



Portable Document Format Reference Manual

Version 1.3

Adobe Systems Incorporated

March 11, 1999

March 11, 1999

© 1993-1999 Adobe Systems Incorporated. All rights reserved. Patents Pending.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written consent of the publisher. Any software referred to herein is furnished under license and may only be used or copied in accordance with the terms of such license. Printed in the United States of America.

This publication and the information herein is furnished AS IS, is subject to change without notice, and should not be construed as a commitment by Adobe Systems Incorporated. Adobe Systems Incorporated assumes no responsibility or liability for any errors or inaccuracies, makes no warranty of any kind (express, implied or statutory) with respect to this publication, and expressly disclaims any and all warranties of merchantability, fitness for particular purposes and noninfringement of third party rights.

PostScript is a registered trademark of Adobe Systems Incorporated. All instances of the name PostScript in the text are references to the PostScript language as defined by Adobe Systems Incorporated unless otherwise stated. The name PostScript also is used as a product trademark for Adobe Systems' implementation of the PostScript language interpreter.

Any references to a "PostScript printer," a "PostScript file," or a "PostScript driver" refer to printers, files, and driver programs (respectively) which are written in or support the PostScript language. The sentences in this book that use "PostScript language" as an adjective phrase are so constructed to reinforce that the name refers to the standard language definition as set forth by Adobe Systems Incorporated.

Adobe, Acrobat, the Acrobat logo, Adobe Garamond, Carta, Distiller, FrameMaker, Illustrator, Minion, Photoshop, the Photoshop logo, Poetica, PostScript, and the PostScript logo are registered trademarks of Adobe Systems Incorporated. AdobePS and InDesign are trademarks of Adobe Systems Incorporated. Apple, Macintosh, Mac, and TrueType are trademarks of Apple Computer, Inc., registered in the United States and other countries. ITC Stone and ITC Zapf Dingbats are registered trademarks of International Typeface Corporation. Helvetica and Times are registered trademarks of Linotype-Hell AG and/or its subsidiaries. Lucida is a registered trademark of Bigelow and Holmes. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Sun is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. UNIX is a registered trademark of The Open Group. All other brand or product names are the trademarks or registered trademarks of their respective holders.

Library of Congress Cataloging-in-Publication Data

Portable document format reference manual / Adobe Systems Incorporated.

p. cm.

Includes bibliographical references (p. 505) and index.

ISBN 0-201-62628-4

1. File organization (Computer science) 2. PostScript (Computer program language)

3. Text processing (Computer science) I. Adobe Systems.

QA76.9.F5P67 199393-8046

005.74—dc20CIP

1 2 3 4 5 6 7 8 9—MA—9796959493

Contents

Chapter 1: Introduction 9

- 1.1 About this book 9
- 1.2 Introduction to the Second Edition—PDF 1.1 11
- 1.3 Introduction to the Third Edition—PDF 1.2 12
- 1.4 Introduction to the Fourth Edition—PDF 1.3 13
- 1.5 Conventions used in this book 14
- 1.6 A note on syntax 15
- 1.7 Copyright permission to use PDF 15

Portable Document Format 17

Chapter 2: Overview 19

- 2.1 What is the Portable Document Format? 19
- 2.2 Using PDF 19
- 2.3 General properties 22
- 2.4 PDF and the PostScript language 25
- 2.5 Understanding PDF 26

Chapter 3: Coordinate Systems 29

- 3.1 Device space 29
- 3.2 User space 30
- 3.3 Text space 31
- 3.4 Character space 31
- 3.5 Image space 31
- 3.6 Form space 31
- 3.7 Pattern space 32
- 3.8 Relationships among coordinate systems 32
- 3.9 Transformations between coordinate systems 33

3.10 Transformation matrices 35

Chapter 4: Objects 37

4.1 Introduction 37

4.2 Booleans 37

4.3 Numbers 37

4.4 Strings and text 37

4.5 Names 39

4.6 Arrays 40

4.7 Dictionaries 40

4.8 Streams 41

4.9 The null object 53

4.10 Indirect objects 53

4.11 Object references 54

Chapter 5: File Structure 55

5.12 PDF files 55

5.13 Header 56

5.14 Body 56

5.15 Cross-reference table 57

5.16 Trailer 59

5.17 Incremental update 61

5.18 Encryption 62

Chapter 6: Document Structure 65

6.1 Introduction 65

6.2 Catalog 67

6.3 Pages tree 71

6.4 Page objects 73

6.5 Thumbnails 81

6.6 Annotations 82

6.7 Outline tree 104

6.8 Actions 107

6.9 Names dictionary 119

6.10 Info dictionary 120

6.11	Articles	121
6.12	File ID	123
6.13	Encryption dictionary	124
6.14	Acrobat Forms	129
6.15	Sounds	147
6.16	Web Capture	148
6.17	Logical Structure in PDF	161
	Chapter 7: Common Data Structures	183
7.1	Rectangle	183
7.2	Date	183
7.3	Destination	184
7.4	File specification	186
7.5	Resources dictionaries	195
7.6	ProcSets	198
7.7	Fonts	198
7.8	CIDFonts	210
7.9	Font encodings	213
7.10	CMaps	215
7.11	Font descriptors	222
7.12	Color spaces	230
7.13	XObjects	246
7.14	Functions	262
7.15	Extended graphics states	272
7.16	Halftones	275
7.17	Patterns	287
7.18	Smooth Shading	294
7.19	Property lists	313
7.20	Name tree	314
7.21	Number tree	316
	Chapter 8: Page Descriptions	319
8.1	Overview	319
8.2	Graphics state	321

8.3	Special Graphics State	322
8.4	General Graphics state	323
8.5	Color	330
8.6	Paths	333
8.7	Text state	339
8.8	External objects (XObject)	348
8.9	In-line image objects	348
8.10	Other operators	351
	Chapter 9: Linearized PDF	357
9.1	Introduction	357
9.2	Background and Assumptions	358
9.3	Linearized PDF document structure specification	360
9.4	Hint Tables	369
9.5	Access Strategies	377
	Optimizing PDF Files	381
	Chapter 10: General Techniques for Optimizing PDF Files	383
10.1	Use short names	383
10.2	Use direct and indirect objects appropriately	384
10.3	Take advantage of combined operators	385
10.4	Remove unnecessary clipping paths	385
10.5	Omit unnecessary spaces	386
10.6	Omit default values	386
10.7	Take advantage of Form XObjects	386
10.8	Limit the precision of real numbers	387
10.9	Write parameters only when they change	387
10.10	Don't draw outside the crop box	387
10.11	Consider target device resolution	388
10.12	Share resources	388
10.13	Store common Page attributes in the Pages object	388
10.14	Use strings for named destinations	389
	Chapter 11: Optimizing Text	391
11.1	Don't produce unnecessary text objects	391

- 11.2 Use automatic leading 392
- 11.3 Take advantage of text spacing operators 394
- 11.4 Don't replace spaces between words 395
- 11.5 Use the appropriate operator to draw text 395
- 11.6 Use the appropriate operator to position text 396
- 11.7 Remove text clipping 396
- 11.8 Consider target device resolution 398

Chapter 12: Optimizing Graphics 399

- 12.1 Use the appropriate color-setting operator 399
- 12.2 Defer path painting until necessary 399
- 12.3 Take advantage of the closepath operator 400
- 12.4 Don't close a path more than once 401
- 12.5 Don't draw zero-length lines 402
- 12.6 Make sure drawing is needed 402
- 12.7 Take advantage of rectangle and curve operators 402
- 12.8 Coalesce operations 403

Chapter 13: Optimizing Images 405

- 13.1 Preprocess images 405
- 13.2 Match image resolution to target device resolution 405
- 13.3 Use the minimum number of bits per color component 405
- 13.4 Take advantage of indexed color spaces 406
- 13.5 Use the DeviceGray color space for monochrome images 406
- 13.6 Use in-line images appropriately 407
- 13.7 Don't compress in-line images unnecessarily 407
- 13.8 Choose the appropriate filters 407
- 13.9 Use predefined spot functions 410

Chapter 14: Clipping 413

- 14.1 Clipping to a path 413
- 14.2 Clipping to text 415
- 14.3 Image masks 417

Appendix A: Example PDF Files 421

- A.1 Minimal PDF file 421

A.2	Simple text string	423
A.3	Simple graphics	425
A.4	Pages tree	428
A.5	Outline	431
A.6	Updated file	435
Appendix B: Summary of Page Marking Operators		443
Appendix C: Predefined Font Encodings		447
C.1	Predefined encodings sorted by character name	448
C.2	Predefined encodings sorted by character code	453
C.3	MacExpert encoding	459
Appendix D: Implementation Limits		463
Appendix E: Obtaining XUIDs and Technical Notes		467
Appendix F: PDF Name Registry		469
Appendix G: Compatibility and Implementation Notes		471
G.1	Version numbers	471
G.2	Dictionary keys	472
G.3	Implementation notes	473
Appendix H: Forms Data Format		485
H.1	File Structure	485
H.2	The FDF Catalog Object	487
H.3	Use of FDF	491
H.4	Sample FDF	492
H.5	FDF for annotations	492
Appendix I: ISO 639 Language Codes		495
Appendix J: ISO 3166 Country Codes		499
Bibliography		505
Index		509
Colophon		517

CHAPTER 1**Introduction**

This book describes the Portable Document Format (PDF), the native file format of the Adobe[®] Acrobat[®] family of products. The goal of these products is to enable users to exchange and view electronic documents easily and reliably, independent of the environment in which they were created. PDF relies on the imaging model of the PostScript[®] language to describe text and graphics in a device-independent and resolution-independent manner. To improve performance for interactive viewing, PDF defines a more structured format than that used by most PostScript language programs. PDF also includes objects, such as annotations and hypertext links, that are not part of the page itself but are useful for interactive viewing.

PDF files are built from a sequence of numbered objects similar to those used in the PostScript language. The text, graphics, and images that make up the contents of a page are represented using operators that are based on those in the PostScript language and that closely follow the Adobe Illustrator[®] 3.0 page description operators.

A PDF file is not a PostScript language program and cannot be directly interpreted by a PostScript interpreter. However, the page descriptions in a PDF file can be converted into a PostScript language program.

1.1 About this book

This book provides a description of the PDF file format, as well as suggestions for producing efficient PDF files. It is intended primarily for application developers who wish to produce PDF files directly. This book also contains enough information to allow developers to write applications that read and modify PDF files. While PDF is independent of any particular application, occasionally PDF features are best explained by the actions a particular application takes when it encounters that feature in a file. Similarly, [Appendix D](#) discusses some implementation limits in the Acrobat viewer applications, even though these limits are not part of the file format itself.

The PDF specification is independent of any particular implementation of a PDF *generator* (an application that creates PDF files) or *consumer* (an application that reads PDF files and interprets their contents). To provide guidance to implementors, however, *Implementation Notes* that accompany the specification and [Appendix G](#) describe the behavior of Acrobat viewers when they encounter the changes documented herein.

This book consists of two sections. The first section describes the file format and the second lists techniques for producing efficient PDF files. In addition, appendices provide example files, detailed descriptions of several predefined font encodings, a summary of PDF page-marking operators, and other information.

Readers are assumed to have some knowledge of the PostScript language, as described in the *PostScript Language Reference Manual, Third Edition* [1]. In addition, some understanding of fonts, as described in the *Adobe Type 1 Font Format* [5], is useful.

The first section of this book, Portable Document Format, includes Chapters 2 through 7 and describes the PDF file format.

[Chapter 2](#) describes the motivation for creating the PDF file format and provides an overview of its architecture. PDF is compared to the PostScript language.

[Chapter 3](#) discusses the coordinate systems and transformations used in PDF files. Because the coordinate systems used in PDF are very much like those used in the PostScript language, users with substantial background in the PostScript language may wish to read this chapter only as a review.

[Chapter 4](#) describes the types of objects used to construct documents in PDF files. These types are similar to those used in the PostScript language. Readers familiar with the types of objects present in the PostScript language may wish to read this chapter quickly as a reminder.

[Chapter 5](#) provides a description of the format of PDF files, how they are organized on disk, and the mechanism by which updates can be appended to a PDF file.

[Chapter 6](#) describes the way that a document is represented in a PDF file, using the object types presented in [Chapter 4](#).

[Chapter 7](#) defines the resources used in a PDF file, including fonts, color spaces, images, and others.

[Chapter 8](#) discusses the operators used in PDF files. These are the operators that actually make marks on a page. Many are similar to one or more PostScript language operators. Readers with PostScript language experience will quickly see the similarities.

The second section of this book, Optimizing PDF Files, includes Chapters 9 through 14 and describes techniques for producing efficient PDF files. Many of the techniques presented can also be used in the PostScript language. The techniques are broken down into four areas: text, graphics, images, and general techniques.

[Chapter 9](#) describes linearized PDF files. Linearized files have been organized to work efficiently in a network environment.

[Chapter 10](#) discusses general optimizations that may be used in a wide variety of situations in PDF files.

[Chapter 11](#) discusses optimizations for text.

[Chapter 12](#) discusses graphics optimizations.

[Chapter 13](#) discusses optimizations that may be used on sampled images.

Finally, [Chapter 14](#) contains techniques for using clipping paths to restrict the region in which drawing occurs and a technique using images to make efficient blends. (*Note: The section on blends has been removed in the Fourth Edition.*)

1.2 Introduction to the Second Edition—PDF 1.1

The second edition of this manual is a revision of the 1993 edition of *Portable Document Format Reference Manual*. It describes version 1.1 of the Portable Document Format.

Implementation note *PDF 1.1 is the native file format of the Adobe Acrobat 2.0 family of products.*

The PDF 1.1 specification, like the PDF 1.0 specification, defines a minimum interchange level of functionality. The Portable Document Format is an extensible format, which means that PDF files may contain objects not defined by this specification. Consumers are expected to implement correctly the semantics of objects that are specified by PDF 1.1 and, as gracefully as possible, to ignore any objects that they do not understand. [Appendix G](#) provides guidance on how a consumer should handle objects it does not understand.

Implementation note *Some Acrobat 2.0 and subsequent products provide an interface that supports plug-ins. These plug-ins can use and/or put private data objects within a PDF file. [Appendix G](#) indicates the kinds of private data that can be used and [Appendix F](#) defines a registry for this data. The registry can be used to avoid conflicts in identifying data from independent plug-ins.*

New features introduced in PDF 1.1 include the following:

- The ability to protect a document with a password and to restrict operations on a document.
- The ability to tie blocks of text together into “articles,” making reading easier.
- The generalization of link and bookmark destinations to “actions,” which include links to other PDF files and foreign files.
- The ability to define new annotation types and to provide additional attributes for existing types.
- The ability to specify default settings and actions when a document is opened.
- Device-independent color.
- An ID included in files to make it easier to verify that a file is the correct file, even under circumstances where the file’s name is incorrect (such as files on some networks).
- A binary option that allows files to be smaller.

- A new date format that allows programmatic comparison of dates.
- The ability to provide additional document information.

Note In PDF 1.1, dictionary key names are often one or two letters in order to conserve space in files. When these keys are described below, they are followed in parentheses by a more descriptive string. However, only the actual one- or two-letter name may be used in a PDF file.

1.3 Introduction to the Third Edition—PDF 1.2

The third edition is a revision of the March 1, 1996, edition of *Portable Document Format Reference Manual*. It describes version 1.2 of the Portable Document Format.

Implementation note PDF 1.2 is the native file format of the Adobe Acrobat 3.0 family of products.

New features introduced in PDF 1.2 include the following:

- Interactive elements with state, such as radio buttons and checkboxes.
- Support for playing movies (from external files) and sounds (either embedded in the PDF file or from external files).
- Interactive, fill-in forms, with PDF-based format for data that can be imported, exported, transmitted, and received from the Web.
- Support for Chinese, Korean, and Japanese text.
- Enhanced borders, highlights, and fully general appearances for annotations.
- Support for responding to mouse-events.
- Essentially unlimited number of hypertext links.
- Support for OPI (Open Prepress Interface).
- Advanced color features, such as halftone screens, transfer functions, patterns, and separation color spaces.

Note PDF is an evolving language, and new editions of this manual will be offered on an ongoing basis to document the changes. The most recent version will be available on the Adobe's Web site (<http://www.adobe.com>).

While many readers will use a printed copy of this manual, others will refer to it online. With that in mind, many changes have been made to the formatting, such as the choice and size of fonts, to make on-screen reading somewhat easier.

1.4 Introduction to the Fourth Edition—PDF 1.3

The fourth edition is a revision of the 1.2 version of the Portable Document Format of October 3, 1996. It describes version 1.3 of the Portable Document Format.

Implementation note *PDF 1.3 is the native file format of the Adobe Acrobat 4.0 family of products.*

PDF 1.3 adds support for the new features of the Adobe imaging model, as defined in PostScript 3, as well as other features. These include the following:

- A facility for storing files converted from the World Wide Web in PDF files. ([Section 6.16, “Web Capture”](#))
- Digital signatures. ([Section 6.14.12, “Signature field”](#))
- A facility for representing the logical structure of a document independent of its graphic structure. ([Section 6.17, “Logical Structure in PDF”](#))
- Alternate representations for a single image. ([Section 7.13.3, “Alternate images”](#))
- Embedding files of any type within a PDF document. ([Section 7.4.3, “Embedded File streams”](#))
- Support for JavaScript. ([Section 6.14.13.4, “JavaScript Action”](#))
- Enhanced page numbering. ([Section 6.2.2, “PageLabel dictionaries”](#))
- Information relating to trapping. ([Section 6.6.16, “PDF Trapping”](#))
- Support for left-to-right and right-to-left reading direction. ([Section 6.2.1, “Viewer preferences”](#))
- A facility for representing color separations. ([Section 6.4.2, “Separated PDF”](#))
- New types of functions. ([Section 7.14, “Functions”](#))
- New colorspaces: **DeviceN** and **ICCBased**. ([Section 7.12, “Color spaces”](#))
- Smooth shading. ([Section 7.18, “Smooth Shading”](#)) A section of Chapter 14 in previous editions, showing how to emulate blends with images, has been removed from the new edition, since blends are now built in to PDF.
- Masked images. ([Section 7.13.5, “Masking images by position”](#))
- Additional support for CID fonts. ([Section 7.8, “CIDFonts”](#))
- Data structures for efficiently mapping strings and numbers to PDF objects. ([Section 7.20, “Name tree”](#) and [Section 7.21, “Number tree”](#))

1.5 Conventions used in this book

Text styles are used to identify various operators, keywords, terms, and objects. Several formatting styles are used in this book:

- PDF operators, PDF keywords, the names of keys in dictionaries, and other predefined names are written in boldface. Examples are **Tf**, **stream**, **Type**, and **MacRomanEncoding**.
- Operands of PDF operators are written in an italic sans serif font. An example is *linewidth*.
- Object types are written with initial capital letters. An example is FontDescriptor.
- The first occurrence of terms and the boolean values *true* and *false* are written in italics. This style is also used for emphasis.
- Samples of code as it would appear in a PDF file are written in a monospaced red font. Examples are: `/MediaBox [0 0 612 792]` and `www.adobe.com`.
- Byte values are written as hex digits either preceded by 0x, as in 0x4F, or enclosed in angle brackets, as in <4F>.

Tables containing dictionary keys are normally organized with the **Type** and **Subtype** keys first, followed by any other keys that are required in the dictionary, followed by any optional keys.

Note The **Type** key can almost always be inferred from context. The operand of the **Tf** operator, for example, must be a Font object, so the **Type** key in a Font dictionary serves primarily as documentation and as information for error-checking. Starting with PDF 1.3, then, the **Type** key is no longer required unless stated in the table; if it is present, however, it must have the correct value. In addition, the value of the **Type** key in any dictionary, even in “private data,” must be either a name defined in this manual or a “registered name”; see [Appendix F](#) for details.

PDF 1.3

The key names in the tables, shown in bold, are sometimes followed by descriptive terms in parentheses, such as **S** (Subtype). Only the key name is used in the table; the descriptive term is not a synonym or alternate key.

Important changes and corrections to the previous edition of this manual are marked with change bars in the left margin (such as the one to the left of this paragraph). Most of the changes are related to the differences between versions of PDF. Those changes are marked with icons in the right margin:

This marks a section that specifically deals with version 1.0, normally indicating a feature that has been superseded in a later version.

PDF 1.0

This marks a section that specifically deals with new features in version 1.1.

PDF 1.1

This marks a section that specifically deals with new features in version 1.2.



This marks a section that specifically deals with new features in version 1.3.



This marks a section for which there is a compatibility note or implementation note in [Appendix G](#).



Generally, new features and attributes are ignored by older viewers. An attribute marked with a version icon may be new with that version, or may have been substantially redefined in that version.

1.6 A note on syntax

Throughout this book, Extended Backus–Naur Form (EBNF) notation is used to describe syntax, as in this example:

```
<xyz> ::=    abc <def> ghi |
             <k> j
```

A token enclosed in angle brackets names a class of document component, while plain text appears verbatim or with some obvious substitution. The grammar rules have two parts. The class name is on the left of the definition symbol (::=). In the example above, the class is *xyz*. On the right of the definition symbol is a set of one or more alternative forms that the class component might take in the document. A vertical bar (|) separates alternative forms.

The right side of the definition may be on one or more lines. With only a few exceptions, these lines do not correspond to lines in the file.

The notation [...] means that the items enclosed in brackets are optional. If an asterisk follows an item, that item may be repeated *zero* or more times. If a plus sign follows an item, that item must be repeated *one* or more times. Braces, { ... }, are used to group items.

When an operator appears in a EBNF specification, it is shorthand for the operator plus its operands. For example, when the operator **m** appears in a EBNF specification, it means *x y m*, where *x* and *y* are numbers.

Note that PDF is case-sensitive. Uppercase and lowercase letters are distinct.

1.7 Copyright permission to use PDF

The general idea of utilizing an interchange format for final-form documents is in the public domain. Anyone is free to devise his or her own set of unique commands and data structures that define an interchange format for final-form documents. Adobe owns the copyright in the data structures, operators, and the written specification for the particular interchange format called the Portable Document Format. These elements may not be copied without Adobe's permission.

Adobe will enforce its copyright. Adobe's intention is to maintain the integrity of the Portable Document Format as a standard. This enables the public to distinguish between the Portable Document Format and other interchange formats for final-form documents.

