values for BitsPerSample=8 and 12, and the corresponding default minimum and maximum default values for the Decode field, calculated using the equations above when PhotometricInterpretation=10.

BitsPer -Sample	Component	ITU-T Rec. T.42 base values		Decode default values	
		Range	Offset	Min	Max
8	L*	100	0	0	100
	a*	170	128	-21760/255	21590/255
	b*	200	96	-19200/255	31800/255
12	L*	100	0	0	100
	a*	170	2048	-348160/4095	347990/4095
	b*	200	1536	-307200/4095	511800/4095

For example, when PhotometricInterpretation=10 and BitsPerSample=8, the default value for Decode is (0, 100, -21760/255, 21590/255, -19200/255, 31800/255).

6.3. Recommended TIFF Fields

See Sections 2.2.3. and 2.2.4.

6.4 Base Color Fax Mode Summary

Recommended fields are shown with an asterisk *

Required fields or values are shown with a double asterisk **. If the double asterisk is on the field name, then all the listed values are required of implementations; if the double asterisks are in the Values column, then only the values suffixed with a double asterisk are required of implementations.

Baseline Fields	Values
BitsPerSample	8**: 8 bits per color sample 12: optional 12 bits/sample
Compression**	7: JPEG
DateTime*	{ASCII}: date/time in 24-hour format "YYYY:MM:DD HH:MM:SS"

FillOrder**	1: most significant bit first 2: least significant bit first
ImageDescription*	{ASCII}: A string describing the contents of the image.
ImageWidth	864, 1024, 1216, 1728**, 2048 2432, 2592, 3072, 3456, 3648 4096, 4864
ImageLength**	n: total number of scanlines in image
NewSubFileType**	2: Bit 1 identifies single page of a multi-page document
Orientation	1**-8, Default 1
PhotometricInterpretation	10**: ITULAB
ResolutionUnit**	2: inch 3: centimeter
RowsPerStrip**	n: number of scanlines per TIFF strip
SamplesPerPixel	1**: L* (lightness) 3: LAB
Software*	{ASCII}: name & release number of creator software
StripByteCounts**	<n>: number of bytes in TIFF strip
StripOffsets**	<n>: offset from beginning of file to each TIFF strip
XResolution	100, 200**, 300, 400 (written in pixels/inch)
YResolution	100, 200**, 300, 400 (must equal XResolution)

+-----+-----+	
Extension Fields	
+-----+-----+	
DocumentName*	{ASCII}: name of scanned document
+-----+-----+	
PageNumber**	n,m: page number followed by total page count
+-----+-----+	
ChromaSubSampling	(1,1), (2, 2)** (1, 1): equal numbers of lightness and chroma samples horizontally and vertically (2, 2): twice as many lightness samples as chroma samples horizontally and vertically
+-----+-----+	
ChromaPositioning	1**: centered
+-----+-----+	
New Fields	
+-----+-----+	
Decode**	minL, maxL, mina, maxa, minb, maxb: minimum and maximum values for L*a*b*
+-----+-----+	
GlobalParametersIFD*	IFD: IFD containing global parameters
+-----+-----+	
ProfileType*	n: type of data stored in TIFF file
+-----+-----+	
FaxProfile*	n: ITU-compatible fax mode
+-----+-----+	
CodingMethods*	n: compression algorithms used in file
+-----+-----+	
VersionYear*	byte sequence: year of ITU std
+-----+-----+	

7. Lossless Color Mode

This section defines the lossless color mode or Profile L of TIFF for facsimile. Implementations of this profile are required to also implement Profiles S and C.

7.1. Overview

This mode, defined in [T.43], uses JBIG to losslessly code three types of color and grayscale images: one bit per color CMY, CMYK and RGB images; a palettized (i.e. mapped) color image; and continuous tone color and grayscale images. The last two are multi-level and use the L*a*b* encoding specified in [T.42].

7.1.1. Color Encoding

While under development, this mode was called T.Palette, as one of its major additions was palette or mapped color images. Baseline TIFF only allows RGB color maps, but ITU-T Rec. T.43 requires L*a*b* color maps, using the encoding specified in ITU-T Rec. T.42. Palette color images are expressed with indices (bits per sample) of 12 bits or less, or optionally 13 to 16 bits, per [T.43].

Enabling T.43 color maps in TIFF requires the extension field Indexed, defined in [TTN1], and the PhotometricInterpretation field value 10, defined in Section 6.2.1. The following table shows the corresponding PhotometricInterpretation, SamplesPerPixel, BitsPerSample and Indexed field values for the different T.43 image types.

Image Type	PhotometricInterpretation	Samples Per Pixel	Bits Per Sample	Indexed
RGB	2=RGB	3	1	0
CMY	5=CMYK	3	1	0
CMYK	5=CMYK	4	1	0
Palette	10=ITULAB	1	n	1
Grayscale	10=ITULAB	1	8, 12	0
Color	10=ITULAB	3	8, 12	0

7.1.2. JBIG Encoding

T.43 uses the single-progression sequential mode of JBIG, defined in ITU-T Rec. T.82. To code multi-level images using JBIG, which is a bi-level compression method, an image is resolved into a set of bit-planes, and each bit-plane is then JBIG compressed. For continuous tone color and grayscale images, Gray code conversion is used. The

Gray code conversion is part of the data stream encoding, and is therefore invisible to TIFF.