However, Adobe desires to promote the use of the Portable Document Format for information interchange among diverse products and applications. Accordingly, Adobe gives copyright permission to anyone to:

- Prepare files in which the file content conforms to the Portable Document Format.
- Write drivers and applications that produce output represented in the Portable Document Format.
- Write software that accepts input in the form of the Portable Document Format and displays the results, prints the results, or otherwise interprets a file represented in the Portable Document Format.
- Copy Adobe's copyrighted list of operators and data structures, as well as the PDF sample code and PostScript language Function definitions in the written specification, to the extent necessary to use the Portable Document Format for the above purposes.

The only condition on such copyright permission is that anyone who uses the copyrighted list of operators and data structures in this way must include an appropriate copyright notice.

This limited right to use the copyrighted list of operators and data structures does not include the right to copy the *Portable Document Format Reference Manual*, other copyrighted material from Adobe, or the software in any of Adobe's products which use the Portable Document Format, in whole or in part, nor does it include the right to use any Adobe patents.



Section I

Portable Document Format

:

March 11, 1999

CHAPTER 2**Overview**

Before examining the detailed structure of a PDF file, it is important to understand what PDF is and how it relates to the PostScript language. This chapter discusses PDF and its relationship to the PostScript language.

[Chapter 3](#) discusses the coordinate systems used to describe various components of a PDF file. Chapters [4](#) and [5](#) discuss the basic types of objects supported by PDF and the structure of a PDF file. Chapters [6](#), [7](#), and [8](#) describe the structure of a PDF document and the operators used to draw text, graphics, and images.

2.1 What is the Portable Document Format?

PDF is a file format used to represent a document in a manner independent of the application software, hardware, and operating system used to create it. A *PDF file* contains a *PDF document* and other supporting data.

A PDF document contains one or more pages. Each page in the document may contain any combination of text, graphics, and images in a device- and resolution-independent format. This is the *page description*. A PDF document may also contain information possible only in an electronic representation, such as hypertext links, sound, and movies.

In addition to a document, a PDF file contains the version of the PDF specification used in the file and information about the location of important structures in the file.

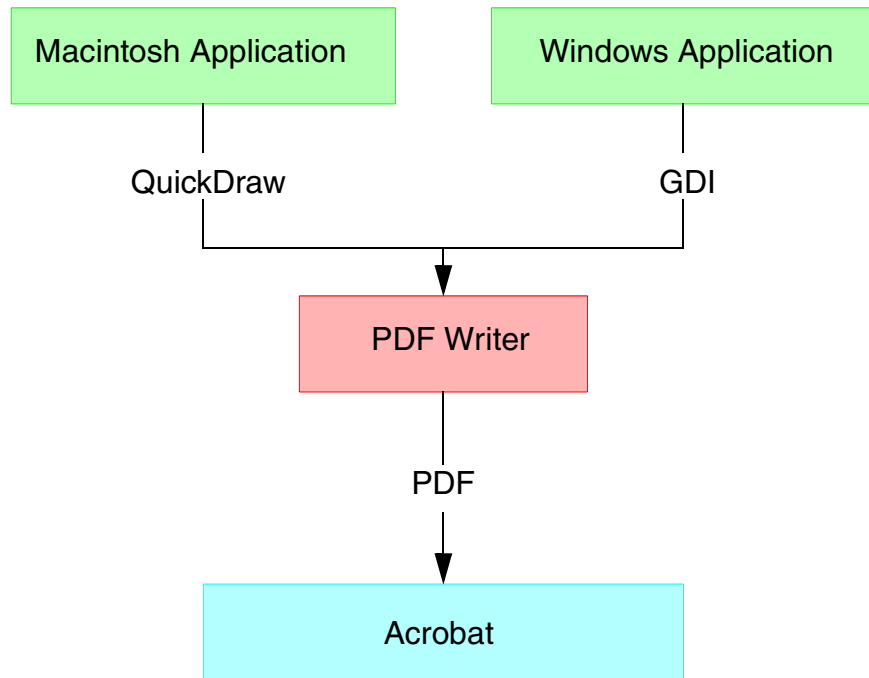
2.2 Using PDF

To understand PDF, it is important to understand how PDF documents will be produced and used. As PDF documents and applications that read PDF files become more prevalent, new ways of creating and using PDF files will be invented. This is one of the goals of this book—to make the file format accessible so that application developers can expand on the ideas behind PDF and the applications that initially support it.

Currently, PDF files may be produced either directly from applications or from files containing PostScript page descriptions.

Many applications can produce PDF files directly. The PDF Writer, available on both Apple® Macintosh® computers and computers running the Microsoft® Windows® environment, acts as a printer driver. A printer driver normally converts operating system graphics and text commands (QuickDraw™ for the Macintosh and GDI for Windows) into commands understood by a printer. The driver embeds these commands in a stream of commands sent to a printer that results in a page being printed. Instead of sending these commands to a printer, the PDF Writer converts them to PDF operators and embeds them in a PDF file, as shown in [Figure 2.1](#).

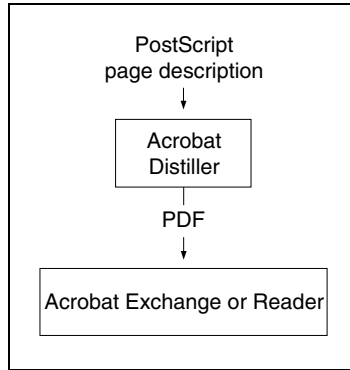
Figure 2.1 *Creating PDF files using PDF Writer*



The resulting PDF files are platform-independent. Regardless of whether they were generated on a Macintosh or Windows computer, they may be viewed by a PDF viewing application on any supported platform.

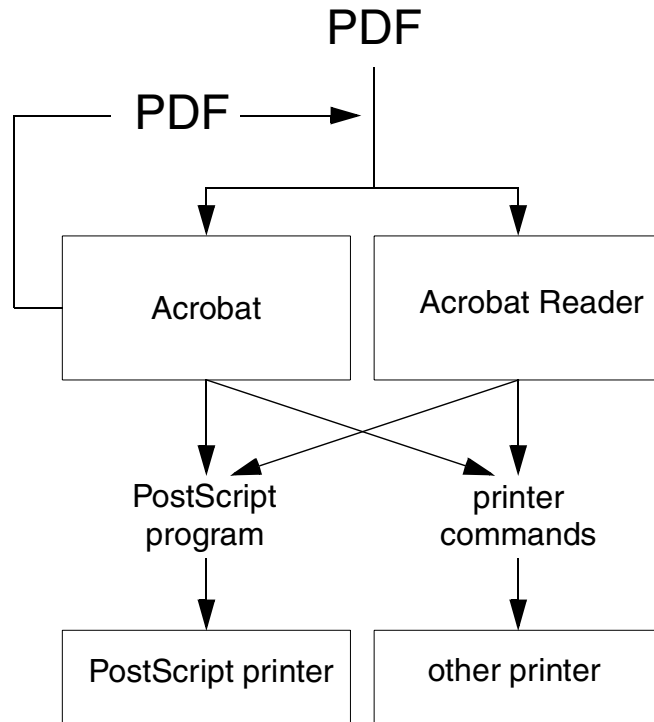
Some applications produce PostScript page descriptions directly because of limitations in the QuickDraw or GDI imaging models or because they run on DOS or UNIX® computers, where there is no system-level printer driver. For these applications, PostScript page descriptions can be converted into PDF files using the Acrobat Distiller® application, as shown in [Figure 2.2](#). The Distiller application accepts any PostScript page description, whether created by a program or hand-coded by a human. The Distiller application produces more efficient PDF files than PDF Writer for some application programs.

Figure 2.2 *Creating PDF files using the Distiller program*



Once a PDF file has been created, the Acrobat viewer can be used to view and print the document contained in the file, as shown in [Figure 2.3](#). Users can navigate through the document using thumbnail sketches, hypertext links, and bookmarks. The document's text may be searched and extracted for use in other applications. In addition, an Acrobat user may modify a PDF document by creating text annotations, hypertext links, thumbnail sketches of each page, and bookmarks that directly access views of specific pages.

Figure 2.3 *Viewing and printing a PDF document*



2.3 General properties

Given the goals and intended use of PDF, its design has several notable properties. This section describes those properties.

2.3.1 Adobe imaging model

PDF represents text and graphics using the Adobe imaging model; this is the same imaging model that is used by the PostScript language. Like a PostScript language program, a PDF page description draws a page by placing “paint” on selected areas.

- The painted figures may be letter shapes, regions defined by combinations of lines and curves, or sampled images such as digitally sampled representations of photographs.
- The paint may be any color.
- Any figure can be clipped to another shape, so that only portions of the figure within the shape appear on the page.
- When a page description begins, the page is completely blank. Various operators in the page description place marks on the page. Each new mark on the page completely obscures any marks it may overlay.

The PDF page-marking operators are similar to the marking operators in the PostScript language. The main reason that the PDF marking operators differ from the PostScript language marking operators is that PDF is not a programming language and does not contain procedures, variables, or control constructs. PDF trades reduced flexibility for improved efficiency. A typical PostScript language program defines a set of high-level operators using the PostScript language marking operators. PDF defines its own set of high-level operators that is sufficient for describing most pages. Because these operators are implemented directly in machine code rather than PostScript language code, PDF page descriptions can be drawn more quickly. Because arbitrary programming constructs are not permitted, applications can more efficiently and reliably locate text strings in a PDF document.

2.3.2 Portability

PDF 1.1

A PDF file is a binary file; the entire 8-bit range of characters may be used.

Unfortunately, some agents treat files that happen to use only the printable subset of the 7-bit ASCII code and whitespace characters as “text,” and take unreasonable liberties with the contents. For example, mail-transmission systems may not preserve certain 7-bit characters and may change line endings. This can cause damage to PDF files.

Therefore, in situations where it is possible to label PDF files as “binary,” we recommend that this be done. One method for encouraging such treatment is to include a few binary characters (codes greater than 127) in a comment near the beginning of the file, as described in [Section 5.12 on page 55](#), *even if the rest of the*

file is ASCII. This ensures that a PDF file will be treated as binary when this is possible, while still allowing it to be transferred through a non-binary channel without damage.

Implementation note Acrobat applications produce PDF files with a comment that includes binary characters.

Note PDF 1.0 files are ASCII files, and all binary data in them must be represented in 7-bit ASCII by means of ASCII filters or escape characters in strings. No line in a PDF 1.0 file may be longer than 255 characters.



2.3.3 Compression

To reduce file size, PDF supports a number of industry-standard compression filters:

- JPEG compression of color and grayscale images.
- CCITT Group 3, CCITT Group 4, LZW (Lempel-Ziv-Welch), and Run Length compression of monochrome images.
- LZW and Flate compression of text, graphics, and indexed image data.

PDF 1.2

Using JPEG compression, color and grayscale images can be compressed by a factor of 10 or more. Effective compression of monochrome images depends upon the compression filter used and the properties of the image, but reductions of 2:1 to 8:1 are common. LZW or Flate compression of text and graphics comprising the balance of the document results in compression ratios of approximately 2:1. All of these compression filters produce binary data, which may then be encoded in the ASCII base-85 encoding to maintain portability.

2.3.4 Font independence

Managing fonts is a fundamental challenge in document exchange. Generally, the receiver of a document must have the same fonts the sender used to create the document. Otherwise, a default font is substituted, producing unexpected and undesirable effects because the default font has different character metrics (widths) than the intended font. Differing font-widths may cause lines to extend into margins or overlap graphics.

The sender could include the fonts with the document, but this can easily make even a short document quite large—a typical two-page memo using four fonts might grow from 10K to 250K. Another possibility is that the sender could convert each page of the document to a fixed-resolution image like a facsimile. Even when compressed, however, the image of a single page can be quite large (45–60K when sampled at 200 dpi). In addition, there is no intelligence left in the file, preventing the receiver from searching for or extracting text from the document.

PDF provides a new solution that makes a document independent of the fonts used to create it. A PDF file contains a *font descriptor* for each font used in a document. The font descriptor includes the font name, character metrics, and style information. This is the information needed to simulate missing fonts and is typically only 1–2K per font.

If a font used in a document is available on the computer where the document is viewed, it is used. If it is not available, a multiple master font is used to simulate on a character-by-character basis the weight and width of the original font, to maintain the overall “color” and formatting of the document. This solution applies to both Adobe Type 1 fonts and fonts in the TrueType™ format [32] developed by Apple Computer, Inc.

Symbolic fonts must be handled in a special way. A symbolic font is any font that does not use the standard ISOLatin1 character set. Fonts such as Carta®, Adobe Caslon™ Swash Italic, Minion® Ornaments, and Lucida® Math fall into this category. It is not possible to simulate a symbolic font effectively.

For symbolic fonts, a font descriptor (including metrics and style information) is not sufficient; the actual character shapes (or glyphs) are required to accurately display and print the document. For all symbolic fonts other than Symbol and ITC Zapf Dingbats®, a compressed version of the Type 1 font program for the font is included in the PDF file. Symbol and ITC Zapf Dingbats, the most widely used symbolic fonts, ship with Acrobat and do not need to be included in a PDF file.

2.3.5 Single-pass file generation

Because of system limitations and efficiency considerations, it may be desirable or necessary for an implementation of a program that produces PDF to create a PDF file in a single pass. This may be, for example, because the application has access to limited memory or is unable to open temporary files. For this reason, PDF supports single-pass generation of files. While PDF requires certain objects to contain a number specifying their length in bytes, a mechanism is provided allowing the length to be located in the file after the object. In addition, information such as the number of pages in the document can be written into the file after all pages have been written into the file.

2.3.6 Random access

Tools that extract and display a selected page from a PostScript language program must scan the program from its beginning until the desired page is found. On average, the time needed to view a page depends not only on the complexity of the page but also on the total number of pages in the document. This is problematic for interactive document viewing, where it is important that the time needed to view a page be independent of the total number of pages in the document.

Every PDF file contains a cross-reference table that can be used to locate and directly access pages and other important objects in the file. The location of the cross-reference table is stored at the end of the file, allowing applications that produce PDF files in a single pass to store it easily and allowing applications that

read PDF files to locate it easily. Using the cross-reference table, the time needed to view a page in a PDF file can be nearly independent of the total number of pages in the document.

2.3.7 Incremental update

Applications may allow users to modify PDF documents, which can contain hundreds of pages or more. Users should not have to wait for the entire file to be rewritten each time modifications to the document are saved. PDF allows modifications to be appended to a file, leaving the original data intact. The addendum appended when a file is incrementally updated contains only the objects that were modified or added, and includes an update to the cross-reference table. Support for incremental update allows an application to save modifications to a PDF document in an amount of time proportional to the size of the modification instead of the size of the file. In addition, because the original contents of the file are still present in the file, it is possible to undo saved changes by deleting one or more addenda.

2.3.8 Extensibility

PDF is designed to be extensible. Undoubtedly, developers will want to add features to PDF that have not yet been implemented or thought of.

The design of PDF is such that not only can new features be added, but applications that understand earlier versions of the format will not completely break when they encounter features that they do not implement. [Appendix G, “Compatibility and Implementation Notes,”](#) specifies how a PDF viewer should behave when it reads a file that does not conform to the specification it was expecting.

2.4 PDF and the PostScript language

The preceding sections mentioned several ways in which PDF differs from the PostScript language. This section summarizes these differences and describes the process of converting a PDF file into a PostScript language program.

While PDF and the PostScript language share the same imaging model, there are some important differences between them:

- A PDF file may contain objects such as hypertext links that are useful only for interactive viewing.
- To simplify the processing of page descriptions, PDF provides no programming language constructs.
- PDF enforces a strictly defined file structure that allows an application to access parts of a document randomly.
- PDF files contain information such as font metrics, to ensure viewing fidelity.


Because of these differences, a PDF file cannot be downloaded directly to a PostScript printer for printing. An application that prints a PDF file to a PostScript printer must carry out the following steps:

1. Insert *procsets*, sets of PostScript language procedure definitions that implement the PDF page description operators.
2. Extract the content for each page. Pages are not necessarily stored in sequential order in the PDF file. Each page description is essentially the script portion of a traditional PostScript language program using very specific procedures, such as “m” for **moveto** and “l” for **lineto**.
3. Decode compressed text, graphics, and image data. Except for data encoded with the Flate filter, this is not required for PostScript Level 2 printers, which can accept compressed data in a PostScript language file.
4. Insert any resources, such as fonts, into the PostScript language file. Substitute fonts are defined and inserted as needed, based on the font metrics in the PDF file.
5. Put the information in the correct order. The result is a traditional PostScript language program that fully represents the visual aspects of the document, but no longer contains PDF elements such as hypertext links, annotations, and bookmarks.
6. Send the PostScript language program to the printer.

2.5 Understanding PDF

PDF is best understood by thinking of it in four parts, as shown in [Figure 2.4](#).

Figure 2.4 *PDF components*



The first component is the set of basic object types used by PDF to represent objects. These types, with only a few exceptions, correspond to the data types used in the PostScript language. [Chapter 4](#) discusses these object types.

The second component is the PDF file structure. The file structure determines how objects are stored in a PDF file, how they are accessed, and how they are updated. This structure is independent of the semantics of the objects. [Chapter 5](#) explains the file structure.

The third component is the PDF document structure. The document structure specifies how the basic object types are used to represent components of a PDF document: pages, annotations, hypertext links, fonts, and more. [Chapter 6](#) explains the PDF document structure.

The fourth and final component is the PDF page description. A PDF page description, while part of a PDF page object, can be explained independently of the other components. A PDF page description has only limited interaction with other parts of a PDF document. This simplifies its conversion into a PostScript language program. [Chapter 8](#) discusses PDF page descriptions.

CHAPTER 3**Coordinate Systems**


Coordinate systems define the canvas on which all drawing in a PDF document occurs; that is, the position, orientation, and size of the text, graphics, and images that appear on a page are determined by coordinate systems.

PDF supports a number of coordinate systems, most of them identical to those used in the PostScript language. This chapter describes each of the coordinate systems used in PDF, how they are related, and how transformations among coordinate systems are specified. At the end of the chapter is a description of the mathematics involved in coordinate transformations. It is not necessary to read this section to use coordinate systems and transformations. It is presented for those readers who wish to gain a deeper understanding of the mechanics of coordinate transformations.

3.1 Device space

The contents of a page ultimately appear on a display or a printer. Each type of device on which a PDF page can be drawn has its own built-in coordinate system, and, in general, each type of device has a different coordinate system. Coordinates specified in a device's native coordinate system are said to be in *device space*. On pixel-based devices such as computer screens and printers, coordinates in device space generally specify a particular pixel.


If coordinates in PDF files were specified in device space, the files would be device-dependent and would accordingly appear differently on different devices. For example, images drawn in the typical device space of a 72 pixel per inch display and on a 600 dpi printer differ in size by more than a factor of 8; an eight-inch line segment on a display would appear as a one-inch segment on the printer. Different devices also have different orientations of their coordinate systems. On one device, the origin of the coordinate system may be at the upper left corner of the page, with the positive direction of the y -axis pointing downward. On another device, the origin may be in the lower left corner of the page with the positive direction of the y -axis pointing upward. [Figure 3.1](#) shows an object that is two units high in device space, and illustrates the fact that coordinates specified in device space are device-dependent.

Figure 3.1 *Device space*

3.2 User space

PDF, like the PostScript language, defines a coordinate system that appears the same, regardless of the device on which output occurs. This allows PDF documents to be independent of the resolution of the output device. This resolution-independent coordinate system is called *user space* and provides the overall coordinate system for a page.

The transformation from user space to device space is specified by the *current transformation matrix* (CTM). Figure 3.2 shows an object that is two units high in user space and indicates that the CTM provides the resolution-independence of the user space coordinate system.

Figure 3.2 *User space*

The user space coordinate system is initialized to a default state for each page of a document. By default, user space coordinates have 72 units per inch, corresponding roughly to the various definitions of the typographic unit of measurement known as the *point*. The positive direction of the *y*-axis points upward, and the positive direction of the *x*-axis to the right. The region of the default coordinate system that is viewed or printed can be different for each page, and is described on [page 73](#).

3.3 Text space

The coordinates of text are specified in *text space*. The transformation from text space to user space is provided by a matrix called the *text matrix*.

3.4 Character space

Characters in a font are defined in *character space*. The transformation from character space to text space is defined by a matrix. For most types of fonts, this matrix is predefined except for an overall scale factor. (For details, see [page 339](#).) This scale factor changes when a user selects the font size for text.

3.5 Image space

All images are defined in *image space*. The transformation from image space to user space is predefined and cannot be changed. All images are one unit by one unit in user space, regardless of the number of samples in the image.

This is different from images in PostScript. In PostScript, an image has an *ImageMatrix* that maps image space (the width and height in samples) to user space (e.g., a rotated 20×35 rectangle in the middle of the page). In PDF, image space always maps to the 1×1 square whose lower left corner is at (0,0). The CTM then maps that wherever it wants (e.g., the rotated 20×35 rectangle in the middle of the page). This makes the image more “portable”: the same image can be used multiple times in the same document.

3.6 Form space

PDF provides an object known as a *Form XObject*, discussed on [page 255](#). Form XObjects contain sequences of operations and are the same as forms in the PostScript language. The space in which a Form XObject is defined is *form space*. The transformation from form space to user space is specified by a matrix contained in the Form XObject.

Note Beginning with version 1.2, PDF defines a feature that allows a PDF document to represent a form, called an AcroForm. By “form” we mean the PDF equivalent of the familiar paper instrument. Any unqualified use of the term “form” refers to an AcroForm described on [page 129](#); the other is referred to explicitly as a Form XObject.

3.7 Pattern space

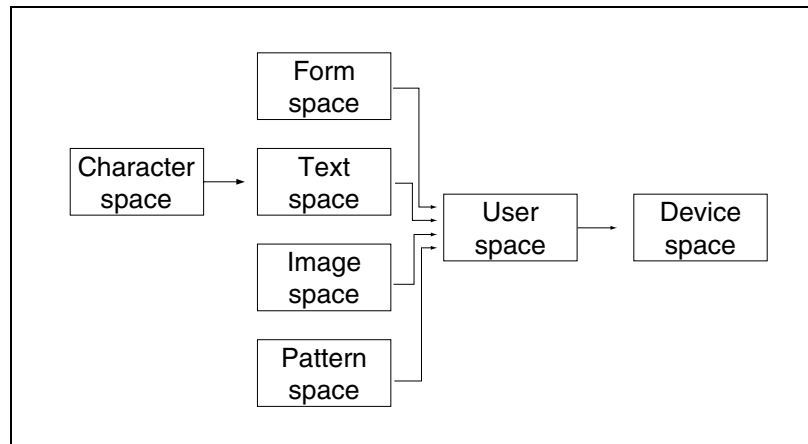
PDF 1.2

PDF defines a type of color known as a *pattern*, discussed on [page 287](#). Like forms, patterns can contain sequences of marking operations; they are the same as patterns in the PostScript language. The space in which a pattern is defined is *pattern space*. The transformation from pattern space to user space is specified by a matrix contained in the pattern.

3.8 Relationships among coordinate systems

PDF defines a number of interrelated coordinate systems, described in the previous sections. [Figure 3.3](#) shows the relationships among the coordinate systems. Each line in the figure represents a transformation from one coordinate system to another. PDF allows modifications to many of these transformations. Such transformations are described in [Section 3.9 on page 33](#).

Figure 3.3 Relationships among PDF coordinate systems



Because PDF coordinate systems are defined relative to each other, changes made to one transformation can affect the appearance of objects drawn in several coordinate systems. For example, a change in the CTM, which defines the map from User space to Device space, will affect forms, text, images, and patterns, since they're “upstream” of User space.

3.9 Transformations between coordinate systems


Transformation matrices specify the relationship between two coordinate systems. By modifying a transformation matrix, objects can be scaled, rotated, translated, or transformed in other ways.

A transformation matrix in PDF, as in the PostScript language, is specified by six numbers, usually in the form of an array containing six elements. This section lists the arrays that specify the most common transformations. The following section contains more mathematical details of transformations, including information on specifying transformations that are combinations of those listed in this section.

- Translations are specified as $[1 \ 0 \ 0 \ 1 \ t_x \ t_y]$, where t_x and t_y are the distances to translate the origin of the coordinate system in x and y , respectively.
- Scaling is obtained by $[s_x \ 0 \ 0 \ s_y \ 0 \ 0]$. This scales the coordinates so that one unit in the x and y directions of the new coordinate system is the same size as s_x and s_y units in the previous coordinate system, respectively.
- Rotations are carried out by $[\cos \theta \ \sin \theta \ -\sin \theta \ \cos \theta \ 0 \ 0]$, which has the effect of rotating the coordinate system axes by an angle θ counterclockwise.
- Skew is specified by $[1 \ \tan \alpha \ \tan \beta \ 1 \ 0 \ 0]$, which skews the x -axis by an angle α and the y -axis by an angle β .


Figure 3.4 shows examples of each transformation. The directions of translation, rotation, and skew shown in the figure correspond to positive values of the array elements.

Figure 3.4 *Effects of coordinate transformations*



Translation


Scaling




If several transformations are applied, the order in which they are applied is important. For example, scaling the x -axis followed by a translation of the x -axis is not the same as first translating the x -axis, then performing the scaling. In general, to obtain the expected results, transformations should be done in the order: translate, rotate, scale/skew.

Figure 3.5 shows the effect of the order in which transformations are applied. The figure shows two sequences of transformations applied to a coordinate system. After each successive transformation, an outline of the letter “n” is drawn. The transformations in the figure are a translation of 10 units in the x -direction and 20 units in the y -direction, a rotation of 30 degrees, and a scaling by a factor of 3 in the x -direction. In the figure, the axes are drawn with a dash-pattern having two units dash, two units gap. In addition, the untransformed coordinate system is drawn in a lighter color in each section. Notice that the scale–rotate–translate ordering results in a distortion of the coordinate system leaving the x - and y -axes no longer perpendicular, while the recommended translate–rotate–scale ordering does not.

Figure 3.5 *Effect of the order of transformations*





3.10 Transformation matrices

Note This section describes the mathematics of transformation matrices, which is identical to that underlying the PostScript language. It is not necessary to read this section to use the transformations discussed in previous sections.

To understand coordinate system transformations in PDF, it is vital to understand two points:

- Transformations in PDF alter coordinate systems, not objects. All objects drawn before a transformation is specified are unchanged by the transformation. Objects drawn after the transformation is specified will be drawn in the transformed coordinate system.
- Transformation matrices in PDF specify the transformation from the transformed (new) coordinate system to the untransformed (old) coordinate system. All coordinates used after the transformation are specified in the transformed coordinate system. PDF applies the transformation matrix to determine the coordinates in the untransformed coordinate system.

Note Many computer graphics textbooks consider transformations of objects instead of coordinate systems. Although these are formally equivalent, some results differ depending on which point of view is taken.

PDF represents coordinates in a two-dimensional space. The point (x, y) in such a space can be expressed in vector form as $[x \ y \ 1]$. Although the third element of this vector (1) is not strictly necessary, it provides a convenient way to specify translations of the coordinate system's origin.

The transformation between two coordinate systems is represented by a 3×3 transformation matrix written as:

$$\begin{bmatrix} a & b & 0 \\ c & d & 0 \\ e & f & 1 \end{bmatrix}$$

Note Because a transformation matrix has only six entries that may be changed, for convenience it is often written as the six-element array $[a \ b \ c \ d \ e \ f]$.

Coordinate transformations are expressed as:

$$\begin{bmatrix} x' & y' & 1 \end{bmatrix} = \begin{bmatrix} x & y & 1 \end{bmatrix} \begin{bmatrix} a & b & 0 \\ c & d & 0 \\ e & f & 1 \end{bmatrix}$$

Because PDF transformation matrices specify the conversion from the transformed coordinate system to the original (untransformed) coordinate system, x' and y' in this equation are the coordinates in the untransformed coordinate system, while x and y are the coordinates in the transformed system. Carrying out the multiplication, we have:

$$x' = ax + cy + e$$

$$y' = bx + dy + f$$

If a series of transformations is carried out, the transformation matrices representing each of the transformations can be multiplied together to produce a single equivalent transformation matrix.

Matrix multiplication is not commutative—the order in which matrices are multiplied is significant. It is not *a priori* obvious in which order the transformation matrices should be multiplied. Matrices representing later transformations could either be multiplied before those representing earlier transformations (premultiplied) or after (postmultiplied).

To determine whether premultiplication or postmultiplication is appropriate, consider a sequence of two transformations. Specifically, apply a scaling transformation to the user space coordinate system, and consider the conversion from this scaled coordinate system to device space. The two transformation matrices in this example are the matrix specifying the scaling (M_S) and the matrix specifying the transformation from user space to device space (the CTM, called M_C here). Recalling that coordinates are always specified in the transformed space, it is clear that the correct order of transformations must first convert the scaled coordinates to those in default user space, and then convert the default user space coordinates to device space coordinates. This can be expressed:

$$X_D = X_U M_C = (X_S M_S) M_C = X_S (M_S M_C)$$

where X_D is the coordinate in device space and X_U is the coordinate in default user space. This shows that when a new transformation is added, the matrix representing it must be *premultiplied* onto the existing transformation matrix.

This result is true in general for PDF—when a sequence of transformations is carried out, the matrix representing the combined transformation (M') is calculated by premultiplying the matrix representing the transformation being added (M_T) onto the matrix representing any existing transformations (M):

$$M' = M_T M$$

CHAPTER 4

Objects

4.1 Introduction

PDF supports seven basic types of objects: booleans, numbers, strings, names, arrays, dictionaries, and streams. In addition, PDF provides a null object. Objects may be labeled so that they can be referred to by other objects. A labeled object is called an *indirect object*.

The following sections describe each object type and the null object. A discussion of creating and referring to indirect objects in PDF files follows.

Note PDF is case-sensitive. Uppercase and lowercase letters are different.

4.2 Booleans

The keywords **true** and **false** represent boolean objects with values *true* and *false*.

4.3 Numbers

PDF provides two types of numbers, integer and real. Integers may be specified by signed or unsigned constants. Reals may only be in decimal format. Throughout this book, *number* means an object whose type is either integer or real.

Note Exponential format for numbers (such as *1.0E3*) is not supported.

4.4 Strings and text

A *string* is a series of unsigned 8-bit bytes, written as a sequence of characters delimited by parentheses. If a string is too long to be conveniently placed on a single line, it may be split across multiple lines by using the backslash (\) character at the end of a line to indicate that the string continues on the following line. When this occurs, the backslash and end-of-line characters are not considered part of the string. Examples of strings are:

```
( This is string number 1? )  
( strangeonium spectroscopy )
```

**(This string is split \
across \
three lines)**

Within a string, the backslash character is used as an escape character to specify unbalanced parentheses, non-printing ASCII characters, and the backslash character itself. This is the same as for PostScript language strings, described in Section 3.2.2 of the *PostScript Language Reference Manual, Second Edition*. [Table 4.1](#) lists the escape sequences for PDF.

Table 4.1 *Escape sequences in strings*

\n	linefeed
\r	carriage return
\t	horizontal tab
\b	backspace
\f	formfeed
\\	backslash
\(left parenthesis
\)	right parenthesis
\ddd	character code <i>ddd</i> (octal)

The **\ddd** escape sequence provides a way to represent characters outside the printable ASCII character set. (This mechanism is not needed in binary PDF files; see [Section 5.12 on page 55](#).) The number *ddd* may contain one, two, or three octal digits. An example of a string with octal characters in it is:

(string with \245two octal characters\307)

It is recommended that three octal characters always be used, in order to prevent any ambiguity with the subsequent characters. For example, **(\0053)** denotes a string that contains two characters, ASCII 5 (control E), and the digit 3, whereas **(\053)** and **(\53)** denote strings that contain a single character, a plus-sign (+).

As in the PostScript language, strings may also be represented in hexadecimal form. A hexadecimal string is written as a sequence of hexadecimal characters (the digits **0–9** and the letters **A–F** or **a–f**) enclosed within angle brackets (**<** and **>**). Each pair of hexadecimal digits defines one byte of the string. If the final digit of a given string is missing—in other words, if there is an odd number of digits—the final digit is assumed to be zero. Whitespace characters (space, tab, carriage return, linefeed, and formfeed) are ignored. For example,

<901fa3>

is a three-byte string consisting of the characters whose hexadecimal codes are 90, 1f, and a3. But

<901fa>

is a three-byte string containing the characters whose hexadecimal codes are 90, 1f, and a0.

Any 8-bit value may appear in a string. In particular, when a document is encrypted (see [page 62](#)), all its strings are encrypted and often contain arbitrary 8-bit values. Note that the backslash character is still required as an escape to specify unbalanced parentheses and the backslash character itself.

Strings can be used for many purposes and can be formatted in different ways. When a string is used for a specific purpose, to represent a date, for example, it is useful to have a standard format for that purpose; see [Section 7.2 on page 183](#). Such formats are conventions for interpreting strings and are not types themselves. The use of a particular format is indicated with the definition of the string object that uses the format.

4.4.1 Text

Certain strings contain information that is intended to be human-readable. These strings include text annotations, bookmark names, article names, document information, etc. They are referred to as *text*. Text is encoded in either **PDFDocEncoding** or Unicode. **PDFDocEncoding**, which is a superset of ISOLatin1, is compatible with Unicode in that all Unicode codes less than 256 match **PDFDocEncoding**.

If text is encoded in Unicode, the first two bytes of the text must be the Unicode Byte Order marker, **<FE FF>**. This sequence collides with the character sequence *thorn ydieresis*, which is not likely to be a meaningful beginning of a word or phrase. The high-order byte of a Unicode character appears first in the text.

The text may also contain an escape sequence to indicate the language of the text. This is useful when the language cannot be determined from the character codes used in the text. The escape sequence uses the Unicode hex value U+001B followed by the two ASCII codes for the language identifiers defined by ISO 639 (see [Appendix I](#)), optionally followed by the two ASCII codes for country defined by ISO 3166 (see [Appendix J](#)), followed by U+001B.

4.5 Names

A name, like a string, is written as a sequence of characters. It must begin with a slash (/) followed by a sequence of ASCII characters in the range ! (<21>) through ~ (<7E>) except %, (,), <, >, [,], {, }, /, and #. Examples of names are:

```
/Name1
/ASomewhatLongerName2
/A;Name_With-various***Characters?.
/1.2
/.notdef
```

Any character except null (<00>) may be included in a name by writing its two-character hex code, preceded by #. Examples:



```
/Adobe#20Green
/PANTONE#205757#20CV
/paired#28#29parentheses
/TheKeyOff#23Minor
/A#42
```

Note that **/A#42** is another way of writing the name **/AB**.




Note In PDF 1.1 (and 1.0), the character # could be used as part of a name, (e.g., **/A#B**), and the specs didn't specifically outlaw blanks (although Adobe generators did not provide a way to write names containing blanks). In PDF 1.2, the character # became an escape character, preceding two hex digits. Thus to create a 3-character name A-space-B, you write **/A#20B**, since 0x20 is the hex code for space. This means that **/A#B** is no longer legal, since the # is not followed by two hex digits. You can create the name A-sharp-B by writing **/A#42B**, since 0x42 is the hex code for sharp (#). There was no escape character for names in PDF 1.1.

Note The maximum number of characters in a name is 127. This limit refers to the internal representation of the name. For example, the name **/A#20B** has four characters (**/**, **A**, space, and **B**), not six.

4.6 Arrays

An array is a sequence of PDF objects. An array may contain a mixture of object types. An array is written as a left square bracket (**[**), followed by a sequence of objects, followed by a right square bracket (**]**). An example of an array is:

```
[0 (Higgs) false 3.14 3 549 /SomeName]
```

4.7 Dictionaries

A dictionary is an associative table containing pairs of objects. The first element of each pair is called the *key* and the second element is called the *value*. The key must be a name (unlike keys in dictionaries in the PostScript language). A value can be any kind of object, including a dictionary. A dictionary is generally used to collect and tie together the attributes of a complex object, with each key–value pair specifying the name and value of an attribute.

A dictionary is written as two left angle brackets (**<<**), followed by a sequence of key–value pairs, followed by two right angle brackets (**>>**). For example:

Example 4.1 Dictionary

```
<< /Type /Example /Key2 12 /Key3 (a string) >>
```


Example 4.2 *Dictionary within a dictionary*

```

<<
  /Type /AlsoAnExample
  /Subtype /Bad
  /Reason (unsure)
  /Version 0.01
  /MyInfo
    <<
      /Item1 0.4
      /Item2 true
      /LastItem (not!)
      /VeryLastItem (OK)
    >>
  >>

```

Dictionary objects are the main building blocks of a PDF document. Many parts of a PDF document, such as pages and fonts, are represented using dictionaries. By convention, the **Type** key of such a dictionary specifies the type of object being described by the dictionary. Its value is always a name. In some cases, the **Subtype** key is used to describe a specialization of a particular type. Its value is always a name. For example, in a font, the value of the **Type** key is **Font** and several subtypes exist, including **Type1**, **MMType1**, **Type3**, **TrueType**, and others. See the note on [page 14](#) for additional restrictions on the **Type** key.

4.8 Streams

A stream, like a string, is a sequence of characters. However, an application can read a small portion of a stream at a time, while a string must be read in its entirety. For this reason, objects with potentially large amounts of data, such as images and page descriptions, are represented as streams.

A stream consists of a dictionary that describes a sequence of characters, followed by the keyword **stream**, followed by zero or more lines of characters, followed by the keyword **endstream**.

```

<stream> ::= <dictionary>
           stream
           { <lines of characters> }*
           endstream

```

All streams must be indirect objects (see [page 53](#)). The stream dictionary must be a direct object. The keyword **stream** that follows the stream dictionary should be followed by a carriage return and linefeed or just a linefeed.

Note Without this restriction, it is not possible to differentiate a stream that uses carriage return as end of line and whose first byte of data is a linefeed from a stream that uses carriage return-linefeed pairs as end of line.

PDF 1.1

The sequence of characters that make up a stream may be found between the **stream** and **endstream** keywords or, in PDF 1.2, may be contained in an external file. If the data is in an external file, the stream dictionary specifies the file. When a stream's data is external, the characters between **stream** and **endstream** are ignored.

PDF 1.2

Table 4.2 shows the attributes of a stream.

Table 4.2 *Stream attributes*

<i>Key</i>	<i>Type</i>	<i>Description</i>
Length	integer	<i>(Required)</i> Number of characters from the first line after the line containing the stream keyword to the endstream keyword.
Filter name or array of names		<i>(Optional)</i> Filters to be applied in processing the stream. The value of the Filter key can be either the name of a single decode filter or an array of decode filter names. Specify multiple filters in the order they should be applied to decode the data. For example, data encoded using LZW and ASCII base-85 filters (in that order) can be decoded by providing the following key and value in the stream dictionary: <code>/Filter [/ASCII85Decode /LZWDecode]</code>
DecodeParms	<i>various</i>	<i>(Optional)</i> Parameters used by the decoding filters specified with the Filter key. The number and types of the parameters supplied must match those needed by the specified filters. For example, if two filters are used, the decode parameters must be specified by an array of two objects, one corresponding to each filter. Use the null object for a filter's entry in the DecodeParms array if that filter does not need any parameters. If none of the filters specified requires any parameters, omit the DecodeParms key.
F (File) File specification		<i>(Optional)</i> The file containing the stream data. If this key is present, the characters between stream and endstream are ignored. However, the Length key should still specify the number of those characters. (Usually there are no characters and the Length is zero.)
FFilter name or array of names		<i>(Optional)</i> Filters to be applied in processing the data found in the stream's external file. The same rules apply as for the Filter key.
FDecodeParms	<i>various</i>	<i>(Optional)</i> Parameters used by the decoding filters specified with the FFilter key.

PDF 1.2

PDF 1.2

PDF 1.2

Streams may be filtered to compress them or to convert binary streams into ASCII form. These filters and their parameters are listed in Table 4.3 and described in the following sections.

Table 4.3 *Standard filters*

<i>Filter name</i>	<i>Parameters</i>	<i>Semantics</i>
ASCIIHexDecode	none	Decodes binary data in an ASCII hexadecimal representation.
ASCII85Decode	none	Decodes binary data in an ASCII base-85 representation.
LZWDecode	dictionary	<i>(Parameters optional)</i> Decompresses text or binary data using LZW adaptive compression method.
RunLengthDecode	none	Decompresses binary data using a byte-oriented run-length decoding algorithm.
CCITTFaxDecode	dictionary	<i>(Parameters optional)</i> Decompresses binary data using a bit-oriented decoding algorithm, the CCITT facsimile standard.
DCTDecode	dictionary	<i>(Parameters optional)</i> Decompresses sampled image data using a discrete cosine transform technique based on the JPEG standard.
FlateDecode	dictionary	<i>(Parameters optional)</i> Decompresses text or binary data using the Flate decompression method.



PDF 1.2

[Example 4.3](#) shows a stream that contains the marking instructions for a page. This stream has been compressed using LZW and then encoded using ASCII85. [Example 4.4](#) shows the same stream without any encoding.

Example 4.3 *Stream that has been LZW and ASCII85 encoded*

```
<<
  /Length 528
  /Filter [/ASCII85Decode /LZWDecode]
>>
stream
J..)6T`?p&<!J9%_[umg"B7/Z7KNXbn'S+,*Q/&"OLT'F
LIDK#!n`$"<Atdi`\Vn%b%)&'cA*VnK\CJY(sF>c!Jnl@
RM]WM;jjH6Gnc75idkL5]+cPZKEBPwDR>FF(kj1_R%W_d
&/jS!;iuad7h?[L-F$+]]0A3Ck*$I0KZ?;<)CJtqi65Xb
Vc3\n5ua:Q/=0$W<#N3U;H,MQKqfg1?:lUpR;6oN[C2E4
ZNr8Udn.'p+?#X+1>0Kuk$bCDF/(3fL5]Oq)^kJZ!C2H1
'TO]Rl?Q:&'<5&iP!$Rq;BXRecDN[IJB`,)o8XJOSJ9sD
S]hQ;Rj@!ND)bd_q&C\g:inYC%)&u#:u,M6Bm%IY!Kb1+
":aAa'S`ViJg1Lb8<W9k6Y1\0McJQkDeLWdPN?9A'jX*
al>iG1p&i;eVoK&juJHs9%;Xomop"5KatWRT"JQ#qYuL,
JD?M$0QP)lKn0611apKDC@qJ4B!!(5m+j.7F790m(Vj8
818Q:_CZ(Gm1%X\N1&u!FKHMB~>
endstream
```

Example 4.4 *Unencoded stream*

(The operators used in this example are described in [Chapter 8](#).)

```

<<
/Length 558
>>
stream
2 J
BT
/F1 12 Tf
0 Tc 0 Tw 72.5 712 TD
[(Unencoded streams can be read easily)65 (,)] TJ
0 -14 TD
[(b)20 (ut generally tak)10
(e more space than \311)] TJ
T* (encoded streams.)Tj
0 -28 TD
[(Se) 25 (v) 15 (eral encoding methods are a) 20 (v)
25 (ailable in PDF)80 (.)] TJ
0 -14 TD
(Some are used for compression and others simply)
Tj
T* [(to represent binary data in an ) 55
(ASCII format.)] TJ
T*
(Some of the compression encoding methods are \
suitable )
Tj
T*
(for both data and images, while others are \
suitable only ) Tj
T* (for continuous-tone images.)Tj
ET
endstream

```

4.8.1 **ASCIIHexDecode filter**

This filter decodes data that has been encoded as ASCII hexadecimal. ASCII hexadecimal encoding and ASCII base-85 encoding (described in the following section) convert binary data such as images to 7-bit data. In general, ASCII base-85 encoding is preferred to ASCII hexadecimal encoding because it is more compact.

ASCII hexadecimal encoding produces a 1:2 expansion in the size of the data. Each pair of ASCII hexadecimal digits (**0–9** and **A–F** or **a–f**) produces one byte of binary data. All whitespace characters are ignored. The right angle bracket (**>**) indicates the end of data (EOD). Any other character causes an error. If the filter encounters the EOD marker after reading an odd number of hexadecimal digits, it behaves as if a zero followed the last digit.

4.8.2 ASCII85Decode filter

This filter decodes data that has been encoded in the ASCII base-85 encoding and produces binary data.