7.2. Required TIFF Fields

This section lists the required fields, in addition to those in Section 2.2.1, and the values they must have to be compatible with ITU-T Rec. T.43.

7.2.1. Baseline Fields

- ImageWidth(256). SHORT or LONG
 Same page widths as the base color mode; see Section 6.2.1.
- NewSubFileType(254) = (Bit 1=1). LONG
 RequiredByTIFFforFAX
 Bit 1 is 1 if the image is a single page of a multi-page document.
 Default = 0 (no subfile bits on, so may not be omitted for fax)
- BitsPerSample(258) = 1, 2-8, 9-16. SHORT
 Count = SamplesPerPixel
 RGB, CMY, CMYK: 1 bit per sample
 Continuous tone (L*a*b*): 2-8 bits per sample, 9-12 bits optional
 Palette color: 12 or fewer bits per sample, 13-16 bits optional
 Note: More than 8 bits per sample is not baseline TIFF.
- ColorMap(320). SHORT
 Count = 3 * number of sample values
 Lossless color fax mode supports palette-color (indexed) images where the single component value is used as an index into a full color lookup table stored in the ColorMap field. The sample value is encoded using the number of bits given by the BitsPerSample field value. However, per [T.43], the number of sample values may be less than 2**BitsPerSample. The color lookup table is only required to have as many entries as there are number of sample values. For palette-color images in lossless color fax mode, the ITULAB encoding with 8 or optionally 12 bits per color map value is supported. To utilize a color map, the TIFF Indexed field must be present. TIFF orders the color map values so that all the L* values come first, followed by all the a* values and then all the b* values. Because ITU-T Rec. T.43 specifies a "chunky" ordering with the L*a*b* components of the first value, followed by those of the second value, and so on, reproducing color map values from a fax data stream in a TIFF file requires reordering values.
- Compression(259) = 10. SHORT
 10: ITU-T Rec. T.43 representation, using ITU-T Rec. T.82 (JBIG) coding

FillOrder(266) = 1 , 2. SHORT
 RequiredByTIFFBaseline
 Profile F readers must be able to read data in both bit orders,
 but the vast majority of facsimile products store data LSB
 first, exactly as it appears on the telephone line.
 1 = Most Significant Bit first.
 2 = Least Significant Bit first

PhotometricInterpretation(262) = 2, 5, 10. SHORT
 2: RGB
 5: CMYK, including CMY
 10: ITULAB
 Image data may also be stored as palette color images, where pixel
 values are represented by a single component that is an index into a
 color map using the ITULAB encoding. This color map is specified by
 the ColorMap field. To use palette color images, set the
 PhotometricInterpretation to 10, SamplesPerPixel to 1, and Indexed to
 1. The color map is stored in the ColorMap field. See Section 7.1.1
 for further discussion on the color encoding.

ResolutionUnit(296) = 2, 3. SHORT
 The unit of measure for resolution. 2 = inch, 3 = centimeter;
 Default = 2 (field may be omitted if this is the value)

SamplesPerPixel(277) = 1, 3, 4. SHORT
 1: Palette color image, or L*-only if Indexed = 0 and
 PhotometricInterpretation is 10 (ITULAB).
 3: RGB, or L*a*b*, or CMY if PhotometricInterpretation is 5 (CMYK).
 4: CMYK.

XResolution(282) = 100, 200, 300, 400. RATIONAL
 YResolution(283) = 100, 200, 300, 400. RATIONAL
 The resolution of the image is expressed in pixels per resolution
 unit. In pixels per inch, allowed XResolution values are: 100, 200,
 300, and 400. The lossless color fax mode requires the pixels to be
 square, hence YResolution must equal XResolution. Base resolution is
 200 pixels per inch. See Section 2.2.2 for inch-metric equivalency.

7.2.2. Extension Fields

Indexed(364) = 0, 1. SHORT
 0: not a palette-color image
 1: palette-color image
 This field is used to indicate that each sample value is an index
 into an array of color values specified in the ColorMap field.
 Lossless color fax mode supports palette-color images with the
 ITULAB encoding. The SamplesPerPixel value must be 1.

7.2.3. New Fields

Decode(433)

SRATIONAL

Decode is used in connection with the ITULAB encoding of image data and color map values; see Section 6.2.3.

7.3. Recommended TIFF Fields

See Sections 2.2.3. and 2.2.4.

7.4. Lossless Color Fax Mode Summary

Recommended fields are shown with an asterisk *.