ASCII base-85 encoding produces five ASCII printing characters from every four bytes of binary data. Each group of four binary bytes ($b_1 b_2 b_3 b_4$) is converted to a group of five encoded characters ($c_1 c_2 c_3 c_4 c_5$) using the relation:

$$\begin{aligned} & (b_1 \times 256^3) + (b_2 \times 256^2) + (b_3 \times 256) + b_4 \\ &= (c_1 \times 85^4) + (c_2 \times 85^3) + (c_3 \times 85^2) + (c_4 \times 85) + c_5 \end{aligned}$$

The five “digits” of the encoded base-85 number are converted to printable ASCII characters by adding 33 (the ASCII code for **!**) to each. The resulting data contains only printable ASCII characters with codes in the range 33 (**!**) to 117 (**u**).

Two special cases occur during encoding. First, if all five encoded digits are zero, they are represented by the character code 122 (**z**), instead of by a series of five exclamation points (**!!!!**). In addition, if the length of the binary data to be encoded is not a multiple of four bytes, the last partial 4-tuple is used to produce a last, partial output 5-tuple. Given n (1, 2, or 3) bytes of binary data, the encoding first appends $4 - n$ zero bytes to make a complete 4-tuple. This 4-tuple is encoded in the usual way, but without applying the special **z** case. Finally, only the first $n + 1$ characters of the resulting 5-tuple are written out. Those characters are immediately followed by the EOD marker, which is the two-character sequence **~>**.

The following conditions are errors during decoding:

- The value represented by a 5-tuple is greater than $2^{32} - 1$.
- A **z** character occurs in the middle of a 5-tuple.
- A final partial 5-tuple contains only one character.

These conditions never occur in the output produced from a correctly encoded byte sequence.

4.8.3 LZWDecode filter

This filter decodes data encoded using the LZW data compression method, which is a variable-length, adaptive compression method. LZW encoding compresses binary and ASCII text data but always produces binary data, even if the original data was ASCII text.

LZW compression can discover and exploit many patterns in its input data, whether that input is text or image data. The compression obtained using the LZW method varies from file to file; the best case (a file of all zeros) provides a

compression approaching 1365:1 for long files, while the worst case (a file in which no pair of adjacent characters appears twice) can produce an expansion of approximately 50%.

Data encoded using LZW consist of a sequence of codes that are 9 to 12 bits long. Each code represents a single character of input data (0–255), a clear-table marker (256), an EOD marker (257), or a table entry representing a multi-character sequence that has been encountered previously in the input (258 and greater).

Initially, the code length is 9 bits and the table contains only entries for the 258 fixed codes. As encoding proceeds, entries are appended to the table, associating new codes with longer and longer input character sequences. The encoding and decoding filters maintain identical copies of this table.

Whenever both encoder and decoder independently (but synchronously) realize that the current code length is no longer sufficient to represent the number of entries in the table, they increase the number of bits per code by one. The first output code that is 10 bits long is the one following creation of table entry 511, and similarly for 11 (1023) and 12 (2047) bits. Codes are never longer than 12 bits, so entry 4095 is the last entry of the LZW table.

The encoder executes the following sequence of steps to generate each output code:

1. Accumulate a sequence of one or more input characters matching a sequence already present in the table. For maximum compression, the encoder looks for the longest such sequence.
2. Output the code corresponding to that sequence.
3. Create a new table entry for the first unused code. Its value is the sequence found in step 1 followed by the next input character.

To adapt to changing input sequences, the encoder may at any point issue a clear-table code, which causes both the encoder and decoder to restart with initial tables and a 9-bit code. By convention, the encoder begins by issuing a clear-table code. It must issue a clear-table code when the table becomes full; it may do so sooner.

The LZW filter can be used to compress text or images. When compressing images, several techniques reduce the size of the resulting compressed data. For example, image data frequently change very little from sample to sample. By subtracting the values of adjacent samples (a process called *differencing*) and LZW-encoding the differences rather than the raw sample values, the size of the output data may be reduced. Further, when the image data contains several color components (red–green–blue or cyan–magenta–yellow–black) per sample, taking the difference between the values of like components in adjacent samples, rather than between different color components in the same sample, often reduces the output data size. In order to control these and other options, the LZW filter accepts several optional parameters, shown in [Table 4.4](#). All values supplied to the decode filter by any optional parameters must match those used when the data was encoded.

Table 4.4 *Optional parameters for LZW filter*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Predictor	integer	If Predictor is 1, the file is decoded assuming that it was encoded using the normal LZW algorithm. If Predictor > 1, decoding is performed assuming that prior to encoding, the data was differenced. The default value is 1. For details on this value, see Section 4.8.5 on page 48 .
Columns	integer	Has an effect only if Predictor > 1. Columns is the number of samples in a sampled row. The first sample in each row is not differenced; all subsequent samples in a row are differenced with the prior sample. Each row begins on a byte boundary. Any extra bits needed to complete a byte at the end of a row (consisting of Columns × Colors × BitsPerComponent bits) are not differenced. The default value is 1.
Colors	integer	Has an effect only if Predictor > 1. Number of interleaved color components per sample in a sampled image. Each color component is differenced with the value of the same color component in the previous sample. Allowed values are 1, 2, 3, and 4. The default value is 1.
BitsPerComponent	integer	Has an effect only if Predictor > 1. BitsPerComponent is the number of bits used to represent each color component in a pixel. Allowed values are 1, 2, 4, and 8. The default value is 8.
EarlyChange	integer	If EarlyChange is 0, increases in the length of the code are postponed as long as possible. If it is 1, they occur one code word early. The value of EarlyChange used in decoding must match that used during encoding. This parameter is included because LZW sample code distributed by some vendors increases the code word length one word earlier than necessary. The default value is 1.

The LZW compression method is the subject of United States patent number 4,558,302 owned by the Unisys Corporation. Adobe Systems has licensed this patent for use in its Acrobat products. However, independent software vendors may be required to license this patent directly from Unisys to develop software using LZW. For information on Unisys licensing policies for products using LZW (GIF, TIFF-LZW, PostScript, Portable Document Format (PDF), V.42bis, etc.), please send E-mail to lzw_info@unisys.com, or visit the Unisys web server at <http://www.unisys.com/>.

4.8.4 FlateDecode Filter

PDF 1.2

This filter decodes data that has been encoded with the Flate compression method, which encodes binary or ASCII data, optionally after transformation by a predictor function. It is based on the public-domain zlib/deflate compression method, which is a variable-length Lempel-Ziv adaptive compression method cascaded with adaptive Huffman coding. See references [21] and [22] for details. The optional Predictor functions are discussed in the following text.

The output produced by Flate encoding is always binary, even if the input is ASCII text.

Table 4.5 *Optional parameters for FlateDecode filter*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Predictor	integer	If Predictor is 1, the file is decoded assuming that it was encoded using the normal Flate algorithm. If Predictor > 1, decoding is performed assuming that prior to encoding, the data was differenced. The default value is 1. For details on this value, see Section 4.8.5, “LZW and Flate predictor functions.”
Columns	integer	Has an effect only if Predictor > 1. Columns is the number of samples in a sampled row. The first sample in each row is not differenced; all subsequent samples in a row are differenced with the prior sample. Each row begins on a byte boundary. Any extra bits needed to complete a byte at the end of a row (consisting of Columns × Colors × BitsPerComponent bits) are not differenced. The default value is 1.
Colors	integer	Has an effect only if Predictor > 1. Colors is the number of interleaved color components per sample in a sampled image. Each color component is differenced with the value of the same color component in the previous sample. Allowed values are 1, 2, 3, and 4. The default value is 1.
BitsPerComponent	integer	Has an effect only if Predictor > 1. BitsPerComponent is the number of bits used to represent each color component in a pixel. Allowed values are 1, 2, 4, and 8. The default value is 8.

4.8.5 LZW and Flate predictor functions

PDF 1.2

LZW and Flate encoding filters compress more compactly if their input data are highly predictable. One way of increasing the predictability of many continuous-tone sampled images is to replace each pixel with the difference between that pixel and some predictor function applied to earlier neighboring pixels. If the predictor function works well, the postprediction data will cluster toward 0.

Two predictor function groups are supported. The first, the TIFF group, consists of the single function that is Predictor 2 in the TIFF standard. (In TIFF 6.0, it applies only to LZW compression, but here it applies to Flate compression as well.) TIFF predictor 2 predicts that each color component of a pixel will be the same as the corresponding color component of the pixel immediately to the left.

The second supported predictor function group, the *PNG* group, consists of the “filters” of the World Wide Web Consortium’s Portable Network Graphics (PNG) recommendation. The term *predictors* is used here instead of *filters*, to avoid confusion. There are five basic PNG predictor algorithms, and a sixth one that invites an optimum hybrid of the first five. The first five are *None*, *Sub* (predicting the same as the pixel to the left), *Up* (predicting the same as the pixel above),

Average (predicting the average of the pixel to the left and the pixel above), and *Paeth* (a nonlinear function of the pixel above, the pixel to the left, and the pixel to the upper left).

The two predictor function groups have some commonalities. Both assume that data are presented in order, from top row to bottom row, and within a row, from left to right. Both assume that a row occupies a whole number of bytes, rounded upward if necessary. Both assume that pixels and their components are packed into bytes from high- to low-order bits. Both assume that all color components of pixels outside the image (which are necessary for predictions near the boundaries) are 0.

Table 4.6 *Predictor values*

1	No prediction (Default value)
2	TIFF predictor 2
10	PNG prediction (on encode, PNG <i>None</i> on all rows)
11	PNG prediction (on encode, PNG <i>Sub</i> on all rows)
12	PNG prediction (on encode, PNG <i>Up</i> on all rows)
13	PNG prediction (on encode, PNG <i>Average</i> on all rows)
14	PNG prediction (on encode, PNG <i>Paeth</i> on all rows)
15	PNG prediction (on encode, PNG optimum)

The two predictor function groups also differ in significant ways. First, the postprediction data for each PNG-predicted row begins with an explicit algorithm tag, so different rows can be predicted with different algorithm to improve compression. TIFF 2 prediction has no such identifier; the same algorithm applies to all rows. Second, the TIFF function group predicts each color component from the prior instance of that color component, without regard to the width of the color component or the number of colors. In contrast, the PNG function group predicts each byte from the corresponding byte of the prior pixel (and/or the same pixel on the prior line and/or the prior pixel on the prior line), regardless of whether there are multiple color components in a byte, or whether a single color component spans multiple bytes. This can yield significantly better speed at a cost somewhat worse compression.

4.8.6 Comparison of LZW and Flate encoding

PDF 1.2

Flate encoding, like LZW encoding, discovers and exploits many patterns in its input data, whether text or images. Thanks to its cascaded adaptive Huffman coding, Flate-encoded output is usually substantially more compact than LZW-encoded output for the same input. Flate and LZW decoding speeds are comparable, but Flate encoding speed is considerably slower than LZW encoding speed. Usually, both Flate and LZW compress their inputs substantially. In the worst case, however, Flate encoding expands its input by no more than 11 bytes, plus the effects of algorithm tags added by PNG predictors. LZW encoding has a worst-case expansion of at least a factor of 1.125, which can increase to a factor of nearly 1.5 in some implementations (plus PNG tags effects, as with Flate encoding).

4.8.7 RunLengthDecode filter

This filter decodes data that has been encoded in a simple byte-oriented, run-length-encoded format. Run-length encoding produces binary data, even if the original data was ASCII text.

The compression achieved by run-length encoding depends on the input data. In the best case, a file of all zeros, a compression of approximately 64:1 is achieved for long files. The worst case, the hexadecimal sequence of alternating <00 FF 00 FF ...>, results in an expansion of 127:128.

The encoded data is a sequence of runs, where each run consists of a *length* byte followed by 1 to 128 bytes of data. If *length* is in the range 0 to 127, the following *length* + 1 (1 to 128 bytes) are copied literally during decompression. If *length* is in the range 129 to 255, the following single byte is to be copied 257 – *length* times (2 to 128 times) during decompression. The value 128 is placed at the end of the compressed data, as an EOD marker.

4.8.8 CCITTFaxDecode filter

This filter decodes image data that has been encoded using either Group 3 or Group 4 CCITT facsimile (fax) encoding. This filter is useful only for bitmap image data, not for color images, grayscale images, or text. Group 3 and Group 4 CCITT encoding produce binary data that may be converted, if necessary, to 7-bit data using either the ASCII hexadecimal or ASCII base-85 encodings, described in previous sections.

The compression achieved using CCITT encodings depends on the data, as well as on the value of various optional parameters. For Group 3 one-dimensional encoding, the best case is a file of all zeros. In this case, each scan line compresses to 4 bytes, and the compression factor depends on the length of a scan line. If the scan line is 300 bytes long, a compression ratio of approximately 75:1 is achieved. The worst case, an image of alternating ones and zeros, produces an expansion of 2:9.

The *Comité Consultatif Internationale de Télégraphie et Téléphonie* (International Coordinating Committee for Telephony and Telegraphy, or CCITT) is an international standards organization that defines CCITT encoding. The encoding is designed to achieve efficient compression of monochrome (1 bit per sample) image data at relatively low resolutions. The algorithm is not described in detail here, but can be found in the CCITT standards, [19] and [20], listed in the Bibliography on [page 506](#).

The fax encoding method is bit-oriented, rather than byte-oriented. This means that, in principle, encoded or decoded data may not end on a byte boundary. The filter addresses this in the following ways:

- Encoded data are ordinarily treated as a continuous, unbroken bit stream. However, the **EncodedByteAlign** parameter (described in [Table 4.7](#)) can be used to cause each encoded scan line to be filled to a byte boundary. Although this is not prescribed by the CCITT standard and fax machines don't do this, some software packages find it convenient to encode data this way.

- When a filter reaches EOD, it always skips to the next byte boundary following the encoded data.

Both Group 3 and Group 4 encoding, as well as optional features of the CCITT standard, are supported. The optional parameters that can be used to control the decoding are listed in Table 4.7. Except as noted, all values supplied to the decode filter by the optional parameters must match those used when the data was encoded.

Table 4.7 *Optional parameters for CCITTFaxDecode filter*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
K	integer	Selects the encoding scheme used. A negative value indicates pure two-dimensional (Group 4) encoding. Zero indicates pure one-dimensional (Group 3, 1-D) encoding. A positive value indicates mixed one- and two-dimensional encoding (Group 3, 2-D) in which a line encoded one-dimensionally can be followed by at most K – 1 lines encoded two-dimensionally. The decoding filter distinguishes between negative, zero, and positive values of K , but does not distinguish between different positive K values. The default value is 0.
EndOfLine	boolean	End-of-line bit patterns are always accepted, but they are required if EndOfLine is <i>true</i> . The default value is <i>false</i> .
EncodedByteAlign	boolean	If <i>true</i> , each encoded line must begin on a byte boundary. The default value is <i>false</i> .
Columns	integer	Specifies the width of the image in samples. If Columns is not a multiple of 8, the width of the unencoded image is adjusted to the next multiple of 8, so that each line starts on a byte boundary. The default value is 1728.
Rows	integer	Specifies the height of the image in scan lines. If this parameter is zero or is absent, the height of the image is not predetermined and the encoded data must be terminated by an end-of-block bit pattern or by the end of the filter's data source. The default value is 0.
EndOfBlock	boolean	If <i>true</i> , the data is expected to be terminated by an end-of-block, overriding the Rows parameter. If <i>false</i> , decoding stops when Rows lines have been decoded or when the data has been exhausted, whichever occurs first. The end-of-block pattern is the CCITT end-of-facsimile-block (EOFB) or return-to-control (RTC) appropriate for the K parameter. The default value is <i>true</i> .
BlackIs1	boolean	If <i>true</i> , causes bits with value 1 to be interpreted as black pixels and bits with value zero to be interpreted as white pixels. The default value is <i>false</i> .
DamagedRowsBeforeError	integer	If DamagedRowsBeforeError is positive, EndOfLine is <i>true</i> , and K is non-negative, then up to DamagedRowsBeforeError rows of data are tolerated before an error is generated. Tolerating a damaged row means locating its end in

the encoded data by searching for an **EndOfLine** pattern, and then substituting decoded data from the previous row if the previous row was not damaged or a white scan line if the previous row was damaged. The default value is 0.

4.8.9 DCTDecode filter

This filter decodes grayscale or color image data that has been encoded in the JPEG baseline format. JPEG encoding produces binary data.

JPEG is a lossy compression method, meaning that some of the information present in the original image is lost when the image is encoded. Because of the information loss, only images (never text) should be encoded in this format. The compression achieved using the JPEG algorithm depends on the image being compressed and the amount of loss that is acceptable. In general, a compression of 15:1 can be achieved without a perceptible loss of information, and 30:1 compression causes little impairment of the image.

During encoding, several optional parameters control the algorithm and the information loss. The values of these parameters are stored in the encoded data, and the decoding filter generally obtains the parameter values it requires directly from the encoded data. A description of the parameters accepted by the encoding filter can be found in Section 3.13.3 of the *PostScript Language Reference Manual, Third Edition* [1].

JPEG stands for the ISO/CCITT Joint Photographic Experts Group, an organization responsible for developing an international standard for compression of color image data. The encoding method uses the discrete cosine transform (DCT). Data to be encoded consists of a stream of image samples, each containing one, two, three, or four color components. The color component values for a particular sample must appear consecutively. Each component value occupies an 8-bit byte.

The details of the encoding algorithm are not presented here but can be found in the references [30] and [34] listed in the Bibliography. Briefly, the JPEG algorithm breaks an image up into blocks of 8×8 samples. Each color component in an image is treated separately. A two-dimensional DCT is performed on each block. This operation produces 64 coefficients, which are then quantized. Each coefficient may be quantized with a different step size. It is the quantization that results in the loss of information in the JPEG algorithm. The quantized coefficients are then compressed.

The amount of loss incurred in JPEG encoding is controlled by the encoding filter, which can reduce the loss by making the step size in the quantization smaller at the expense of reducing the amount of compression achieved by the algorithm. The JPEG filter implementation in the Acrobat products does not support features of the JPEG standard that are not relevant. In addition, certain choices regarding reserved marker codes and other optional features of the standard have been made.

In addition to the baseline JPEG format, in PDF 1.3 the DCTDecode filter supports the progressive JPEG extension. This extension does not entail any additional attributes in the DCTDecode parameter dictionary; the distinction between baseline and progressive JPEG is represented in the encoded data.



Note There is no benefit from using progressive JPEG for stream data that is embedded in a PDF file. Decoding progressive JPEG is slower and consumes more memory than baseline JPEG. The purpose of this feature is to enable a stream to refer to an external file whose data happens to be encoded in progressive JPEG already.



4.9 The null object

The keyword **null** represents the null object.

Note The value of a dictionary key can be specified as **null**. A simpler but equivalent way to express this is to omit the key from the dictionary.

4.10 Indirect objects

A *direct object* is a boolean, number, string, name, array, dictionary, stream, or null, as described in the previous sections. An *indirect object* is an object that has been labeled so that it can be referenced by other objects. Any type of object may be labeled as an indirect object. Indirect objects are very useful; for example, the value of the stream's **Length** key may be specified as an indirect object that is stored in the file after the stream. This is useful for applications that generate PDF in a single pass.

An indirect object consists of an object identifier, a direct object, and the **endobj** keyword. The *object identifier* consists of an integer *object number*, an integer *generation number*, and the **obj** keyword:

```
<indirect object> ::=
    <object ID>
    <direct object>
    endobj

<object ID> ::= <object number>
                <generation number>
                obj
```

The combination of object number and generation number serves as a unique identifier for an indirect object. Throughout its existence, an indirect object retains the object number and generation number it was initially assigned, even if the object is modified.

Each indirect object has a unique object number, and indirect objects are often but not necessarily numbered sequentially in the file, beginning with 1. Until an object in the file is deleted, all generation numbers are 0.

4.11 Object references

Any object used as an element of an array or as a value in a dictionary may be specified by either a direct object or an indirect reference. An *indirect reference* is a reference to an indirect object, and consists of the indirect object's object number, generation number, and the **R** keyword:

```
<indirect reference> ::=
    <object number>
    <generation number>
    R
```

Using an indirect reference to the stream's length, a stream could be written in this way:

Example 4.5 *Indirect reference*

```
7 0 obj
<<
/Length 8 0 R           An indirect reference to object 8.
>>
stream
BT
/F1 12 Tf
72 712 Td (A stream with an indirect Length) Tj
ET
endstream
endobj

8 0 obj
64                       This is the Length of the previous stream.
endobj
```

Note An indirect reference to an undefined object is not an error; it is treated as a reference to the null object. For example, if a PDF file contains the indirect reference (**12 0 R**) but does not contain the definition (**12 0 obj ... endobj**), then the indirect reference is null.

CHAPTER 5

File Structure

This chapter describes the overall organization of a PDF file. A PDF file provides a structure that represents a document. This structure provides a way to rapidly access any part of a document and a mechanism for updating it.

The body of a PDF file contains a sequence of PDF objects that are used to construct a document.

A particular arrangement of the object in a PDF file, optimized for incremental access in network environments, is known as *Linearized PDF*. It is described in [Chapter 9](#).

PDF 1.2


5.12 PDF files

A canonical PDF file consists of four sections: a one-line header, a body, a cross-reference table, and a trailer. [Figure 5.1](#) shows this structure:

```
<PDF file> ::= <header>
                <body>
                <cross-reference table>
                <trailer>
```

A line in a PDF file is terminated by a carriage return (<0D>), a linefeed (<0A>), or a carriage return followed by a linefeed. Updates may be appended to a PDF file, as described in [Section 5.17](#), “[Incremental update](#).”

PDF files with binary data may have arbitrarily long lines. However, to increase compatibility with other applications that process PDF files, all lines that are not part of stream object data shall be no longer than 255 characters. An exception is made in the case of the Contents string of a Signature dictionary in PDF 1.3. (See [page 141](#).)

Figure 5.1 *Structure of a PDF file that has not been updated*

5.13 Header

The first line of a PDF file specifies the version number of the PDF specification to which the file adheres. The current version is 1.3; the first line of a 1.3-conforming PDF file should be **%PDF-1.3**. However, a PDF file that conforms to an earlier version also conforms to 1.3, so an application that understands PDF 1.3 also accepts a file whose first line is either **%PDF-1.0**, **%PDF-1.1**, or **%PDF-1.2**.

`<header> ::= <PDF version>`

5.14 Body

The body of a PDF file consists of a sequence of indirect objects representing a document. The objects, which are of the basic types described in Chapter 4, represent components of the document such as fonts, pages, and sampled images.

Comments can appear anywhere in the body section of a PDF file. Comments have the same syntax as those in the PostScript language; they begin with a **%** character and may start at any point on a line. All text between the **%** character and the end of the line is treated as a comment. Occurrences of the **%** character within strings or streams are not treated as comments.

5.15 Cross-reference table

The cross-reference table contains information that permits random access to indirect objects in the file, so that the entire file need not be read to locate any particular object. For each indirect object in the file, the table contains a one-line entry describing the location of the object in the file.

A PDF file contains one cross-reference table, consisting of one or more *sections*. If no updates have been appended to the file, the cross-reference table contains a single section. One section is added each time updates are appended to the file.

The cross-reference section is the only part of a PDF file with a fixed format. This permits random access to entries in the cross-reference table. The section begins with a line containing the keyword **xref**. Following this line are one or more cross-reference subsections:

```
<cross-reference section> ::=
    xref
    <cross-reference subsection>+
```

Note If no updates have been appended to the file, the cross-reference section may contain only one subsection.

Each subsection contains entries for a contiguous range of object numbers. The organization of the cross-reference section into subsections is useful for incremental updates, because it allows a new cross-reference section to be added to the PDF file, containing entries only for objects that have been added or deleted. Each cross-reference subsection begins with a header line containing two numbers: the first object number in that subsection and the number of entries in the subsection. Following the header are the entries, one per line:

```
<cross-reference subsection> ::=
    <object number of first entry in subsection>
    <number of entries in subsection>
    <cross-reference entry>+
```

Each entry is exactly 20 bytes long, including the end-of-line marker. There are two formats for cross-reference table entries: one for objects that are in use and another for objects that have been deleted and so are free:

```
<cross-reference entry> ::= <in-use entry> | <free entry>
```

For an object that is in use, the entry contains a byte offset specifying the number of bytes from the beginning of the file to the beginning of the object, the generation number of the object, and the **n** keyword:

```
<in-use entry> ::= <byte offset> <generation number> n <end-of-line>
```

The byte offset is a ten-digit number, padded with leading zeros if necessary. It is separated from the generation number by a single space. The generation number is a five-digit number, also padded with leading zeros if necessary. Following the generation number is a single space and the **n** keyword. Following the keyword is the end-of-line sequence. If the end-of-line is a single character (either a carriage return or linefeed), it is preceded by a single space. If the end-of-line sequence is two characters (a carriage return followed by a linefeed), it is not preceded by a space.

```
<end-of-line> ::= <space> <carriage return>
                | <space> <linefeed>
                | <carriage return> <linefeed>
```

For an object that is free, the entry contains the object number of the next free object, a generation number, and the **f** keyword:

```
<free entry> ::=
                <object number of next free object>
                <generation number> f <end-of-line>
```

The entry has the same format as that for an object that is in use: a ten-digit object number, a space, a five-digit generation number, a space, the **f** keyword, and an end-of-line sequence.

The free objects in the cross-reference table form a linked list, with the entry for each free object containing the object number of the next free object. The first entry in the table (object number 0) is always free and has a generation number of 65535. It is the head of the linked list of free objects. The last free entry in the cross-reference table (the tail of the linked list) uses 0 as the object number of the next free object.

When an indirect object is deleted, its cross-reference entry is marked free, and the generation number in the entry is incremented by one to record the generation number to be used the next time an object with that object number is created. Each time the entry is reused, its generation number is incremented. The maximum generation number is 65535. Once that number is reached, that entry in the cross-reference table will not be reused.

The cross-reference table (comprising the original table and all the updates) must contain one entry for each object number between zero and the maximum object number used in the file, even if one or more object numbers are never used in the file.

[Example 5.1](#) shows a cross-reference section containing a single subsection with six entries; four that are in use (object numbers 1, 2, 4, and 5) and two that are free (object numbers 0 and 3). Object number 3 has been deleted, and the next object created with an object number of 3 will be given the generation number of 7.

Example 5.1 *Cross-reference section with a single subsection*

```

xref
0 6
0000000003 65535 f
0000000017 00000 n
0000000081 00000 n
0000000000 00007 f
0000000331 00000 n
0000000409 00000 n

```

[Example 5.2](#) shows a cross-reference section with four subsections containing a total of five entries. The first subsection contains one entry, for object number 0, which is free. The second subsection contains one entry, for object number 3, which is in use. The third subsection contains two entries, for objects number 23 and 24, both of which are in use. Object number 23 has been reused, as can be seen from the fact that it has a generation number of 2. The fourth subsection contains one entry, for object number 30, which is in use.

Example 5.2 *Cross-reference section with multiple subsections*

```

xref
0 1
0000000000 65535 f
3 1
0000025325 00000 n
23 2
0000025518 00002 n
0000025635 00000 n
30 1
0000025777 00000 n

```

[Appendix A](#) contains a more extensive example of the structure of a PDF file after several updates have been made to it.

5.16 Trailer

The trailer enables an application reading a PDF file to quickly find the cross-reference table and certain special objects. Applications should read a PDF file from its end. The last line of a PDF file contains only the end-of-file marker, **%%EOF**. The two preceding lines contain the keyword **startxref** and the byte offset from the beginning of the file to the beginning of the word **xref** in the last cross-reference section in the file. The *trailer dictionary* precedes this line.

The trailer dictionary, shown in [Table 5.1](#), consists of the keyword **trailer** followed by a set of key–value pairs enclosed in double angle brackets:

```

<trailer> ::= trailer
<<
<trailer key–value pair>+

```

```

>>
startxref
<cross-reference table start address>
%%EOF

```

Table 5.1 *Trailer attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Size	integer	<i>(Required)</i> Total number of entries in the file's cross-reference table, including the original table and all updates.
Prev	integer	<i>(Present only if the file has more than one cross-reference section)</i> Byte offset from the beginning of file to the location of the previous cross-reference section. If the file has never been updated, it will not contain the Prev key.
Root	dictionary	<i>(Required; must be indirect reference)</i> Catalog object for the document, described on page 67 .
Info	dictionary	<i>(Optional; must be indirect reference)</i> Info dictionary for the document, described on page 120 .
ID	array	<i>(Optional)</i> An array of two strings, each of which is an ID. The first ID is established when the file is created and the second ID is changed each time the file is updated. IDs are described on page 123 .
Encrypt	dictionary	<i>(Required if document is encrypted)</i> Information used to decrypt a document, described on page 124 .

PDF 1.1

PDF 1.1

An example trailer for a file that has not been updated is shown in [Example 5.3](#). The fact that the file has not been updated is determined from the absence of a **Prev** key in the trailer dictionary.

Example 5.3 *Trailer*

```

trailer
<<
  /Size 22
  /Root 2 0 R
  /Info 1 0 R
>>
startxref
18799
%%EOF

```

5.17 Incremental update

The contents of a PDF file can be updated without rewriting the entire file. Changes can be appended to the end of the file, leaving completely intact the original contents of the file. When a PDF file is updated, any new or changed objects are appended, a cross-reference section is added, and a new trailer is inserted. The resulting file has the structure shown in [Figure 5.2](#):

```

<Updated PDF file> ::=
    <PDF file>
    {<update>}*

<update> ::= <body>
             <cross-reference section>
             <trailer>


```

A complete example of an updated file is shown in [Appendix A](#).

The cross-reference section added when a PDF file is updated contains entries only for objects that have been changed, replaced, or deleted, plus the entry for object 0. Deleted objects are left unchanged in the file, but are marked as deleted in their cross-reference entries. The trailer that is added contains all the information in the previous trailer, as well as a **Prev** key specifying the location of the previous cross-reference section. As shown in [Figure 5.2](#), after a file has been updated several times it contains several trailers, as well as several **%%EOF** lines.

Because updates are appended to PDF files, it is possible to end up with several copies of an object with the same object ID (object number and generation number) in a file. This occurs, for example, if a text annotation is changed several times, with the file being saved between changes. Because the text annotation object is not deleted, it retains the same object number and generation number. Because it has been changed, however, an updated copy of the object is included in the update section added to the file. The cross-reference section added includes a pointer to this new changed version, overriding the information contained in the original cross-reference section. When a program such as Acrobat reads the file, it must build cross-reference information in such a way that the most recent version of an object is accessed in the file.

Figure 5.2 *Structure of a PDF file after changes have been appended several times*



5.18 Encryption

PDF 1.1

Documents can be encrypted to protect their contents from unauthorized access. Access to a protected document's contents is controlled by the security handler specified in the Encryption dictionary. The Encryption dictionary is the value of the **Encrypt** key in the trailer dictionary. [Section 6.13, "Encryption dictionary,"](#) describes the Encryption dictionary and security handlers.

Strings and streams in a protected document are encrypted. Other data types (such as integers and booleans) that are used primarily for structural information in a PDF file are not encrypted. This combination protects a document's contents, while allowing random access to the objects within a PDF file.



All strings and streams in a protected document, except those in the Encryption dictionary, are encrypted using the RC4 encryption algorithm. This prevents unauthorized users from simply removing the password from a PDF file to gain access to it. Strings in the Encryption dictionary are encrypted and decrypted by the security handler itself, using whatever encryption algorithm it chooses.

Note External stream data is not encrypted, since it is not part of the PDF document. However, if an external stream refers to a file descriptor for a file that is embedded within the PDF file, the embedded file stream will be encrypted.

RC4 is a copywritten, proprietary algorithm of RSA Data Security, Inc. Adobe Systems has licensed this algorithm for use in its Acrobat products. Independent software vendors may be required to license RC4 to develop software that encrypts or decrypts PDF documents. For further information please send e-mail to products@rsa.com or visit the RSA web server at <http://www.rsa.com/>.

Encryption details

Protection of data in a PDF file consists of two steps: computation of a key to be used to encrypt data, and encryption of the data. The key is simply a string of five bytes. (The key is restricted in length to five bytes (40 bits) to satisfy current U.S. cryptographic export requirements.) The key can be computed in any number of ways, more or less cryptographically secure, while encryption of data based on the key is always performed in the same way. The Encryption dictionary identifies the security handler that computes the key. PDF includes one built-in method for computing the key, called “Standard.” The remainder of this section explains how data is encrypted given a key. [Section 6.13, “Encryption dictionary](#), discusses security handlers, including the Standard security handler.

Once the key for a document is computed, strings and streams are encrypted using [Algorithm 5.3](#):

Algorithm 5.3 *Encrypting string and stream data*

1. Extend the key by five bytes using the string or stream’s object identifier. (See [Section 4.10, “Indirect objects.”](#)) If the string or stream is a direct object, the identifier of the indirect object containing it is used. An object identifier consists of an object number and a generation number. These are treated as binary integers. The low-order three bytes of the object number and the low-order two bytes of the generation number are concatenated to the key in that order, low-order byte first.
2. The resulting ten-byte string is used as input for the MD5 hash function.
3. The first ten bytes of the output of the MD5 function are used as input to the RC4 function, along with the string or stream data to be encrypted. The output of the RC4 function is the encrypted data that is stored in the PDF file.

RC4 is a symmetric stream cipher—the same algorithm is used for both encryption and decryption, and the algorithm does not change the length of the data. The PDF encryption algorithm is also symmetric. Given a key, the three steps described above can also be used to decrypt data.

Stream data is encrypted after all stream encoding filters have been applied (and is decrypted before the stream decoding filters are applied). Decryption of strings, other than those in the Encryption dictionary, is done after escape-sequence processing and hex decoding as appropriate to the string representation described in [Section 4.4, “Strings and text.”](#)

CHAPTER 6

Document Structure

PDF provides an electronic representation of a document—a series of pages containing text, graphics, and images, along with other information such as thumbnails (miniature images of the pages), text annotations, hypertext links, and outline entries (also called *bookmarks*). Previous chapters lay the groundwork for understanding the PDF representation of a document, but do not describe the representation itself.

The body of a PDF file consists of a sequence of objects that collectively represent a PDF document. This chapter focuses exclusively on the contents of the body section of a PDF file and contains a description of each type of object that may be contained in a PDF document. Following each description is an example showing the object as it might appear in a PDF file. Complete example PDF files appear in [Appendix A](#).


6.1 Introduction

A PDF document can be described as a hierarchy of objects contained in the body section of a PDF file. [Figure 6.1](#) shows the structure of a PDF document. Most objects in this hierarchy are dictionaries. Parent, child, and sibling relationships are represented by key–value pairs whose values are indirect references to parent, child, or sibling objects. For example, the Catalog object, which is the root of the hierarchy, contains a **Pages** key whose value is an indirect reference to the object that is the root of the Pages tree.

Each page of the document includes references to its contents, its thumbnail, and any annotations that appear on the page. The PDF file’s standard trailer, described in [Section 5.16](#), “Trailer,” specifies the location of the Catalog object as the value of the trailer’s **Root** key. In addition, the trailer specifies the location of the document’s Info dictionary, a structure that contains general information about the document, as the value of the trailer’s **Info** key.

Note In many of the tables in this chapter, certain key–value pairs contain the notation “must be an indirect reference” or “indirect reference preferred.” Unless one of these is specified in the description of the key–value pair, objects that are the value of a key can either be specified directly or using an indirect reference, as described in [Section 4.11](#), “Object references.”

Figure 6.1 Structure of a PDF document



6.2 Catalog

The Catalog is a dictionary that is the root node of the document. It contains a reference to the tree of pages contained in the document, a reference to the tree of objects representing the document's outline, a reference to the document's article threads, and the list of named destinations. In addition, the Catalog indicates whether the document's outline or thumbnail page images should be displayed automatically when the document is viewed and whether some location other than the first page should be shown when the document is opened. [Example 6.1](#) shows a sample Catalog object.

PDF 1.1

Example 6.1 *Catalog*

```
1 0 obj
<<
  /Type /Catalog
  /Pages 2 0 R
  /Outlines 3 0 R
  /PageMode /UseOutlines
>>
endobj
```

[Table 6.1](#) shows the attributes for a Catalog.

Table 6.1 *Catalog attributes*

Key	Type	Semantics
Type	name	Object type. Always Catalog .
Dests	dictionary	<i>(Required in PDF 1.1 if the document has named destinations; must be an indirect reference)</i> A dictionary of names and corresponding destinations; see page 184 .
Outlines	dictionary	<i>(Required if the document has an outline; must be an indirect reference)</i> The Outlines object that is the root of the document's outline tree, described on page 104 .
Pages	dictionary	<i>(Required, must be an indirect reference)</i> Pages object that is the root of the document's Pages tree; see page 71 .
Threads	array	<i>(Required if the document has any threads; must be an indirect reference)</i> An array of threads as described on page 121 .
AcroForm	dictionary	<i>(Optional)</i> A dictionary that defines the AcroForm database contained in this file. See page 129 .
Names	dictionary	<i>(Optional)</i> A dictionary that contains lists of various types of names and strings to be referenced within the document. See page 119 .
OpenAction	array or dictionary	<i>(Optional)</i> If the value of this key is an array, it must be a destination; see page 184 . If the value is a dictionary, it must be an action; see page 107 . If this key is omitted, the top of the first page appears at the default zoom level.

PDF 1.1

PDF 1.3

PDF 1.2

PDF 1.1

PageMode	name	<i>(Optional)</i> How the document should appear when opened. Allowed values:	
			<p>UseNone Open document with neither outline nor thumbnails visible. This is the default value.</p> <p>UseOutlines Open document with outline visible.</p> <p>UseThumbs Open document with thumbnails visible.</p> <p>FullScreen Open document in full-screen mode. In full-screen mode, there is no menu bar, window controls, nor any other window present.</p>
URI	dictionary	<i>(Optional)</i> Contains document-level information for Uniform Resource Identifier annotations; see page 113 .	
ViewerPreferences			
	dictionary	<i>(Optional)</i> Specifies a dictionary that contains opening display options for this document; see page 69 . If this key is omitted, viewers behave in accordance with any current user preferences. The name of the key reflects the fact that this dictionary is not part of the document structure itself, but represents a set of viewer-level options for displaying this document. A given viewer implementation may or may not support the options in this dictionary.	
PageLabels	Number tree	<i>(Optional)</i> Specifies the page labeling for the entire document. The keys in this number tree are page indices, and the values are PageLabel dictionaries (see page 70). The page index indicates the first page to which the related PageLabel dictionary applies. If the PageLabels tree is present, there must be a 0 entry.	PDF 1.3
PageLayout	name	<i>(Optional)</i> Specifies the layout for the page when the document is opened. If this attribute is not present, viewers behave in accordance with the current user preference. Allowed values:	
			<p>SinglePage Display the pages one page at a time.</p> <p>OneColumn Display the pages in one column.</p> <p>TwoColumnLeft Display the pages in two columns, with odd-numbered pages on the left.</p> <p>TwoColumnRight Display the pages in two columns, with odd-numbered pages on the right.</p>
		<i>Note</i> In a previous version of the manual, the <i>PageLayout</i> key was documented as being in the <i>Viewer Preferences</i> dictionary. The <i>PageLayout</i> key was actually implemented in the <i>Catalog</i> .	
JavaScript	number	<i>(Optional)</i> Default is 1.2. The minimum version of the JavaScript interpreter necessary to execute JavaScript actions found in this document.	PDF 1.3
StructTreeRoot	dictionary	<i>(Optional)</i> The structure-tree root object for the document. See page 163 .	PDF 1.3
SpiderInfo	dictionary	<i>(Optional)</i> Specifies a dictionary that contains state information for Web Capture (also known as AcroSpider); see Table 6.78 on page 157 .	PDF 1.3

Note In PDF 1.2, the **AA** key was also defined as an Additional Actions attribute. It was never implemented. In PDF 1.3, the **AA** key is obsolete and should be ignored.

PDF 1.3

6.2.1 Viewer preferences

The viewer preferences dictionary has the following attributes:

Table 6.2 *Viewer Preferences*

PDF 1.2

Key	Type	Semantics
HideToolbar	boolean	<i>(Optional)</i> Specifies that the viewer's toolbar should be hidden whenever the document is active. This attribute defaults to <i>false</i> .
HideMenubar	boolean	<i>(Optional)</i> Specifies that the viewer's menubar should be hidden whenever the document is active. This attribute defaults to <i>false</i> .
HideWindowUI	boolean	<i>(Optional)</i> Specifies that the user interface elements in the document's window should be hidden. This attribute defaults to <i>false</i> .
FitWindow	boolean	<i>(Optional)</i> Specifies that the viewer should resize the window displaying the document to fit the size of the first displayed page of the document. This attribute defaults to <i>false</i> .
CenterWindow	boolean	<i>(Optional)</i> Specifies that the viewer should position the window displaying the document in the center of the computer's monitor. This attribute defaults to <i>false</i> .
NonFullScreenPageMode	name	<i>(Optional)</i> Specifies how the document should be displayed after exiting full-screen mode if the value of the PageMode key in the Catalog is FullScreen . This key is ignored if the value of the PageMode key in the Catalog is not FullScreen . Allowed values and semantics are the same as for the PageMode key in the Catalog, except that a value of FullScreen is not allowed.
Direction	name	<i>(Optional)</i> This key defines the predominant reading order for text. If present the values must be either R2L or L2R , indicating right-to-left or left-to-right reading order, respectively. R2L should be supplied for documents that use languages or writing systems that use text that is read from right to left such as Arabic, Hebrew, and Chinese, Korean, and Japanese that use vertical writing. If present this entry will be used when displaying or printing multiple pages. The default value is L2R .

PDF 1.3

Note In a previous version of the manual, the **PageLayout** key was documented as being in the Viewer Preferences dictionary. The **PageLayout** key was actually implemented in the Catalog.

In languages such as Arabic, Hebrew, or vertical writing in Chinese, Korean and Japanese, text is read from right to left. The pages in documents that use these languages are also read starting from the right-hand page. The **Direction** key can be used to identify documents that use right-to-left reading order. It is used by the Acrobat viewer to determine the sequence of pages where more than one page is

displayed side by side on the screen at one time. It can also be used in *n-up* printing. Reading direction has no effect on the contents of a PDF file, or on the page numbering.

6.2.2 PageLabel dictionaries

PDF 1.3

The term *page label* to refer to the short, generally numeric string used to identify a page in the viewer user interface. The term *page index* refers to the position of a page relative to the first page of the document, which has a page index of 0. For example, the first page of a document which begins using Roman numbers would have a page label of “i” and a page index of 0. In the same document the first page of the first chapter would have a page label of “1” and a page index greater than 0.

A *labeling range* is a contiguous set of pages which use the same numbering system and prefix. All of the pages in a labeling range are numbered in ascending order.

A page label is composed of a *numeric portion* and an optional *label prefix*. The numeric portion of each label is automatically generated based on the position of the page within its labeling range. The label prefix is a string that is prefixed to the numeric portion to form the complete page label. For example, the pages in an appendix may be labeled with decimal numeric portions prefixed with the string “A-”. The resulting page labels would be “A-1”, “A-2”, etc.

A **PageLabel** dictionary specifies the labeling system to be used for a labeling range.

Table 6.3 *PageLabel dictionary entries*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Always PageLabel .
S (Style)	name	<p>(<i>Optional</i>) Specifies the numbering system to use for the numeric portion of each label in this range of pages. Possible values are:</p> <p>D for decimal numbers R for upper-case Roman numbers r for lower-case Roman numbers A for upper-case alphabetic numbers a for lower-case alphabetic numbers</p> <p>If this key is not present, the labels for this range will not have a numeric portion. The display will consist solely of the prefix string with no number appended. This allows pages to be named without numbers. For example,</p>

<< /P (Contents) >>

will cause a page’s label to be just the string “Contents”. If the prefix string is also missing or empty, the page’s label will be an empty string. When a valid PageLabel dictionary is present, the viewer never uses default numbering; the PageLabel dictionary specifies every label of every page.

P (Prefix)	Text	(Optional) Specifies a string to be prefixed to the numeric portion to create the page label.
St (StartAt)	integer	(Optional) Specifies the value to use when generating the numeric portion of the first label in this range. This number must be greater than or equal to 1. The default value is 1.