Baseline Fields	Values
BitsPerSample	1: Binary RGB, CMY(K) 8: 8 bits per color sample 9-16: optional
ColorMap	n: LAB color map
Compression	10: JBIG, per T.43
DateTime*	{ASCII}: date/time in the 24-hour format "YYYY:MM:DD HH:MM:SS"
FillOrder**	1: Most significant bit first 2: Least significant bit first
ImageDescription*	{ASCII}: A string describing the contents of the image.
ImageWidth	864, 1024, 1216, 1728**, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864
ImageLength**	n: total number of scanlines in image
NewSubFileType	2: Bit 1 identifies single page of a multi-page document
Orientation	1**-8, Default 1
PhotometricInterpretation	2: RGB 5: CMYK 10: ITULAB

ResolutionUnit	2: inch 3: centimeter
RowsPerStrip	n: number of scanlines per TIFF strip
SamplesPerPixel	1: L* (lightness) 3: LAB, RGB, CMY 4: CMYK
Software*	{ASCII}: name & release number of creator software
StripByteCounts	<n>: number of bytes in TIFF strip
StripOffsets	<n>: offset from beginning of file to each TIFF strip
XResolution	100, 200, 300, 400 (written in pixels/inch)
YResolution	equal to XResolution (pixels must be square)
Extension Fields	
DocumentName*	{ASCII}: name of scanned document
PageNumber	n,m: page number followed by total page count
Indexed	0: not a palette-color image 1: palette-color image
New Fields	
Decode	minL, maxL, mina, maxa, minb, maxb: minimum and maximum values for L*a*b*
GlobalParameters IFD*	IFD: global parameters IFD

ProfileType*	n: type of data stored in TIFF file
FaxProfile*	n: ITU-compatible fax mode
CodingMethods*	n:compression algorithms used in file
VersionYear*	byte sequence: year of ITU fax std

8. Mixed Raster Content Mode

This section defines the Mixed Raster Content mode or Profile M of TIFF for facsimile. Implementations of this profile are required to implement Profiles S and C, and may optionally implement Profiles F, J and L.

8.1. Overview

Unlike previous fax modes, which use a single coding method and spatial resolution for an entire fax page, the Mixed Raster Content mode [T.44] enables different coding methods and resolutions within a single page. For example, consider a page that contains black-and-white text, which is best coded with MMR or JBIG, a color bar chart, best coded with JBIG, and a scanned color image, best coded with JPEG. Similarly, while spatial resolution of 400 pixels per inch may be best for the black-and-white text, 200 pixel per inch is usually sufficient for a color image.

Rather than applying one coding method and resolution to all elements, MRC allows multiple coders and resolutions within a page. By itself, MRC does not define any new coding methods or resolutions. Instead it defines a 3-layer image model for structuring and combining the scanned image data. The MRC 3-layer model has been applied here using the TIFF format to yield a data structure which differs from [T.44] though it applies the same coding methods, uses the same compressed image data stream and is consistent with the TIFF principle of a single IFD per image.

8.1.1. MRC 3-layer model

The 3 layers of the MRC model are Foreground and Background, which are both multi-level, and Mask, which is bi-level. Each layer may appear only once on a page and is coded independently of the other two. In our earlier example, the black-and-white text could be in the Mask layer, the color chart in the Foreground layer, and the color image in the Background layer.

Each layer is an image and, when present, is represented by at least one IFD in a TIFF file. This is consistent with TIFF, which provides fields to define the attributes, such as resolution, image size, bits per sample, etc., of a single image or layer. The distribution of content among layers is determined by the writer, as is the choice of coding method, color encoding and spatial resolution for a layer.

The final image is obtained by using the Mask layer to select pixels from the other two layers. When the Mask layer pixel value is 1, the corresponding pixel from the Foreground layer is selected; when it is 0, the corresponding pixel from the Background layer is selected. Details are given in the Introduction of [T.44].

Not all pages, and not all parts of a page, require 3 layers. If there is only one layer present, then that layer is the primary image or IFD. If there is more than one layer, then the Mask must be one of the layers, in which case it is the primary image and it must be page size.

MRC allows a page to be split into strips, with a variable number of scanlines in a strip. A strip can have 1, 2 or 3 layers. A single, stripped layer may be stored as a single, stripped image in an IFD, e.g., all strips associated with the Background layer may be treated as a single image. Alternatively, each strip associated with a layer may be stored as a separate image or IFD, e.g., the Background layer can be composed of several images that are offset vertically with respect to the page. In this case, there can be no overlap between images associated with a single layer. According to [T.4] Annex G, strips having more than 1 layer SHOULD NOT be more than 256 lines in length unless the capability to receive longer strips has been negotiated.

Furthermore, color fax also requires the spatial resolutions of Background and Foreground images to be legal fax values that are also integer factors of the Mask image resolution. For example, if the Mask Layer resolution is 400 pixels per inch, then allowed resolutions for the Foreground and Background layers are 100, 200 or 400 pixels per inch; if the Mask is at 300 pixels per inch, then allowed values are 100 and 300. The Foreground and Background layer resolutions can be independently set.

8.1.2. A TIFF Representation for the MRC 3-layer model

In the TIFF representation of the 3-layer MRC model, each page is represented by a single IFD, called the Primary IFD, that represents the Mask layer (unless the Foreground or Background is the single layer present), and a set of child IFDs that are referenced through the SubIFDs extension field [TTN1]. To distinguish MRC-specific

SubIFDs from other SubIFDs, the NewSubFileType field MUST have Bit 4 ON, indicating an MRC-related IFD. A new ImageLayer field is also introduced that consists of two values that identify the layer (Foreground, Background, or Mask) and the order within the layer (first, second, ... image of the layer); see Section 8.2.3.


Because MRC allows strips with variable numbers of scanlines, a reader MUST support StripRowCounts field because a writer may use it in place of the RowsPerStrip field in this mode. The StripRowCounts field allows each layer, with a variable number of scanlines in each strip, to be represented by a single IFD, when the coding parameters are the same for all strips in the layer. The MRC standard [T.44] allows the Foreground and Background layers to have strips with different coding parameters. In this case, a separate IFD is required to represent the strips which use different coding parameters; see text in next paragraph. In all cases, the Mask layer is required to be represented by a single IFD and a single set of coding parameters.