For ranges with an **S** key, the value of the numeric portion of the first label in the range is specified by the **St** key, which defaults to 1. Subsequent pages in the range will be numbered in ascending order. The label is constructed by expressing the numeric portion as a string in the numbering system specified by the **S** key and then appending the prefix string, if any, to the front.

Labels for ranges with no **S** key consist solely of the prefix string.

Alphabetic numbering uses 'a' through 'z' for 1 through 26, 'aa' through 'zz' for 27 through 52, etc.

The following example shows a document with pages labeled

i, ii, iii, iv, 1, 2, 3, A-8, A-9, ...

Example 6.2 *Document with pages labeled*

```

1 0 obj <<
  /Type /Catalog
  /PageLabels <<
    A "number tree" with 3 Page Label dictionaries
    /Nums [
      0 << /S /r >>
      4 << /S /D >>
      7 << /S /D /P (A-) /St 8 >>
    ]
  >>
  ...
  >>
endobj

```

6.3 Pages tree

The pages of a document are accessible through a tree of nodes known as the Pages tree. This tree defines the ordering of the pages in the document.

To optimize the performance of viewer applications, the Acrobat Distiller program and Acrobat PDF Writer construct balanced trees. (For further information on balanced trees, see reference [15] in the Bibliography on page 506.) The tree structure allows applications to quickly open a document containing thousands of pages using only limited memory. Applications should accept any sort of tree

structure as long as the nodes of the tree contain the keys described in [Table 6.4](#). The simplest structure consists of a single Pages node that references all the page objects directly.

Note The structure of the Pages tree for a document is unrelated to the content of the document. In a PDF file for a book, for example, there is no guarantee that a chapter is represented by a single node in the Pages tree. Applications that consume or produce PDF files are not required to preserve the existing structure of the Pages tree.

The root and all interior nodes of the Pages tree are dictionaries, whose minimum contents are shown in [Table 6.4](#).

Table 6.4 Pages attributes

Key	Type	Semantics
Type	name	(Required) Object type. Always Pages .
Kids	array	(Required) List of indirect references to the immediate children of this Pages node.
Count	integer	(Required) Specifies the number of leaf nodes (imageable pages) under this node. The leaf nodes do not have to be immediately below this node in the tree, but can be several levels deeper in the tree.
Parent	dictionary	(Required; must be indirect reference) Pages object that is the immediate ancestor of this Pages object. The root Pages object has no Parent .

The following illustrates the Pages object for a document with three pages. [Appendix A](#) contains an example showing the Pages tree for a document containing 62 pages.

Example 6.3 Pages tree for a document containing three pages

```

2 0 obj
<<
  /Type /Pages
  /Kids [4 0 R 10 0 R 24 0 R]
  /Count 3
>>
endobj

```

6.3.1 Inheritance of Attributes

A Pages object may contain additional keys that provide values for Page objects that are its descendants. Such values are said to be “inherited.” For example, a document may specify a **MediaBox** for all pages by defining one in the root Pages object. An individual page in the document could override the **MediaBox** in this example by specifying a **MediaBox** in the Page object for that page.

Attributes that may be inherited are indicated in [Table 6.5](#). If a *required* key that may be inherited is omitted from a Page object, then a value must be supplied in one of its ancestors. If an *optional* key that may be inherited is omitted, then a value may be supplied in one of its ancestors; barring that, the default value is used.

[Example 6.4](#) demonstrates inheritance by showing a tree of Pages objects and Page objects. Pages 1, 2, and 4 are rotated 90°. Page 3 is rotated 270°. Pages 5 and 7 are not rotated (rotated 0°). Page 6 is rotated 180°.

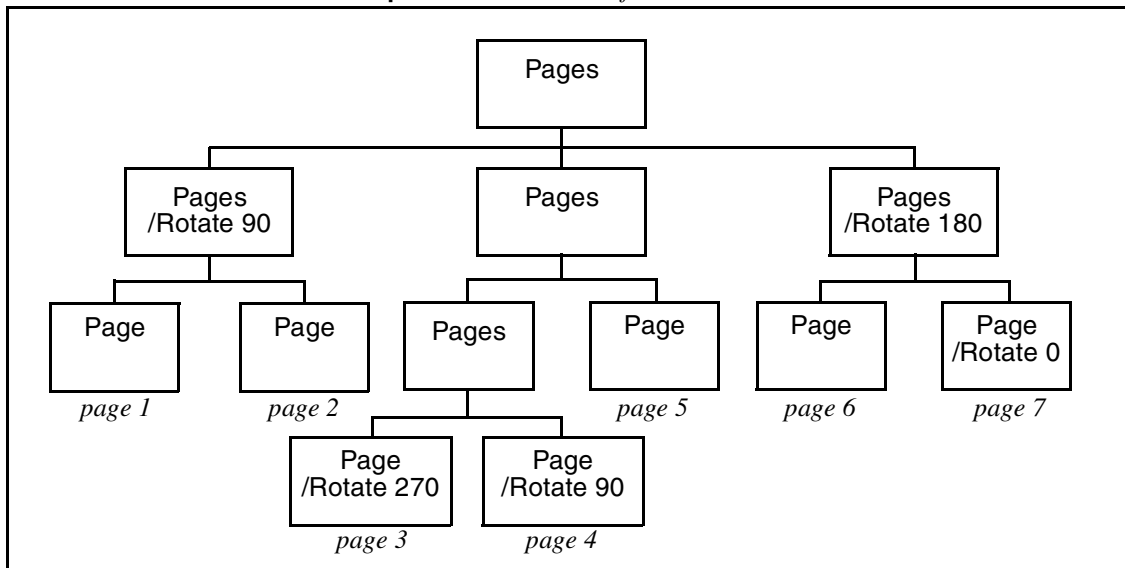
6.4 Page objects

A Page object is a dictionary that describes a single page containing text, graphics, and images. A Page object is a leaf of the Pages tree, and has the attributes shown in the following table.

Table 6.5 Page attributes

Key	Type	Semantics
Type	name	(Required) Object type. Always Page .
MediaBox	Rectangle	(Required; may be inherited) Rectangle specifying the “natural size” of the page, for example the dimensions of an A4 sheet of paper. The coordinates are measured in default user space units. This rectangle includes any extended area surrounding the finished page for bleed, printing marks, or other similar purpose. (Also see the CropBox , BleedBox , ArtBox , and TrimBox keys.)
Parent	dictionary	(Required; must be indirect reference) The Pages object that is the immediate ancestor of this page.

Example 6.4 Inheritance of attributes



Resources	dictionary	<i>(Required; may be inherited)</i> Resources required by this page, described in Chapter 7 . If the page requires no resources, this value should be an empty dictionary, written as <code><< >></code> . Omitting this value, or specifying a null value, indicates that the value is to be inherited from an ancestor Pages object.	
Contents	stream or array	<i>(Optional; must be indirect reference)</i> The page description (contents) for this page, described in Chapter 8 . Contents may be a single stream or an array of streams. If it is an array of streams, the effect is equivalent to a single stream formed by concatenating all the elements of the array, in order. (This allows a program that is creating a PDF file to create image objects and other resources as they occur, even though they interrupt the page description.) If Contents is absent, the page is empty.	
		<i>Note</i> <i>The arrangement of streams in the array is unrelated to the content on the page. Applications that consume or produce PDF files are not required to preserve the existing structure of a Contents array.</i>	
CropBox	Rectangle	<i>(Optional; may be inherited)</i> Rectangle specifying the default clipping region for the page when displayed or printed. The default is the value of the MediaBox .	
Rotate	integer	<i>(Optional; may be inherited)</i> Specifies the number of degrees the page should be rotated clockwise when it is displayed or printed. This value must be zero (the default) or a multiple of 90.	
Thumb	stream	<i>(Optional; must be indirect reference)</i> Object that contains a thumbnail sketch of the page, described in Section 6.5 , “ Thumbnails .”	
Annots	array	<i>(Optional)</i> An array of objects, each representing an annotation on the page, described in Section 6.6 , “ Annotations .” Omit the Annots key if the page has no annotations.	
B (Beads)	array	<i>(Recommended if the page contains article beads)</i> An array whose elements are indirect references to each article bead on the page, in drawing order (the same order as the Annots array). Articles are described on page 121 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.1</div>
Dur (Duration)	real	<i>(Optional; may be inherited)</i> Specifies the “advance timing” (display duration) of a page. By default, the page does not advance automatically. See page 77 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.1</div>
Hid (Hidden)	boolean	<i>(Optional; may be inherited)</i> If <i>true</i> , the page should be hidden (not displayed) during a presentation. The default is <i>false</i> . See page 77 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.1</div>
Trans (Transition)	dictionary	<i>(Optional; may be inherited)</i> A Transition dictionary, containing information about transitions between pages. See page 77 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.1</div>
AA	dictionary	<i>(Optional; may not be inherited)</i> An additional-actions dictionary, providing defaults for the entire page. See page 108 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.2</div>
PieceInfo	dictionary	<i>(Optional)</i> A PieceInfo dictionary. PieceInfo is described on page 261 .	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.3</div>
LastModified	Date	<i>(Optional unless PieceInfo key is present)</i> Time of the most recent alteration of the content of the Page.	<div style="border: 1px solid black; padding: 2px; display: inline-block;">PDF 1.3</div>

SeparationInfo				PDF 1.3
	dictionary	<i>(Optional)</i>	If this page is a separation, this dictionary describes its output characteristics as described in the Separation Info dictionary (page 81).	
ArtBox	Rectangle	<i>(Optional; may not be inherited)</i>	Rectangle specifying an area of the page to be used when placing PDF content into another application. The default is the value of the CropBox .	PDF 1.3
TrimBox	Rectangle	<i>(Optional; may not be inherited)</i>	Rectangle specifying the intended “finished size” of the page (for example, the dimensions of an A4 sheet of paper). In some cases, the MediaBox will be a larger rectangle, which includes printing instructions, cut marks, or other content. The default is the value of the CropBox .	PDF 1.3
BleedBox	Rectangle	<i>(Optional; may not be inherited)</i>	Rectangle specifying the region to which all page content should be clipped if the page is being output in a production environment. In such environments, a “bleed area” is desired, to accommodate physical limitations of cutting, folding, and trimming equipment. The actual printed page may include printer’s marks that fall outside the bleed box. The default is the value of the CropBox .	PDF 1.3
ID	string	<i>(Optional; may not be inherited. Indirect reference preferred.)</i>	An ID that can be used to locate a Web Capture Content Set array in the IDS name tree. See page 160 .	PDF 1.3
PZ (Preferred Zoom)	number	<i>(Optional; may not be inherited)</i>	A scaling factor by which the page can be scaled to achieve the “natural” zoom level. See page 160 .	PDF 1.3

Note that some Page attributes may be inherited; see [Section 6.3.1 on page 72](#). No optional Page attributes defined in PDF 1.3 or later versions may be inherited.

Note The intersection between the page’s media box and the crop box is the region of the default user space coordinate system that is viewed or printed. Typically, the crop box is located entirely inside the media box, so that the intersection is the same as the crop box itself.

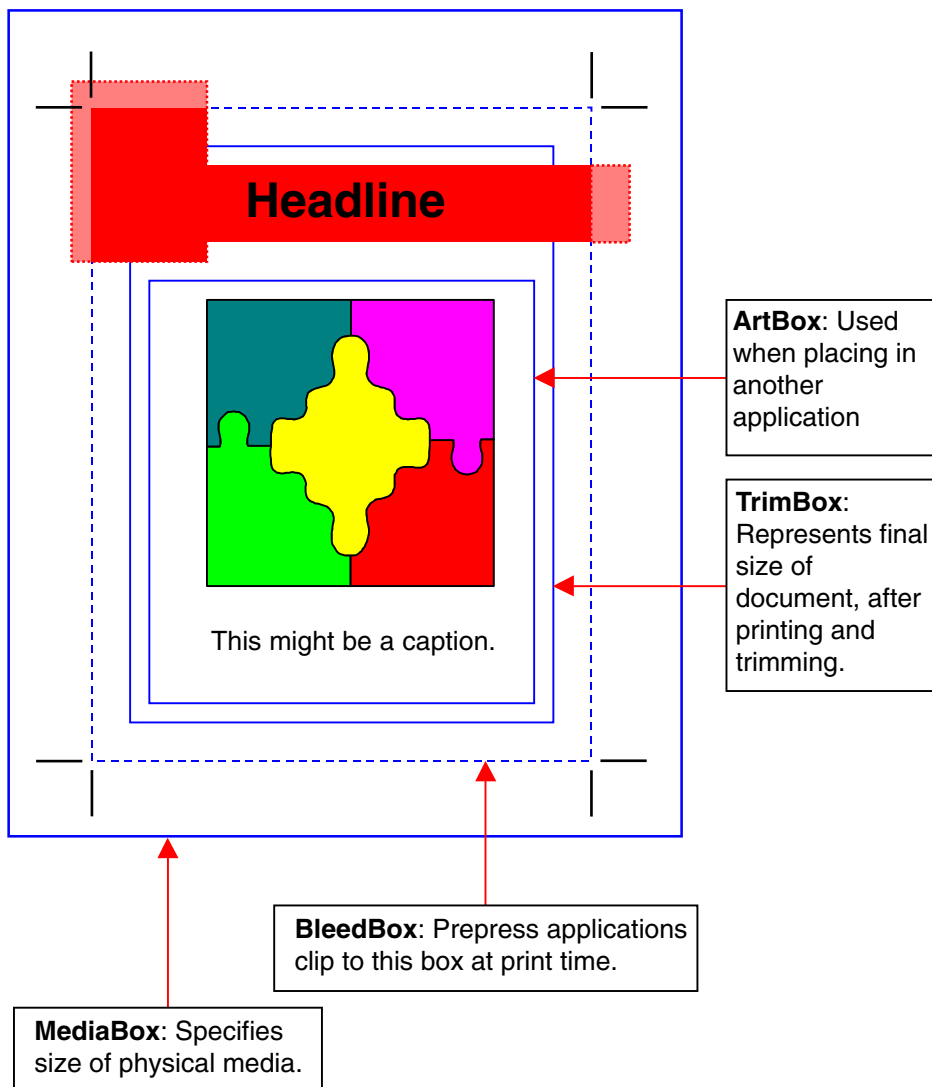
[Figure 6.2 on page 77](#) shows the distinction between the media box, art box, trim box, and bleed box. The crop box, which is not explicitly shown in the figure, is the default clipping region for viewing and printing. However, an application may use one of the other, more specific boxes for cropping, as appropriate. For example, the figure shows printer’s crop marks at the bleed box.

[Example 6.5 on page 76](#) shows a Page object with a thumbnail and two annotations. In addition, the Resources dictionary is specified as a direct object, and shows that the page makes use of three fonts, with the names **F3**, **F5**, and **F7**.

Example 6.5 *Page with thumbnail, annotations, and Resources dictionary*

```
3 0 obj
<<
  /Type /Page
  /Parent 4 0 R
  /MediaBox [0 0 612 792]
  /Resources <<
    /Font << /F3 7 0 R /F5 9 0 R /F7 11 0 R >>
    /ProcSet [/PDF] >>
    /Thumb 12 0 R
    /Contents 14 0 R
    /Annots [23 0 R 24 0 R]
  >>
endobj
```

Figure 6.2 *Page object's media box and crop box*



6.4.1 Presentation mode


PDF 1.1

A Page dictionary may contain three keys, **Dur**, **Hid**, and **Trans**, that provide information that is intended to be used when displaying a PDF document as a “presentation” or “slide show.” They are otherwise ignored. A PDF viewer is not required to provide a presentation mode. If such a mode is provided by the viewer or a plug-in, however, then these keys define its behavior.

6.4.1.1 Duration

The **Dur** key in a Page dictionary specifies the *advance timing* of the page. The advance timing is intended to be used only when a presentation is being played in a non-interactive mode. It describes the maximum amount of time the page is displayed before the viewer automatically turns to the next page; the user can advance the page manually before the time is up. If no **Dur** key is specified for a Page object or any of its Pages ancestors, the page does not advance automatically.

The advance timing is defined as the amount of time between the *end* of the last transition and the *beginning* of the next one, as shown in the time-line below:



6.4.1.2 Hidden

The **Hid** (Hidden) key in a Page dictionary specifies that the page is not to be displayed during the presentation. If the user attempts to turn to a hidden page from the previous or following page during a presentation, the page is skipped and the next visible page is displayed. If the page is the destination of a link or thread, the Hidden attribute is ignored and the page is displayed.

The Hidden attribute of a page hides the page only during a presentation; other aspects of the user interface ignore the Hidden attribute.

6.4.1.3 Transition

The **Trans** key in a Page dictionary specifies a Transition dictionary, which describes the effect to use when going *to* that page, and the amount of time the transition should take. For example, a transition effect in the Transition dictionary of page two executes whenever the user goes to page two, regardless of the previous page. The following table defines keys for all Transition dictionaries; they may contain additional keys that control specific transition effects.

Table 6.6 *Transition attributes*

Key	Type	Semantics
Type	name	Object type. Always Trans .
S (Subtype)	name	(<i>Optional</i>) Describes the transition effect. If this key is omitted, there is no transition effect to that page (the page is displayed normally), and the D key in the Transition dictionary is ignored. The allowed values are described in the following table.

D (Duration)	number	<i>(Optional)</i> The duration (in seconds) of the transition effect. The default duration is 1 second.
---------------------	--------	---

6.4.1.4 Transition effects


All implementations of presentation mode support the transition effects shown in the following table. Some of these effects include optional parameters that control the appearance of the effect. The parameters are described in [Table 6.8](#).

Table 6.7 *Transition Effects*

<i>Subtype</i>	<i>Parameters</i>	<i>Description</i>
Split	Dm, M	Two lines sweep across the screen revealing the new page image. The lines can be either horizontal or vertical, as determined by the Dm key, and can move from the center out or from the edges in as determined by the M key.
Blinds	Dm	Multiple lines, evenly distributed across the screen, appear and synchronously sweep in the same direction to reveal the new page. The lines are either horizontal or vertical, as determined by the Dm key. Horizontal lines move down; vertical lines move to the right.
Box	M	A box sweeps from the center out or from the edges inward, as determined by the M key, revealing the new page image.
Wipe	Di	A single line sweeps across the screen from one edge to the other, revealing the new page image. Possible values for Di include 0, 90, 180, and 270.
Dissolve	<i>(none)</i>	The old page image “dissolves” in a piecemeal fashion to reveal the new page.
Glitter	Di	Similar to Dissolve , except the effect sweeps across the image in a wide band moving from one side of the screen to the other. Supported directions are 0, 270, and 315.
R	<i>(none)</i>	(Replace) The effect is simply to replace the old page with the new page; i.e., there is no “transition” per se. This is the default effect if the S key is omitted from the transition dictionary, but it may be explicitly specified as a way to override any default transition that may be used for full-screen mode.

Table 6.8 *Effect parameters*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Di (Direction)	number	The direction of movement, specified in degrees, increasing in a counterclockwise direction. A value of 0 points to the right, indicating that the effect proceeds from left to right. A value of 90 points upward, indicating that the effect moves from bottom to top.



Note This is different from the page rotation, where the degrees increase in a clockwise direction.

Dm (Dimension)	name	For those effects which can be performed either horizontally or vertically, the Dm key specifies which dimension to use. Possible values are H (horizontal) or V (vertical).
M (Motion)	name	For those effects which can be performed either from the center out or the edges in, the M key specifies which direction to use. Possible values are I (In) or O (Out).

The following example shows a page that, in presentation mode, would be displayed for 5 seconds before advancing to the following page. Before the page is displayed, there is a 3-second transition in which two vertical lines sweep across the screen, from the center outwards.

Example 6.6 *A page with information for presentation mode*

```
<<
  /Type /Page
  /Parent 4 0 R
  /Contents 16 0 R
  /Dur 5
  /Trans <<
    /Type /Trans
    /S /Split
    /D 3.5
    /M /O
    /Dm /V >>
>>
```

6.4.2 Separated PDF

PDF 1.3

High-end printing workflows make use of documents in a pre-separated form, where each printed color page is represented by several gray pages, each of which is intended to be printed with a different color. When such a document is converted to PDF, the following information must be extracted from the PDF file:

- The pages, if any, to be printed as separations.
- The color needed for each separated page.
- The pages which are to be grouped together for output as a single PDF page.

In PostScript, this information is typically provided by the **%%PlateColor** and **%%Page** comments. This section describes how corresponding information is embedded in PDF files.

6.4.2.1 Separation Info Dictionary

Table 6.9 describes the attributes in a Separation Info dictionary. See Table 6.5 for more information about the page attribute dictionary keys.

Table 6.9 *Separation Attributes*

Key	Type	Semantics
DeviceColorant	name or string	<i>(Required)</i> The name of the device colorant assigned to this separation. Examples: /Cyan , (PANTONE 35 CV) .
ColorSpace	color space	<i>(Optional)</i> A color space object of type Separation or DeviceN . This provides additional information about the color specified by DeviceColorant . If the color space is of type Separation , then the DeviceColorant key must match the colorant name of the separation color space. If the color space is of type DeviceN , then the DeviceColorant key must match one of the colorant names in the color space.
Pages	array	<i>(Required)</i> An array of indirect references to page objects which are the PDF pages that comprise a single separated document page. The current page object must be one of the pages referenced by the array, and all of the pages referenced by the array must have Separation Info dictionaries, which must have Pages arrays identical to this one.

6.5 Thumbnails

A PDF document may include thumbnail sketches of its pages. They are not required, and even if some pages have them, others may not.

The thumbnail image for a page is the value of the **Thumb** key of the page object. The structure of a thumbnail is very similar to that of an Image resource (see Section 7.13.1 on page 246). The only difference between a thumbnail and an Image resource is that a thumbnail does not include **Type**, **Subtype**, and **Name** keys.

Example 6.7 *Thumbnail*

```

12 0 obj
<<
  /Filter [/ASCII85Decode /DCTDecode]
  /Width 76
  /Height 99
  /BitsPerComponent 8
  /ColorSpace /DeviceRGB
  /Length 13 0 R
>>
stream
s4IA>!"M;*Ddm8XA,1T0!!3,S!/(=R!<E3%!<N<(!WrK*!WrN,

```

... image data omitted ...

```
endstream
endobj
```

```
13 0 obj
4298
endobj
```

6.6 Annotations

Annotations are notes or other objects that are associated with a page but are separate from the page description itself. PDF supports several kinds of annotations:

- Text notes (see [page 89](#))
- Hypertext links ([page 90](#))
- Movies ([page 90](#))
- Sounds ([page 93](#))
- Widgets ([page 133](#))
- Trap networks ([page 98](#))
- Various mark-up annotations, such as rubber stamps, lines, squares, circles, strikethrough text, highlights, and underlines.