The use of SubIFDs to store child IFDs is described in [TTN1]. An example is shown graphically below. The Primary IFD associated with page 1 (PrimaryIFD 0) points to page 2 (PrimaryIFD 1) with the nextIFD offset. The Primary IFD, corresponding to the Mask layer (ImageLayer=[2,1]), contains a SubIFDs field that points to a list of child IFDs. The first child IFD represents one image of the Background layer, i.e., ImageLayer=[1,1]. This child IFD points to the second child IFD via the nextIFD offset. This child represents the second Background layer image, ImageLayer=[1,2]. Finally, the second child points to the third child, which corresponds to the single Foreground layer image, ImageLayer=[3,1]. The next IFD offset associated with this Foreground image is 0, indicating no more child IFDs exist. Each primary IFD has the NewSubFileType set to 18, indicating the IFD is MRC-specific (bit 4) and that it is a single page of a multi-page document (bit 1). Each child IFD has the NewSubFileType set to 16, indicating the IFD is MRC-specific. The 'V' character should be read as a down-pointing arrow.

```

                                (nextIFD)
PRIMARY IFD 0 -----> PRIMARY IFD 1--> ...
  ImageLayer = [2,1]
  NewSubFileType = 18
  SubIFDs
    |
    v
  Child IFD
    ImageLayer = [1,1]
    NewSubFileType = 16
    |
    |(nextIFD)

```



In the example above, the SubIFDs field of the Primary IFD points to the first IFD in a list of child IFDs. TIFF allows the SubIFDs field to point to an array of IFDs, each of which can be the first of a list of IFDs. An MRC-enabled TIFF reader must scan all available child IFDs to locate and identify IFDs associated with MRC layers.

In the case where the Background or Foreground layers are described with multiple IFDs, the XPosition and YPosition TIFF fields specify the offset to the upper-left corner of the IFD with respect to the Mask layer; see Section 8.2.2. When there is only a single layer (Mask, Foreground, or Background), it is stored as the Primary IFD.

8.2. Required TIFF Fields

This section describes the TIFF fields required, in addition to those in Section 2.2.1, to represent MRC mode fax images. Since MRC mode stores fax data as a collection of images corresponding to layers or parts of layers, the coding methods, color encodings and spatial resolutions used by previous modes apply to MRC. Therefore, the descriptions here will typically reference the appropriate earlier section. Fields and values specific to MRC mode are pointed out.

8.2.1. Baseline Fields

ImageWidth(256). SHORT or LONG
 Same page widths as the base color mode; see Section 6.2.1.
 In the MRC mode, the width of a Foreground or Background image in the coded data stream may be less than the page width. In this case, the image width in the coded data stream is used to interpret the coded data, and the value of this field is used as the page width.

NewSubFileType(254) = 16, 18. LONG
 For MRC fax mode, the NewSubFileType field has two bits that are required.
 Bit 1 indicates a single page of a multi-page document and must be set for the Primary IFD;
 Bit 4 indicates MRC imaging model as described in ITU-T Recommendation T.44 [T.44], and must be set for Primary IFDs and all MRC-specific child IFDs.

BitsPerSample(258) = 1, 2-8, 9-16 SHORT
 Compression(259) = 3, 4, 7, 9, 10. SHORT
 SamplesPerPixel(277) = 1, 3, 4. SHORT

FillOrder(266) = 1, 2. SHORT
 RequiredByTIFFBaseline
 Profile F readers must be able to read data in both bit orders, but the vast majority of facsimile products store data LSB first, exactly as it appears on the telephone line.
 1 = Most Significant Bit first.
 2 = Least Significant Bit first

ResolutionUnit(296) = 2, 3. SHORT
 PhotometricInterpretation(262) = 0, 1, 2, 5, 10. SHORT
 For Mask layer, see Sections 4.2.1 and 5.2.1.
 For Foreground and Background layers, see Sections 6.2.1 and 7.2.1.

ColorMap(320). SHORT
 Count = 3 * (2**BitsPerSample)
 Used when Foreground or Background layer is a palette-color image; see Section 7.2.1.

XResolution(282) = 100, 200, 300, 400. RATIONAL
 YResolution(283) = 100, 200, 300, 400. RATIONAL
 The resolution of the image is expressed in pixels per resolution unit. In pixels per inch, allowed XResolution values for all layers are: 100, 200, 300, and 400. MRC color fax mode requires the pixels to be square, hence YResolution must equal XResolution for all layers. The resolution of Background and Foreground layers must each be an integer factor of the Primary image, which is the Mask layer, when it is present; see Section 8.4.
 See Section 2.2.2 for inch-metric equivalency.

8.2.2. Extension Fields

ChromaSubSampling(530). SHORT
 ChromaPositioning(531). SHORT
 For Foreground and Background layers, see Section 6.2.2.

Indexed(346) = 0, 1. SHORT
 For Foreground and Background layers: 1 indicates a palette-color image, see Section 7.2.2.

T4Options(292) = 0, 1, 4, 5. SHORT

T6Options(293) = 0. SHORT
 For Mask layer, see Section 4.2.2.

SubIFDs(330). IFD
 Count = number of child IFDs
 Each value is an offset from the beginning of the TIFF file to a child IFD [TTN1].

XPosition(286). RATIONAL

YPosition(287). RATIONAL

Specifies the horizontal and vertical offsets of the top-left of the IFD from the top-left of the Primary IFD in page resolution units. For example, if the Primary IFD is at 400 pixels per inch, and a foreground layer IFD is at 200 pixels per inch and located at pixel coordinate (345, 678) with respect to the Primary IFD, the XPosition value is 345/400 and the YPosition value is 678/400. Color fax does not currently allow overlap of any component images within a single layer. Default values for XPosition and YPosition are 0.

8.2.3. New Fields

Decode(433). SRATIONAL
 For Foreground and Background layers, see Section 6.2.3.

DefaultImageColor(434). SHORT
 Count = SamplesPerPixel
 In areas where no image data is available, a default color is needed to specify the color value. If the StripByteCounts value for a strip is 0, then the color for that strip must be defined by a default image color.

The DefaultImageColor field uses the same encoding as the image data, and its value is therefore interpreted using the PhotometricInterpretation, SamplesPerPixel, BitsPerSample, and Indexed fields. If the fax data stream requires a different encoding, then transferring the default color value between a TIFF file and fax data stream requires a color conversion. For the Foreground layer image, the default value for the DefaultImageColor field is black. For other cases, including the Background layer image, the default value is white.

StripRowCounts(559). LONG

Count = number of strips

The number of scanlines stored in a strip. MRC allows each fax strip to store a different number of scanlines. For strips with more than one layer there is a maximum strip size of 256 scanlines or full page size. The 256 maximum SHOULD be used unless the capability to receive longer strips has been negotiated. This field replaces RowsPerStrip for IFDs with variable-sized strips. Only one of the two fields, StripRowCounts and RowsPerStrip, may be used in an IFD.

ImageLayer (34732). SHORT or LONG.

Count = 2

Image layers are defined such that layer 1 is the Background layer, layer 3 is the Foreground layer, and layer 2 is the Mask layer, which selects pixels from the Background and Foreground layers. The ImageLayer tag contains two values, describing the layer to which the image belongs and the order in which it is imaged.

ImageLayer[0] = 1, 2, 3.

- 1: Image is a Background image, i.e., the image that will appear whenever the Mask contains a value of 0. Background images typically contain low-resolution, continuous-tone imagery.
- 2: Image is the Mask layer. In MRC, if the Mask layer is present, it must be the Primary IFD and be full page in extent (no gaps.)
- 3: Image is a Foreground image, i.e., the image that will appear whenever the Mask contains a value of 1. The Foreground image generally defines the color of text or lines, but may also contain high-resolution imagery.

ImageLayer[1]:

- 1: first image to be imaged in this layer,
- 2: second image to be imaged in this layer,
- 3: ...

Value describing the image order. In MRC, this may be considered the strip number. Since MRC mode currently does not allow overlap between images within a layer, the order value does not have any visual effect.

In MRC fax mode, it is possible that only a single layer is transmitted. For example, if a page contains only a single continuous-tone photograph, then only the Background layer may be transmitted. In this case, the Background layer will be stored as the Primary IFD. ImageLayer[0] will be 1 indicating Background; ImageLayer[1] will be 1 since there can be no other IFDs associated with that layer. No Mask layer will exist.

8.3. Recommended TIFF Fields

See Sections 2.2.3. and 2.2.4.

8.4. Rules and Requirements for Images

The MRC mode defines a fundamental set of rules for images in the 3-layer representation.

1. If more than one layer exists, then the binary Mask layer SHALL be present and be the primary image. The Mask layer SHALL support the encoding defined in Section 3 and MAY support the encodings defined in Sections 4 and 5. If only one layer exists, then the image corresponding to that layer is the primary image.
2. When the binary Mask layer is the Primary IFD, the Primary IFD defines and extends to the entire page boundary; all attached model images cannot extend beyond the Primary image. Resolution differences may cause some pixels to "hang over" the page boundary, but no new pixels should exist completely beyond the page extent. When the Foreground or Background layer is the Primary IFD, the Primary IFD may not be page width.
3. The Background and Foreground images SHALL support the color encoding defined in Section 6 and MAY support the color encoding defined in Section 7. These images MAY optionally cover only a portion of the strip or page.
4. Each Primary IFD and each MRC-specific SubIFD must have an ImageLayer field to specify which layer the IFD belongs to, and the imaging order of that IFD within the layer.
5. Each Primary IFD must have a NewSubFileType field value set to 18, indicating a single page of a multi-page document (bit 1) and MRC mode (bit 4).
6. Each MRC-specific child IFD must have a NewSubFileType field value set to 16, indicating MRC mode (bit 4).
7. In MRC mode, each layer is transmitted as a sequence of strips. It is possible that each strip of each layer can be stored as a separate IFD. In this case, the SubIFDs structure pointed to by the Primary IFD will contain several IFDs that have an ImageLayer field with the layer identified as either Background (layer 1) or Foreground (layer 3). There may be no overlap in the vertical direction between IFDs associated with a single layer, although

there may be a gap from one of these images to the next. The TIFF XPosition and YPosition fields are used to indicate the placement of these images with respect to the primary image.