In the future, PDF may support additional types.

If a page includes annotations, they are stored in an array as the value of the **Annots** key of the Page object. Each annotation is a dictionary. As shown in [Table 6.10](#), all annotations must provide a core set of keys, including **Type**, **Subtype**, and **Rect**. Certain other keys, indicating an annotation's color, title, modification date, border, and other information, are also defined for all annotations but are optional.

Note All coordinates and measurements in text annotations, link annotations, and outline entries are specified in default user space units.

Note As an alternative to the simple border and color characteristics defined in PDF 1.0 and 1.1, Annotations in PDF 1.2 can have one or more appearances attached to them. An appearance is a PDF Form XObject that is rendered inside the annotation's bounding box.

PDF 1.2

Table 6.10 Annotation attributes (common to all annotations)

Key	Type	Semantics
Type	name	Object type. Always Annot .



Subtype	name	<i>(Required)</i> Annotation subtype.	
Rect	Rectangle	<i>(Required)</i> Rectangle specifying the location of the annotation.	
Border	array	<i>(Optional)</i> In PDF 1.0, this is an array of three numbers, specifying the horizontal corner radius, the vertical corner radius, and the width of the border of the annotation. The default values are 0, 0, and 1, respectively. No border is drawn if the width is 0.	PDF 1.0
		The array may have a fourth element, a <i>dash array</i> that allows specification of solid and dashed borders. The dash array contains “on” and “off” stroke-lengths for drawing dashes, in the same format as the d marking operator (see page 325). An example of a border with a dash array is <code>[0 0 1 [3]]</code> .	PDF 1.1
		<i>Note</i> This key has been superseded in PDF 1.2 with the BS key, described below.	
C (Color)	array	<i>(Optional)</i> The annotation color. For links, this is the border color. For text annotations, it is the background color of a closed annotation’s icon, the title bar color of an active open annotation’s window, and the window frame color of an inactive open annotation. A color is specified as an array of three numbers in the range 0 to 1, representing a color in DeviceRGB .	PDF 1.1
T (Title)	Text	<i>(Optional)</i> An arbitrary text label associated with the annotation. It is displayed in an active open text annotation’s title bar and can be edited from the annotation’s properties dialog.	PDF 1.1
M (ModDate)	Date	<i>(Optional)</i> The last time an annotation was modified. A text annotation’s modification date is updated each time the text is changed. The ModDate must be a string; the preferred format described on page 183 , but viewers should accept and display any string.	PDF 1.1
F (Flags)	integer	<i>(Optional)</i> Flags. The binary value of the integer is interpreted as a collection of flags that define various characteristics of the annotation. All undefined bits are reserved and must be set to 0. The default value for this key is 0.	PDF 1.1
		<i>Note</i> Bit-positions are numbered starting at 1, which is the least significant bit.	
	bit 1	The <i>Invisible</i> flag specifies how an annotation is displayed when the corresponding annotation handler is not available. If this flag’s value is 1, and the viewer does not provide a handler for the annotation’s Subtype, the annotation is not displayed, even if the annotation has an appearance dictionary (AP key). If this flag’s value is 0, and the viewer does not provide a handler for the annotation’s Subtype, the annotation appears as an unknown annotation unless the annotation has an appearance (AP key), in which case the appearance is displayed.	
	bit 2	The <i>Hidden</i> flag determines whether the annotation is to be shown. If the value of this flag is 1, then the annotation is hidden, in which case there is no user interaction, display, or printing of the annotation, regardless of the existence of an annotation handler. The ability to hide and show annotations selectively, combined with appearances (see the AP key), is especially useful in cases where screen real estate is limited. They can then be used to display pop-up auxiliary data similar in function to online help systems.	PDF 1.2

bit 3	The <i>Print</i> flag indicates whether the annotation should be printed. If the value of this flag is 1, then the annotation is printed, regardless of the existence of an annotation handler. Corresponding to the definition of appearances for annotations, there may be some instances where the author of a document wishes to create an annotation for display purposes only. That is, the annotation should display its appearance while in the viewer but that appearance should not print. This is typical for annotations that act like push buttons whose presence on the printed page would be distracting.	PDF 1.2
bit 4	The <i>NoZoom</i> flag indicates that if the annotation has an appearance that the appearance should not zoom as the page is magnified in the viewer.	PDF 1.3
bit 5	The <i>NoRotate</i> flag indicates that if the annotation has an appearance that the appearance should not be rotated with respect to the page before displaying the appearance. This is useful for certain annotation types (e.g., text notes) whose appearances do not rotate with the page.	PDF 1.3
bit 6	The <i>NoView</i> flag indicates that if the annotation has an appearance, then that appearance should be used only for printing. A viewer should consider this annotation hidden for display and user-interaction purposes. This bit may be useful in some instances where an author wants the annotation to print (a watermark, for example) but not to show up on the screen, as it might appear distracting.	PDF 1.3
bit 7	The <i>Read-only</i> flag indicates that the annotation should display but not interact with the user. This is equivalent to the <i>Read-only</i> flag for form fields (see page 132) but applicable to other types of annotations as well.	PDF 1.3
H (Highlight)	name <i>(Optional)</i> In PDF 1.1, the visual effect of clicking on a link annotation produced an inversion of the colors inside the bounding box of the annotation. In PDF 1.2, additional highlight modes, specified by this key, have been added for both Link and Widget annotations, as described in Table 6.11 . The default for this key is I (Invert).	PDF 1.2
BS (Border Style)	dictionary <i>(Optional)</i> This key overrides the Border key. The value of this key is a dictionary, defined on page 86 , that specifies several attributes related to the border of the annotation. This key is ignored by the Stamp , Text , StrikeOut , Highlight , Underline , Sound , FileAttachment , and Popup annotations.	PDF 1.2
AA (Additional Actions)	dictionary <i>(Optional)</i> An additional-actions dictionary. See page 108 .	PDF 1.2
AP (Appearance)	dictionary <i>(Optional)</i> Specifies that one or more appearances are available for this annotation. See page 87 .	

AS (Appearance State)

name *(Required if more than one appearance is possible)* If the **N**, **R**, or **D** keys in the **AP** dictionary are dictionaries instead of Form XObjects, then this key indicates which entry in the dictionary is to be used in each instance. This allows for the specification of items such as checkboxes and radio buttons whose appearance may change.

PDF 1.2

Contents

Text The text to be displayed, usually in a pop-up window. Text can be separated into paragraphs using carriage returns.

Popup

dictionary *(Optional)* An indirect reference to an annotation that displays the **Contents** text of this annotation in an associated pop-up window. The pop-up window is used for text entry and editing.

PDF 1.3


P

dictionary *(Optional; defined only for annotations in PDF files, not in FDF files)* An indirect reference to the **Page** object for the page on which the annotation appears.

PDF 1.3

The *NoRotate* and *NoZoom* flags indicate that the annotation’s appearance will not change with respect to the view in some manner. In these cases, it is important to define a reference point as to where the annotation will show in the view. For example, if a page is rotated and the annotation is not, about which point should the annotation’s appearance be unrotated? In all cases, the upper left corner of the annotation bounding box is considered to be this reference point. [Figure 6.3](#) shows a Text annotation that specifies the *NoRotate* flag and the result of a rotating the page 90 degrees clockwise. The upper left corner is “pinned” to the page.

Figure 6.3 Text annotation with the *NoRotate* flag



For the **Stamp, Line, Square, Circle, Text, StrikeOut, Highlight, Underline, and Ink** annotations, the *Read-only* flag is interpreted in the following manner: if the flag is on, the annotation icon may be selected but not moved, nor may any of its properties be changed. The associated **Popup** annotation, if any, may be displayed as a result of interacting with the parent annotation. This pop-up window can be moved and resized, but the text in the pop-up may not be edited.

6.6.1 Annotation borders

The border is drawn completely inside the annotation's bounding box. If neither the **BS** key nor the **Border** key is present in the annotation dictionary, the border style is drawn as a solid line with a width of one point.

Table 6.11 *Border Style attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always Border .
S (Subtype)	name	<p><i>(Optional)</i> One of the following names:</p> <p>S (Solid) The border is drawn as a solid line. This is the default.</p> <p>D (Dashed) The border is drawn with a dashed line. The dash pattern is specified by the D attribute (see below).</p> <p>B (Beveled) The border is drawn in a beveled style (faux three-dimensional) such that it looks as if it is pushed out of the page (opposite of Inset).</p> <p>I (Inset) The border is drawn in an inset style (faux three-dimensional) such that it looks as if it is inset into the page (opposite of Beveled).</p> <p>U (Underlined) The border is drawn as a line on the bottom of the annotation rectangle.</p> <p>Other subtypes may be defined in the future.</p>
W (Width)	number	<i>(Optional)</i> The border width in points. The default is 1 point.
D (Dash array)	array	<i>(Optional)</i> If the Subtype is D , this array contains numbers representing on and off stroke lengths for drawing dashes, in the same format as the d marking operator. The default for this key is [3] .

6.6.2 Annotation highlighting

An annotation is highlighted when any of the following events is triggered.

PDF 1.2



PDF 1.2



1. If the user clicks the mouse button down inside the active area of the annotation, the highlight is shown.
2. If the mouse button is down and the user moves the cursor out of the active area of the annotation, the highlight is removed.
3. If the mouse button is down and the user moves the cursor back into the active area of the annotation, the highlight is shown.
4. If the mouse button is released then the highlight is removed.

A highlight is *removed* by reverting the visual appearance of the document back to what it was before the highlight occurred.

Table 6.12 *Highlight Modes*

<i>Name</i>	<i>Semantics</i>
I (Invert)	The rectangle specified by the bounding box of the annotation is inverted.
N (None)	No highlighting is done.
O (Outline)	The border of the annotation is inverted. If no active color is defined, the border is inverted.
P (Push)	If the annotation specifies a Down appearance (i.e., if the Appearance dictionary defines a D key), that appearance is drawn. If the annotation does not specify a Down appearance, then the region underneath the bounding box of the annotation is drawn inset into the page.

The use of the highlighting works in conjunction with the Down sub-appearance specified in the annotation appearance dictionary (if any) in the following manner: If the highlight value is not **P** (push), then the specified highlight mode is respected regardless of the fact that a down sub-appearance exists.

Note *The active behavior of highlighting is applicable only to annotations whose annotation handler is present. The highlighting behavior is logically the responsibility of the handler. Plug-in designers are encouraged to use the same PDF extensions for specifying highlighting for their annotations.*

6.6.3 Annotation appearances

The appearance dictionary defines the following three keys.

PDF 1.2



Table 6.13 *Appearance dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
N (Normal)	dictionary or stream	<i>(Required)</i> This is the appearance that is used when the annotation is in a normal state, that is, when the user is not interacting with the annotation. This is also the appearance that is used when printing the annotation.
R (Rollover)	dictionary or stream	<i>(Optional)</i> appearance. This is the appearance that is used when the user moves the mouse cursor inside the annotation. If the cursor leaves the annotation then the Normal appearance is re-displayed. The default is that no alternative appearance is displayed when the rollover occurs.
D (Down)	dictionary or stream	<i>(Optional)</i> This is the appearance that is used when the user clicks the mouse down inside the annotation. When the user releases the mouse, the viewer displays the Rollover appearance if that is defined, or the Normal appearance otherwise. The default is that no alternative appearance is displayed when the mouse-down event occurs.

If the **N**, **R**, or **D** values are dictionaries, then they contain one or more sub-appearances. The appearance state (**AS**) key in the annotation dictionary determines which sub-appearance to use as the current appearance.

Consider a checkbox with two states: checked and unchecked. The appearance dictionary, containing only a normal appearance (**N**), would look like this:

```

/AP <<
  /N <<
    /Checked FormXObject
    /Unchecked FormXObject
  >>
>>

```

The appearance state key would determine which appearance is the current normal appearance. For example, if the checkbox were currently checked, the value of the **AS** key would be the name **Checked**.

```

/AS /Checked

```

If the **N**, **R**, or **D** values are not dictionaries, then they must be Form XObjects (see [page 255](#)), which are represented as streams, and the appearance does not depend on the **AS** key.

The Form XObject's graphics state when rendering the appearance should have the origin translated to the lower left corner of the annotation bounding box and scaled to fit inside and clipped to the annotation bounding box. All other graphics state values are set to their defaults.

Some annotation subtypes defined in PDF (e.g., **Movie**) and all annotation subtypes defined by third parties are implemented through plug-ins. If the plug-in that implements a particular annotation subtype is not available, the viewer displays that annotation's Normal appearance.

Note The viewer should attempt to provide reasonable behavior (e.g., display nothing) if the sub-appearance named by the **AS** key is not present in the **N**, **R**, or **D** dictionaries.

6.6.4 Text annotations

A text annotation contains a **Text** in its **Contents** key, which is required. When the annotation is open, the text is displayed. A PDF viewer chooses the size and typeface of the text. The following table shows the attributes of the Text annotation dictionary.

Table 6.14 Text annotation attributes (in addition to those in Table 6.10)

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always Text .
Open	boolean	(Optional) If <i>true</i> , specifies that the annotation should initially be displayed opened. The default is <i>false</i> (closed).
Name	name	(Optional) The name of the icon to be displayed for the annotation. This is used by the viewer to construct an appearance if the AP key is omitted. The default value is Note .
AP	dictionary	(Optional) An indirect reference to an appearance dictionary. If this key is omitted, an appearance will be constructed using the Name key.

PDF 1.3

PDF 1.3

Implementation note The Acrobat viewer will use a predefined appearance for the following values of **Name**: **Comment**, **Insert**, **Note**, **Paragraph**, **NewParagraph**, **Key**, and **Help**.

Example 6.8 Text annotation

```

22 0 obj
<<
  /Type /Annot
  /Subtype /Text
  /Rect [266 116 430 204]
  /Contents (text for two)
>>
endobj

```

6.6.5 Link annotations

A link annotation, when activated, displays a *destination* or performs an *action*. A destination is a view of another location, possibly on a different page, with a different zoom factor, or in a different file. The following table shows the contents of the link annotation dictionary.



Table 6.15 *Link annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always Link .
Dest	array or name	(Optional; not permitted if the A key is present) The view to go to, represented as a destination. See page 184 .
A (Action)	dictionary	(Optional; not permitted if the Dest key is present) The action to be performed on activating this link annotation. See page 107 .
PA (Previous Action)	dictionary	(Optional) As a part of its normal operation, Web Capture (see Section 6.16) will change the actions associated with some link annotations from URI actions (see Section 6.8.6) to GoTo actions (see Section 6.8.2). For such an annotation, it desirable that the data for the original URI action be available even after the action has been changed to a GoTo action. This allows Web Capture to change the action for the link annotation back to the original URI action if the target page for the GoTo action is deleted. The PA key contains the former action to perform upon activating this link.

PDF 1.1

PDF 1.3

If neither the **A** or **Dest** key is present, the **AA** key can be used to describe the annotation's behavior.

Example 6.9 *Link annotation*

```

93 0 obj
<<
  /Type /Annot
  /Subtype /Link
  /Rect [71 717 190 734]
  /Border [16 16 1]
  /Dest [3 0 R /FitR -4 399 199 533]
>>
endobj

```

6.6.6 Movie Player annotations

A Movie Player annotation describes the static display and playing of movies and sounds within PDF documents. These annotations appear to be embedded in the document, like links. The activation area may be invisible, bordered in the manner

PDF 1.2



of a link button, and it may have the movie’s Poster frame displayed. There are several authoring options that control the way a movie is to be displayed and played.


Table 6.16 *Movie Player annotation attributes (in addition to those in Table 6.10)*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	name	<i>(Required)</i> Annotation subtype. Always Movie .
Movie	dictionary	<i>(Required)</i> A description of the static characteristics of the Movie; see Table 6.17 .
A (Activation)	boolean	<p><i>(Optional)</i> A flag that indicates whether the movie should be shown by clicking in the annotation rectangle. Possible values are:</p> <p>false Do not play the movie when clicked.</p> <p>true Play the movie with the default activation values. (This is the default value for the A key.)</p> <p>or</p>
	dictionary	<i>(Optional)</i> Directions for playing the Movie; see Table 6.17 .

The Movie dictionary contains information needed to locate the movie data and to display the poster (if requested) in the annotation rectangle.

Table 6.17 *Movie dictionary attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
F (File) File specification		<p><i>(Required)</i> A self-describing movie file.</p> <p><i>Note</i> The format of a “self-describing movie file” is left unspecified, and there is no guarantee of portability.</p>
Aspect	array	<i>(Optional)</i> If the movie is visible, the horizontal and vertical sizes of the movie’s bounding box in pixels: [horiz vert] . An “invisible movie” is one with no video: it has only sound.
Rotate	integer	<p><i>(Optional)</i> Specifies the number of degrees clockwise the movie must be rotated, relative to the rotation of the page. This value must be a multiple of 90. The default is 0.</p>
Poster	boolean	<p><i>(Optional)</i> A flag indicating whether the poster is to be retrieved from the movie file for display. Possible values are:</p> <p>false Do not show a poster image. (This is the default if the Poster key is omitted.)</p> <p>true Show the poster image from the movie file.</p> <p>or</p>
	stream	<i>(Optional)</i> An Image resource (see Section 7.13.1 on page 246) that is to be displayed as the poster.



The Movie Activation dictionary contains information needed to control the dynamics of playing the movie.

Table 6.18 *Movie Activation attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
ShowControls	boolean	<i>(Optional)</i> If this key is <i>true</i> , a Movie Controller bar is shown when the movie is being played. The default is <i>false</i> .
Mode	name	<p><i>(Optional)</i> The playing mode for the movie. Currently defined values are:</p> <p>Once Show the movie once and stop. Open Show the movie and leave the controller open. Repeat Repeat the movie from the beginning until stopped. Palindrome Play the movie back and forth until stopped.</p> <p>If Mode is omitted or unrecognized, the mode is Once.</p>
Synchronous	boolean	<i>(Optional)</i> If this key is <i>true</i> , the player does not return to Acrobat until the movie is completed or dismissed by the user. The default is <i>false</i> .
Start	number, string, or array	<p><i>(Optional)</i> The starting time of the movie segment to play. If omitted, the movie is played from the beginning.</p> <p>Movie time values are based on 64-bit integers. If the start time is representable in 32 bits, the key value should be an integer. If not, the key value should be an 8-byte string, with the most significant byte of the value first, that is treated as a 64-bit integer. If the time scale for the starting time is not the same as the Movie time scale, the start time is represented as an array of two values: the first element is the time value (integer or string), and the second element is the time scale (integer). The time scale is measured in events per second.</p>
Duration	number, string, or array	<p><i>(Optional)</i> The duration of the movie segment to play. If omitted, the movie is played to the end.</p> <p>Movie time values are based on 64-bit integers. If the duration is representable in 32 bits, the key value should be an integer. If not, the key value should be an 8-byte string, with the most significant byte of the value first, that is treated as a 64-bit integer. If the time scale for the duration is not the same as the Movie time scale, the duration is represented as an array of two values: the first element is the time value (integer or string), and the second element is the time scale (integer). The time scale is measured in events per second. A negative duration means the movie is to be played backwards.</p>
Rate	number	<i>(Optional)</i> The initial speed at which the movie is played. The default speed is 1.0. A negative speed means the movie is to be played backwards with respect to Start and Duration .

Volume	number	<i>(Optional)</i> The initial volume setting for the movie. This number must be between -1 and 1. Negative settings are muted. A higher number denotes increased volume. The default volume is 1.
FWScale	array	<i>(Optional)</i> For floating play windows, the magnification at which to play the movie. FWScale is an array of two integers: <code>[numerator denominator]</code> representing a rational magnification factor for the movie. The final window size for the movie is $(\text{numerator} \div \text{denominator}) \times \text{Aspect pixels}$. The presence of the FWScale key implies that the movie is to be played in a floating window. The absence of the FWScale key implies that the movie is to be played in the annotation rectangle.
FWPosition	array	<i>(Optional)</i> For floating play windows, the position within the screen at which to play the movie. FWPosition is an array of two numbers: <code>[horizontal vertical]</code> representing a relative position of the movie window with respect to the left and top of the screen. Each number must be in the range 0 to 1, with <code>[0.5 0.5]</code> meaning “center the movie window on the screen.”

6.6.7 Sound annotations

PDF 1.2

A Sound annotation is analogous to a Text annotation, except that it contains sound, recorded from the computer’s microphone or imported from a file. The annotation behaves like a Text annotation in most ways, with a different icon (by default, a speaker) to indicate that it is a Sound annotation. The **Contents** key, which is optional for this annotation, specifies the text to be displayed as the description of the sound.

Table 6.19 *Sound annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	<i>(Required)</i> Annotation subtype. Always Sound .
Sound	stream	<i>(Required)</i> A Sound object (see Section 6.15 on page 147).
Name	name	<i>(Optional)</i> The name of the icon to be displayed for the annotation. This is used by the viewer to construct an appearance if the AP key is omitted. The default value is Speaker .
AP	dictionary	<i>(Optional)</i> An indirect reference to an appearance dictionary. If this key is omitted, an appearance will be constructed using the Name key.

PDF 1.3

PDF 1.3

Implementation note Acrobat uses a predefined appearance when the value of **Name** is **Speaker** or **Microphone**.

6.6.8 Free Text annotations

PDF 1.3

The Free Text annotation displays text on top of the associated page using the specified appearance attributes. The text is specified by the **Contents** key, which is required for this annotation. There is no pop-up window for this annotation, since the text is always visible.

Table 6.20 Free text annotation attributes (in addition to those in Table 6.10)

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always FreeText .
DA	string	(Required) The default appearance string (see Table 6.60).
AP	dictionary	(Optional) An indirect reference to an appearance dictionary (see Section 6.6.3 on page 87). If this key is omitted, then the appearance is generated based on values of the Contents and DA keys.

6.6.9 Rubber Stamp annotations

PDF 1.3

The Rubber Stamp annotation displays an appearance as specified by its **AP** key. The associated **Contents** text is displayed in a pop-up window when the annotation is opened.