8. When the Mask image is present, the resolution of Background and Foreground images must each be an integer factor of the Mask image. For example, if the Mask image is 400 pixels/inch, then the Background or Foreground image may be at 400 pixels/inch (400/1), 200 pixels/inch (400/2) or 100 pixels/inch (400/4).

8.5. MRC Fax Mode Summary

Recommended fields are shown with an asterisk *

Baseline Fields	Values
BitsPerSample	1: binary mask 8: 8 bits per color sample 9-16: optional 12 bits/sample
ColorMap	n: LAB color map
Compression	3: Modified Huffman and Modified Read 4: Modified Modified Read 7: JPEG 9: JBIG, per T.85 10: JBIG, per T.43
DateTime*	{ASCII}: date/time in the 24-hour format "YYYY:MM:DD HH:MM:SS"
FillOrder**	1: Most significant bit first 2: Least significant bit first
ImageDescription*	{ASCII}: A string describing the contents of the image.
ImageWidth	864, 1024, 1216, 1728**, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864
ImageLength**	n: total number of scanlines in image
NewSubFileType	16, 18: Bit 1 indicates single page of a multi- page document on Primary IFD Bit 4 indicates MRC model

Orientation	1**-8, Default 1
PhotometricInterpretation	0: WhiteIsZero 1: BlackIsZero 2: RGB 5: CMYK 10: ITULAB
ResolutionUnit	2: inch 3: centimeter
RowsPerStrip	n: number or scanlines per strip
SamplesPerPixel	1: L* (lightness) 3: RGB, LAB, CMY 4: CMYK
Software*	{ASCII}: name & release number of creator software
StripByteCounts	<n>: number or bytes in each strip
StripOffsets	<n>: offset from beginning of file to each TIFF strip
XResolution	100, 200, 300, 400 (written in pixels/inch)
YResolution	equal to XResolution (pixels must be square)
Extension Fields	
T4Options	0: required if Compression is Modified Huffman, EOLs not byte aligned 1: required if Compression 2D Modified Read, EOLs are not byte aligned 4: required if Compression Modified Huffman, EOLs byte aligned 5: required if Compression 2D Modified Read, EOLs are byte aligned
T6Options	0: required if Compression is 2D Modified Modified Read

DocumentName*	{ASCII}: name of scanned document
PageNumber	n,m: page number followed by total page count
ChromaSubSampling	(1,1), (2, 2)** (1, 1): equal numbers of lightness and chroma samples horizontally & vertically (2, 2): twice as many lightness samples as chroma horizontally and vertically
ChromaPositioning	1: centered
Indexed	0: not a palette-color image 1: palette-color image
SubIFDs	<IFD>: byte offset to fg/bg IFDs
XPosition	horizontal offset in primary IFD resolution units
YPosition	vertical offset in primary IFD resolution units
New Fields	
Decode	minL, maxL, mina, maxa, minb, maxb: minimum and maximum values for L*a*b*
DefaultImageColor	<n>: background color
StripRowCounts	<n>: number of scanlines in each strip
ImageLayer	n, m: layer number, imaging sequence (e.g., strip number)
GlobalParameters IFD*	IFD: global parameters IFD
ProfileType*	n: type of data stored in TIFF file
FaxProfile*	n: ITU-compatible fax mode
CodingMethods*	n: compression algorithms used in file
ModeNumber*	n: version of ITU fax standard

```

+-----+
| VersionYear*      | byte sequence: year of ITU fax standard |
+-----+

```

9. MIME content-type image/tiff

[TIFF-REG] describes the registration of the MIME content-type image/tiff to refer to TIFF encoded image data. When transported by MIME, the TIFF content defined by this document must be encoded within an image/tiff content type. In addition, an optional "application" parameter is defined for image/tiff to identify a particular application's subset of TIFF and TIFF extensions for the encoded image data, if it is known. Typically, this would be used to assist the recipient in dispatching a suitable rendering package to handle the display or processing of the image file.

9.1 Refinement of MIME content-type image/tiff for Facsimile Applications

Since this document defines facsimile specific profiles of TIFF, it is useful to note an appropriate application parameter for the image/tiff MIME content-type.

The two values of the image/tiff application parameter as defined for facsimile are shown below, separated by a comma:

```
faxbw, faxcolor
```

The "faxbw" application parameter is suitable for use by applications that can process one or more TIFF for facsimile profiles or subsets used for the encoding of black and white facsimile data.

The "faxcolor" application parameter is suitable for use by applications that can process one or more TIFF for facsimile profiles or subsets that can be used for the encoding of black and white, AND color facsimile data.

Since this document defines several profiles of TIFF for facsimile, the following rules should be followed when setting the application parameter value. For TIFF image data which is encoded for the profiles of TIFF for Facsimile that support black-and-white image data (Profiles S, F or J), applications which use one of these profiles or a subset should set the value of the application parameter to "faxbw". For TIFF image data which is encoded for the defined profiles of TIFF for Facsimile that support color image data (Profiles C, L or M), as well as black-and-white image data, applications which use one of these profiles or a subset should set the value of the application parameter to "faxcolor".

An example of the use of the image/tiff MIME Content-type with the application parameter set with the value 'faxbw' follows:

```
Content-type: image/tiff; application=faxbw
```

In this example, use of this parameter value will enable applications to identify the content as being within a profile or subset of TIFF for Facsimile that is suitable for encoding black and white image data, Before attempting to process the image data.

In a similar respect, an example of the image/tiff MIME Content-type with the application parameter setting suitable for handling a color subset or profile of TIFF for facsimile is shown below:

```
Content-type: image/tiff; application=faxcolor
```

10. Security Considerations

This document describes a file format for Internet fax, which is a series of profiles of TIFF for facsimile. As such, it does not create any security issues not already identified in [TIFF-REG], in its use of fields as defined in [TIFF]. There are also new TIFF fields defined within this specification, but they are of a purely descriptive nature, so that no new security risks are incurred.

Further, the encoding specified in this document does not in any way preclude the use of any Internet security protocol to encrypt, authenticate, or non-repudiate TIFF-encoded facsimile messages.

11. References

[REQ] Bradner, S, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[T.4] ITU-T Recommendation T.4, Standardization of group 3 facsimile apparatus for document transmission, October 1997

[T.6] ITU-T Recommendation T.6, Facsimile coding schemes and coding control functions for group 4 facsimile apparatus, November 1988

[T.30] ITU-T Recommendation T.30 - Procedures for Document Facsimile Transmission in the General Switched Telephone Network, June 1996

[T.42] ITU-T Recommendation T.42, Continuous-tone colour representation method for facsimile, February 1996

[T.43] ITU-T Recommendation T.43, Colour and gray-scale image representations using lossless coding scheme for facsimile, February 1997

[T.44] ITU-T Recommendation T.44, Mixed Raster Content (MRC), October 1997.

[T.81] ITU-T Recommendation T.81, Information technology - Digital compression and coding of continuous-tone still images - Requirements and guidelines, September 1992

[T.82] ITU-T Recommendation T.82, Information technology - Coded representation of picture and audio information - Progressive bi-level image compression, March 1995

[T.85] ITU-T Recommendation T.85, Application profile for Recommendation T.82 - Progressive bi-level image compression (JBIG coding scheme) for facsimile apparatus, August 1995

[TIFF] Tag Image File Format, Revision 6.0, Adobe Developers Association, June 3, 1992,
<ftp://ftp.adobe.com/pub/adobe/devrelations/devtechnotes/pdffiles/tiff6.pdf>

The TIFF 6.0 specification dated June 3, 1992 specification (c) 1986-1988, 1992 Adobe Systems Incorporated. All Rights Reserved.

[TIFF-FY] Parsons, G. and J. Rafferty, "Tag Image File Format (TIFF) - F Profile for Facsimile", RFC 2306, March 1998.

[TIFF-F0] TIFF Class F specification, Apr 28, 1990,
ftp://ftp.faximum.com/pub/documents/tiff_f.txt

[TIFF-REG] Parsons, G., Rafferty J. and S. Zilles, "Tag Image File Format (TIFF) - image/tiff MIME Sub-type Registration", RFC 2302, March 1998.

[TTN1] Adobe PageMaker 6.0 TIFF Technical Notes, Sept. 14, 1995,
<http://www.adobe.com/supportservice/devrelations/PDFS/TN/TIFFPM6.pdf>

[TTN2] Draft TIFF Technical Note 2, Replacement TIFF/JPEG specification, March 17, 1995,
<ftp://ftp.sgi.com/graphics/tiff/TTN2.draft.txt>

[VPIM2] Vaudreuil, G. and G. Parsons, "Voice Profile for Internet Mail - version 2", work in progress, <draft-ema-vpim-06.txt>

The ITU-T Recommendations are available at <http://www.itu.ch>.

12. Authors' Addresses

Lloyd McIntyre
Xerox Corporation
Mailstop PAHV-305
3400 Hillview Ave.
Palo Alto, CA 94304 USA
Voice: +1-650-813-6762
Fax: +1-650-845-2340
Email: lmcintyre@adoc.xerox.com

Robert Buckley
Xerox Corporation
Mailstop 0128-30E
800 Phillips Road
Webster, NY 14580, USA
Voice: +1-716-422-1282
Fax: +1-716-422-6117
Email: Rob_Buckley@wb.xerox.com

Glenn S. Parsons
Northern Telecom
P.O. Box 3511, Station C
Ottawa, ON K1Y 4H7, Canada
Phone: +1-613-763-7582
Fax: +1-613-763-2697
Email: Glenn.Parsons@Nortel.ca

Stephen Zilles
Adobe Systems Inc.
Mailstop W14
345 Park Avenue
San Jose, CA 95110-2704, USA
Voice: +1-408-536-4766
Fax: +1-408-536-4042
Email: szilles@adobe.com

Dennis Venable
Xerox Corporation
Mailstop 0128-27E
800 Phillips Road
Webster, NY 14580, USA
Voice: +1-716-422-8009
Fax: +1-716-422-6117
Email: venable@wrc.xerox.com

James Rafferty
Human Communications
12 Kevin Drive
Danbury, CT 06811-2901, USA
Phone: +1-203-746-4367
Fax: +1-203-746-4367
Email: Jrafferty@worldnet.att.net

Annex A: Summary of TIFF Fields for Internet Fax

This annex includes tables which list by mode the TIFF fields used in the proposed fax file format. The fields are organized into 3 categories:

- 1) TIFF Baseline Fields
- 2) TIFF Extension Fields
- 3) New Fields.

The tables include the allowed values for each fax mode. Entries other than explicit numbers are described by:

- n - single number
- n, m - 2 numbers
- a, b, c - 3 numbers
- r - rational number
- <n> - array of numbers
- - byte sequence
- {ASCII} - string
- IFD - IFD byte offset
- <IFD> - array of IFD byte offsets

A blank entry in the table indicates that the field is not used by that particular fax mode.

Table A.1 TIFF Baseline Fields

TIFF Field	Fax Mode/Profile					
	Minimal B&W	Extended B&W	JBIG B&W	Lossy Color	Lossless Color	Mixed Raster Content
	S	F	J	C	L	M
BitsPerSample	1	1	1	8, 12	1, 2-8 9-16	1, 2-8 9-16
ColorMap					<n>	<n>
Compression	3	3, 4	9	7	10	3, 4, 7 9,10
DateTime		{ASCII}	{ASCII}	{ASCII}	{ASCII}	{ASCII}
FillOrder	2	1, 2	1, 2	1, 2	1, 2	1,2

ImageDescription		{ASCII}	{ASCII}	{ASCII}	{ASCII}	{ASCII}
ImageLength	n	n	n	n	n	n
ImageWidth	1728	1728, 2048, 2432 2592, 3072, 3456 3648, 4096, 4864		864, 1024, 1216, 1728, 2048, 2432, 2592, 3072, 3456, 3648, 4096, 4864		
NewSubFileType	2	2	2	2	2	16, 18
Orientation	1	1-8	1-8	1-8	1-8	1-8
PhotometricInterpretation	0	0, 1	0, 1	10	2, 5, 10	0, 1, 2, 5, 10
ResolutionUnit	2	2, 3	2, 3	2, 3	2, 3	2, 3
RowsPerStrip	n	n	n	n	n	n
SamplesPerPixel	1	1	1	1, 3	1, 3, 4	1, 3, 4
Software		{ASCII}	{ASCII}	{ASCII}	{ASCII}	{ASCII}
StripByteCounts	n	<n>	<n>	<n>	<n>	<n>
StripOffsets	n	<n>	<n>	<n>	<n>	<n>
XResolution	204 200	200, 204, 300 400, 408		100, 200, 300, 400		
YResolution	98, 196 100, 200	98, 196, 100, 200 300, 391, 400		100, 200, 300, 400		

Table A.2 TIFF Extension Fields

TIFF Field	Fax Mode/Profile					
	Minimal B&W S	Extended B&W F	JBIG B&W J	Lossy Color C	Lossless Color L	Mixed Raster Content M
Chroma-Positioning				1		1
Chroma-SubSampling				<1, 1> <2, 2>		<1, 1> <2, 2>
Document-Name		{ASCII}	{ASCII}	{ASCII}	{ASCII}	{ASCII}
Indexed					0,1	0,1
Page-Number	n, m	n, m	n, m	n, m	n, m	n, m
SubIFDs						<IFD>
T4Options	0, 4	0, 1, 4, 5				0, 1, 4, 5
T6Options		0				0
XPosition						r
YPosition						r

Table A.3 New Fields

TIFF Field	Fax Mode/Profile					
	Minimal B&W S	Extended B&W F	JBIG B&W J	Lossy Color C	Lossless Color L	Mixed Raster Content M
BadFax-Lines		n				
CleanFax-Data		0, 1, 2				
Coding-Method			n	n	n	n
ConsecutiveBad-FaxLines		n				
Decode				<r>	<r>	<r>
Default-ImageColor						<n>
Fax-Profile			n	n	n	n
Global-ParametersIFD		IFD	IFD	IFD	IFD	IFD
Image-Layer						n, m
Mode-Number						n
Profile-Type			n	n	n	n

Strip- RowCounts						<n>
Version- Year						

Annex B. IANA Registration for image/tiff Application Parameter Values used for facsimile

To: IANA@isi.edu

Subject: Registration of new Application parameter values for image/tiff

MIME media type name: image/tiff

Optional parameters: Application

New Value(s): faxbw, faxcolor

Description of Use:

faxbw - The "faxbw" application parameter is suitable for use by applications that can process one or more TIFF for facsimile profiles or subsets used for the encoding of black-and-white facsimile data. The definition of the use of this value is contained in Section 9 of this document (TIFFPLUS).

Faxcolor - The "faxcolor" application parameter is suitable for use by applications that can process one or more TIFF for facsimile profiles or subsets that can be used for the encoding of black and white, AND color facsimile data. The definition of the use of this value is contained in Section 9 of this document (TIFFPLUS).

Security Considerations:

Security considerations related to use of the TIFF subsets described by the "faxbw" and "faxcolor" values of the Application parameter are identified in Section 10 of this document (TIFFPLUS).

Persons & email addresses to contact for further information:

Glenn W. Parsons (Glenn.Parsons@Nortel.ca)
James Rafferty (Jrafferty@worldnet.att.net)
Stephen Zilles (szilles@adobe.com)

Change Controller: Stephen Zilles

INFORMATION TO THE SUBMITTER:

The accepted registrations will be listed in the "Assigned Numbers" series of RFCs. The information in the registration form is freely distributable.

Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