Table 6.21 Rubber Stamp annotation attributes (in addition to those in Table 6.10)

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always Stamp .
Name	name	(Optional) The name of the rubber stamp appearance to be displayed for the annotation. This is used by the viewer to construct an appearance if the AP key is omitted. The default value is Draft .
AP	dictionary	(Optional) An indirect reference to an appearance dictionary. If this key is omitted, an appearance will be constructed using the Name key.

Implementation note Acrobat uses a predefined appearance if **Name** is any one of the following: **Approved, AsIs, Confidential, Departmental, Draft, Experimental, Expired, Final, ForComment, ForPublicRelease, NotApproved, NotForPublicRelease, Sold, and TopSecret.**

6.6.10 Line annotation

PDF 1.3

A Line annotation displays an appearance as specified by its **AP** key which is in the form of a single line. The associated **Contents** text is displayed in a pop-up window when the annotation is opened.

Table 6.22 *Line annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	<i>(Required)</i> Annotation subtype. Always Line .
BS (Border Style)	dictionary	<i>(Optional)</i> The border style (see Table 6.11 on page 86) specifies the characteristics of the line: its width and its dash pattern.
L	array	<i>(Required)</i> An array of four numbers, [x_1 y_1 x_2 y_2] , specifying the starting and ending points of a line, in default user coordinates.
AP	dictionary	<i>(Optional)</i> An indirect reference to an appearance dictionary. If this key is omitted, then an appearance is generated from the BS and L keys.

6.6.11 Square and Circle annotations

PDF 1.3

The Square and Circle annotations display an appearance as specified by their **AP** keys which are in the form of either a rectangle or an ellipse respectively (despite the names Square and Circle). The associated text (in the **Contents** key) is displayed in a pop-up window when the annotation is opened. The square and circle are drawn inside the bounding box specified by the **Rect** key, which is described in [Table 6.10](#).


Table 6.23 *Square and Circle annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	<i>(Required)</i> Annotation subtype. Always Square or Circle .
BS (Border Style)	dictionary	<i>(Optional)</i> The border style (see Table 6.11 on page 86) specifies the characteristics of the line: its width and its dash pattern.
AP	dictionary	<i>(Optional)</i> An indirect reference to an appearance dictionary. If this key is omitted, then an appearance is generated from the BS and Rect keys.


6.6.12 StrikeOut, Highlight, and Underline annotations

PDF 1.3

The StrikeOut, Highlight, and Underline annotations display an appearance as specified by their **AP** keys which is in the form of a one or more disjoint horizontal line segments. The associated text (**Contents** key) is displayed in a pop-up window when the annotation is opened.

Figure 6.4 *Square and Circle annotations*


Square annotation with a border size of 18 points, shown inscribed inside the bounding box of the annotation (represented by the dashed line).



Circle annotation with a border size of 18 points, shown inscribed inside the bounding box of the annotation (represented by the dashed line).

Table 6.24 *StrikeOut, Highlight, and Underline annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always StrikeOut , Highlight , or Underline .
QuadPoints	array	<p>(Required) An array of $8 \times n$ numbers defining the coordinates of n quadrilaterals, in default user space. Each 8 numbers are in the following order:</p> $x_1 \ y_1 \ x_2 \ y_2 \ x_3 \ y_3 \ x_4 \ y_4$ <p>specifying four points of a convex quadrilateral, proceeding counterclockwise from (x_1, y_1). Each quadrilateral encompasses a word or group of contiguous words in the text underlying the annotation. The text is oriented with respect to the edge from (x_1, y_1) to (x_2, y_2).</p>
AP	dictionary	(Optional) An indirect reference to an appearance dictionary. If this key is omitted, then an appearance is generated from the QuadPoints key.

*Implementation note*

In Acrobat 4.0, the text is oriented with respect to the vertex with the minimum y value (or the leftmost of those, if there are two such points) and the next vertex in a counterclockwise direction, regardless of whether those are the first two points in the **QuadPoints** array.

6.6.13 Ink annotations**PDF 1.3**

The Ink annotation displays an appearance as specified by its **AP** key which is in the form of a one or more disjoint paths. The associated text (**Contents** key) is displayed in a pop-up window when the annotation is opened.

Table 6.25 *Ink annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always Ink .
InkList	array	(Required) An array of n arrays, representing n stroked paths. Each path is represented by an array of alternating x and y coordinates, specifying points along the path, connected by straight lines or curves, forming a continuously drawn gesture. All coordinates are in default user space.
BS (Border Style)	dictionary	(Optional) The border style (see Table 6.11 on page 86) specifies the width and dash patterns of the paths in the InkList .
AP	dictionary	(Optional) An indirect reference to an appearance dictionary. If this key is omitted, then an appearance is generated from the BS and InkList keys.

Implementation note *The Acrobat 4.0 viewer always connects the points along each path with straight lines.*

6.6.14 File Attachment annotations

PDF 1.3

A File Attachment annotation allows the user to embed a file within a PDF file for later extraction. For example, a user may want to embed a spreadsheet file next to the table that shows the data in the spreadsheet.

Table 6.26 *File attachment annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always FileAttachment .
FS File specification		(Required) A file specification that refers to an embedded file (not the embedded file itself).
Name	name	(Optional) The name of the icon to be displayed for the annotation. This is used by the viewer to construct an appearance if the AP key is omitted. The default value is PushPin .
AP	dictionary	(Optional) An indirect reference to an appearance dictionary. If this key is omitted, an appearance will be constructed using the Name key.

Implementation note *Acrobat uses a predefined appearance when the value of **Name** is **Graph**, **Paperclip**, **PushPin**, or **Tag**.*

6.6.15 Pop-up annotations

PDF 1.3


Pop-up annotations display text in a pop-up window. The window can be used for entry and editing of the text. Typically, the pop-up annotation is associated with another “parent” annotation (e.g., Text, Ink) and does not appear alone in the document.

A Pop-up annotation has no appearance or associated actions.

Table 6.27 *Pop-up annotation attributes (in addition to those in Table 6.10)*

Key	Type	Semantics
Subtype	name	(Required) Annotation subtype. Always Popup .
Parent	dictionary	(Optional) The parent annotation for this pop-up. The parent uses the pop-up annotation to display its Contents text for viewing and editing. If this key is specified, then the parent’s C (color), T (title), Contents , and M (modification date) keys override those specified in the Pop-up dictionary.
Open	boolean	(Optional) Specifies whether the annotation should be opened (displayed) initially. The default value is <i>false</i> .

Figure 6.5 *Pop-up annotation for a Text annotation*



6.6.16 PDF Trapping

PDF 1.3

Beginning with PDF 1.3, a PDF file may contain a *trap network*.

This section does not specify how trapping instructions are generated for PDF files. That is explained in the Portable Job Ticket Format (PJTF) Specification; see the Bibliography [9].

6.6.16.1 Discussion of Trapping

Trapping refers to the process of adding marks to pages along color borders in order to avoid unattractive visual artifacts when those pages are printed. These marks are calculated by software that interprets the page description language (in this case, PDF) and determines where the marks should be added.

There are two aspects to trapping: trapping instructions, and trap networks.

Trapping instructions determine how the trapping engine should calculate traps for a page. Specifically, trapping instructions identify zones on the page which should be trapped (the *trap zones*), and prescribe the set of parameters (the *trap styles*) which the trapping engine should use when calculating traps for each zone.

Trapping instructions are encoded in **TrappingParameters** and **TrapRegion** objects as defined in the Portable Job Ticket Format specification (version 1.1).

A *trap network annotation* is the collection of all trap calculations for a single page. A single trap style may be used on each region, or different trap styles may be used on different regions of the same page. However, each trap region may have only one trap style specified for it at any given time.


6.6.16.2 Terminology

The following terms are used in this section:

<i>Trap Zone</i>	A closed region on a page which is to be trapped. A trap zone is trapped using exactly one trap style (see below).
<i>Trap Style</i>	A set of parameters which define how traps should be calculated. Trap styles are equivalent to the contents of the trapping parameters dictionary described in section 6.3.3 of the <i>PostScript Language Reference Manual, Third Edition</i> [1].
<i>Trapping Parameters</i>	A PJTF object which specifies a Trap Style.
<i>Trap Region</i>	A PJTF object [9] which defines a trap zone, and which identifies the Trapping Parameters object which is to be used to trap the zone.
<i>Trap Network</i>	A collection of PDF marking operations which comprise the traps for one or more Trap Regions on a page.
<i>Current Trap Network</i>	The trap network that the annotation is currently displaying or printing. (A trap network annotation is a PDF annotation of type TrapNet as defined in Section 6.6.16.4 on page 100.)

Trap Engine

Software which calculates trap networks for pages.

Figure 6.6 *Trap Network Annotation***1 Trap Network Annotation = 1 PDF Document Page****6.6.16.3 Trap Network Definition**

This section specifies an annotation type to record, display, and print trap networks. Each page may have a single TrapNetwork annotation. The annotation can include multiple trap networks for a page. Each trap network is represented as a separate entry in the Appearance Dictionary for the TrapNetwork annotation.

6.6.16.4 TrapNetwork Annotation Attributes

Each page may have at most one annotation of type **TrapNet**.

The **TrapNet** annotation must occur as the last entry in the **Annots** array for the page. This ensures that the trap network is printed *after* all other page contents.



Trap networks are represented as Form XObjects that occur as entries in the Appearance Dictionary of **TrapNet** annotations. Trap network appearances which are to be printed must occur in the **N** (Normal) sub-dictionary of the annotation **AP** dictionary. Entries may be provided for the **R** (Rollover) or **D** (Down) keys in the **AP** dictionary, but they cannot be printed.

More than one trap network may be calculated for a page, based on different trapping styles (presuming different output devices). The sub-keys used inside the **N** (Normal), **R** (Rollover) and **D** (Down) dictionaries are arbitrary.

Additional keys are provided in the trap network appearance objects which allow the identification of the trap styles and zones which were used to calculate them. See [TrapNetwork Appearances Attributes on page 102](#) for more details.

The **Version** key provides a mechanism to identify the elements which made up the page description when the trap network was calculated. This is accomplished by retaining a reference to page contents and resources which might be changed by an editing application, and would thus modify the appearance of the page, invalidating the trap network. This array is unordered; to validate the trap network, an application must enumerate the page elements which comprise the array and verify that the results of that enumeration exactly match the elements of the **Version** array.

Only one **Version** key is provided per **TrapNet** annotation (and thus per page). For this reason, an application which calculates trap networks must verify the validity of existing trap networks when calculating new trap networks. Trap networks that are valid may be retained. Invalid trap networks must be deleted, then replaced.

The **FontFauxing** key is used to identify **Fonts** which were fauxed during the generation of traps for the page.

Table 6.28 *TrapNetwork Annotation Attributes (in addition to those in Table 6.10)*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	name	(Required) Annotation subtype. Always TrapNet
F (Flags)	integer	(Required) The flag field must be present, with the bit values specified below: <ul style="list-style-type: none"> bit 1 The <i>Invisible</i> flag must be 0. bit 2 The <i>Hidden</i> flag must be 0. bit 3 The <i>Print</i> flag must be 1. bit 4 The <i>NoZoom</i> flag must be 0. bit 5 The <i>NoRotate</i> flag must be 0. bit 6 The <i>NoView</i> flag must be 0. bit 7 The <i>ReadOnly</i> flag must be 1.

AP (Appearance)	dictionary	<i>(Required)</i> Holds one or more trap networks for this page.
AS (Appearance State)	name	<i>(Required if more than one appearance is present)</i> This value specifies which trap network is “current” and will be displayed and printed.
Version	array	<i>(Required)</i> An unordered array of the objects that were present in the page description when the trap network was calculated. For example, this array includes all stream objects which occur in the Page key. The entries in this array represent all the objects which (if changed) could affect the appearance of the page. A complete list of the objects which must be included in this array is provided in Section 6.6.16.6 on page 103 .
AnnotStates	array	<i>(Required)</i> An ordered array of names. Each entry is the value of the AS key for the corresponding annotation in the Annots array of the page. For an annotation with no AS key, null should be specified as the array element. No entry should be included for the TrapNet annotation.
FontFauxing	array	<i>(Optional)</i> An array of fonts which were “fauxed” (replaced by substitution fonts) during trapping. The entries in this array are the Font Resource objects which were fauxed.

6.6.16.5 TrapNetwork Appearances Attributes

Each trap network appearance stores the marking operations calculated by the trapping engine.

When the trapping instructions and the geometry of the trap zones are specified by objects which occur in a Portable Job Ticket which is embedded in the PDF file, the trap network appearances may contain references to the **PJTF TrapRegion** objects which were used to by the engine to calculate those traps.

Note that **TrapRegion** objects (as defined in the Portable Job Ticket Format) provide the geometry of the trap region, as well as reference (by name) to the **TrappingParameters** used.

When the trapping instructions and the trap zone geometry are specified by an external job ticket, or by some other means, the **TrapRegions** key is not present.

Whether or not **TrapRegions** is present, a textual description of the trapping styles used to calculate the trap network may be provided using the **TrapStyles** key.

Table 6.29 *TrapNetwork Appearance attributes (in addition to those in Table 6.28)*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
TrapRegions	array	(Optional) An array of references to TrapRegion objects (as specified in the Portable Job Ticket Specification).
TrapStyles	Text	(Optional) Text which an application can use to describe a trap network to users (for example, if the application provides the ability to switch between trap networks).
PCM	name	(Required) Identifies the Process Color Model which was assumed when this network was created (as documented in the <i>PostScript Language Reference Manual, Third Edition</i> [1]). Legal values are DeviceGray , DeviceRGB , DeviceCMYK , DeviceCMY , DeviceRGBK , and DeviceN .
SeparationColorNames	array	(Optional) Identifies the colorants which were assumed when this network was created. If absent, the colorants implied in the ProcessColorModel are assumed. Colorants implied by PCM are always included.

6.6.16.6 Version Array

The following objects must be included in the **Version** array:

Table 6.30 *Page elements to include in the Version array*

<i>Element</i>	<i>Subtype</i>	<i>Resources to include</i>
Contents	--	Include the Contents stream, or all streams in the Contents array.
	Allothers	None.

In addition, the following resources must be included:

Table 6.31 *Page Resources to include in the Version array stream, or all stream*

<i>Type</i>	<i>Subtype</i>	<i>Resources to include</i>
XObject	all	Include all XObject streams and all OPI dictionaries contained in the XObject.
ExtGState		Include all ExtGState dictionaries.
Font	any	Include all Font dictionaries.
Pattern		Include all Pattern stream objects.
ProcSet		None.

Properties		Include each Property List (see Section 7.19 , “Property lists”).
Shading		Include each Shading dictionary.
ColorSpace	any	Include all ColorSpace arrays.

In addition, for Form XObjects, include the following entries from the XObject’s Resources dictionary:

Table 6.32 *XObject Resources to include in the Version array*

<i>Type</i>	<i>Subtype</i>	<i>Resources to include</i>
ExtGState		Include all ExtGState dictionaries.
Font	any	Include all Font dictionaries.
Pattern		Include all Pattern stream objects.
ProcSet		None.
Properties		Include each Property List dictionary.
Shading		Include each Shading dictionary.
ColorSpace	any	Include all ColorSpace arrays.

6.6.16.7 Separated PDF and trapping

Separated PDF files (see [page 80](#)) cannot be trapped. This is because traps are defined along the borders of adjacent objects that have different colors, compensating for the small gaps that may occur when printing. A single separated file uses only one color. Logically, separation must occur after trapping.

If a composite trapped PDF file is separated, it is the responsibility of the application which separates the file to calculate new **Version** keys for the separated trap networks.

6.7 Outline tree

An outline allows a user to access views of a document by name. As with a link annotation, activation of an outline entry (also called a *bookmark*) brings up a new view based on the destination description. Outline entries form a hierarchy of elements. An entry may be one of several at the same level in the outline, it may be a sub-entry of another entry, and it may have its own set of child entries. An outline entry may be *open* or *closed*. If it is open, its immediate children are visible when the outline is displayed. If it is closed, they are not.

If a document includes an outline, it is accessed from the **Outlines** key in the Catalog object. The value of this key is the Outlines object, which is the root of the outline tree. The contents of the Outlines dictionary appear in [Table 6.33](#) and [Example 6.10](#). The top-level outline entries are contained in a linked list, with **First** pointing to the head of the list and **Last** pointing to the tail of the list. When displayed, outline entries appear in the order in which they occur in the linked list.

Table 6.33 *Outlines attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Count	integer	<i>(Required if document has any open outline entries, otherwise optional)</i> Total number of open entries in the outline. This includes the total number of items open at all outline levels, not just top-level outline entries. If the count is zero, this key should be omitted.
First	dictionary	<i>(Required if document has any outline entries; must be indirect reference)</i> Reference to the outline entry that is the head of the linked list of top-level outline entries.
Last	dictionary	<i>(Required if document has any outline entries; must be indirect reference)</i> Reference to the outline entry that is the tail of the linked list of top-level outline entries.

Example 6.10 *Outlines object with six open entries*

```

21 0 obj
<<
  /Count 6
  /First 22 0 R
  /Last 29 0 R
>>
endobj

```

Each outline entry is a dictionary, whose contents are shown in the following table.

Table 6.34 *Outline entry attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Title	Text	<i>(Required)</i> The text that appears in the outline for this entry.
Dest	Destination	<i>(Optional; not permitted if the A key is present)</i> A destination, as described on Section 7.3.1 on page 184 .
A (Action)	dictionary	<i>(Optional, not permitted if the Dest key is present)</i> The action to be performed when this bookmark is activated; see page 107 .
Parent	dictionary	<i>(Required; must be indirect reference)</i> Specifies the entry for which the current entry is a sub-entry. The parent of the top-level entries is the Outlines object.

PDF 1.1

Prev	dictionary	<i>(Required if the entry is not the first of several entries at the same outline level; must be indirect reference)</i> Specifies the previous entry in the linked list of outline entries at this level.	
Next	dictionary	<i>(Required if the entry is not the last of several entries at the same outline level; must be indirect reference)</i> Specifies the next entry in the linked list of outline entries at this level.	
First	dictionary	<i>(Required if an entry has sub-entries; must be indirect reference)</i> Specifies the outline entry that is the head of the linked list of sub-entries of this outline item.	
Last	dictionary	<i>(Required if an entry has sub-entries; must be indirect reference)</i> Specifies the outline entry that is the tail of the linked list of sub-entries of this outline item.	
Count	integer	<i>(Required if an entry has sub-entries)</i> If positive, Count specifies the number of open descendants the entry has. This includes not just immediate sub-entries, but sub-entries of those entries, and so on. If the value is negative, the entry is closed, and the absolute value of Count specifies how many entries would appear if the entry were reopened. If an entry has no descendants, the Count key is omitted.	
SE	dictionary	<i>(Optional; must be indirect reference)</i> The structural element to which the outline entry refers. See Section 6.17.3 on page 164 .	PDF 1.3

Note In PDF 1.2, the **AA** (Additional Actions) key was also defined as an Outline entry attribute. It was never implemented. In PDF 1.3, the **AA** key is obsolete and should be ignored. **PDF 1.3**

In PDF 1.3, an outline entry may refer to a structural element (e.g., the start of Chapter 14) as well as to a destination (e.g., a certain rectangle on page 432) or action. For backwards compatibility, however, such outline items should also specify a destination as in PDF 1.2 and earlier. If a structural element is supplied, the destination should correspond to a page area where the content of the structural element is displayed. **PDF 1.3**

The following example shows an outline entry. An example of a complete outline tree can be found in [Appendix A](#).

Example 6.11 *Outline entry*

```

22 0 obj <<
  /Parent 21 0 R
  /Dest [3 0 R /XYZ 0 792 0]
  /Title (Document)
  /Next 29 0 R
  /First 25 0 R
  /Last 28 0 R
  /Count 4
>>
endobj

```

6.8 Actions

PDF 1.1



Beginning in PDF 1.1, it is possible to specify an *action* to be performed when a Link annotation or Outline entry is activated, or when a document is opened. PDF defines several subtypes of actions.

Table 6.35 *Types of actions*

<i>Action type</i>	<i>Description</i>
GoTo	Change the current page view to a specified page and zoom factor. See page 111 .
GoToR	(“GoTo Remote”) Open another PDF file at a specified page and zoom factor. See page 111
Launch	Launch an application, usually to open a file. See page 112 .
Thread	Begin reading an article thread. See page 121 .
URI	Resolve the specified Uniform Resource Identifier (URI). See page 113 .
Sound	Play a sound. See page 115 .
Movie	Play a movie. See page 116
SetState	Store a value in the appearance-state (AS key) of an annotation’s appearance dictionary. See page 117 .
Hide	Set or clear the Hidden flag for an annotation. See page 117 .
Named	Execute an action predefined by the viewer. See page 118 .
SubmitForm	Send data to a URL. See page 143 .
ResetForm	Set field values to their defaults. See page 144 .
ImportData	Import field values from a file. See page 145 .
JavaScript	Execute a JavaScript script. See page 145

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.2

PDF 1.3

Note It is intended that plug-in extensions may add new actions, as described in [Appendix G](#).



Note The **NOP** (*no-op*) action, defined in PDF 1.2, is obsolete in PDF 1.3.

An action is represented as a dictionary. Every action must contain an **S** (Subtype) key. Other keys may be present, depending on the action type.

Table 6.36 *Action attributes (common to all actions)*

Key	Type	Semantics
Type	name	Object type. Always Action .
S (Subtype)	name	<i>(Required)</i> Action subtype.
Next	dictionary or array	<i>(Optional)</i> The action, or sequence of actions, to be performed after the current action.

PDF 1.2

In PDF 1.1, only single actions can be triggered as the result of a user action. PDF 1.2 allows actions to be chained together to provide a richer action model. For example, if a user clicks on a particular link, it might play a sound, then go to a new page, and then start up a movie.

PDF 1.2

The value of the **Next** key is either an action or an array of actions, each of which may in turn contain a **Next** key. Note that this data structure allows the expression of the actions as a tree. Actions are executed in the following order:

1. The current action is performed.
2. If the current action has a **Next** key whose value is an action, that action is executed.
3. If the current action has a **Next** key whose value is an array of actions, then each action in the array is executed, in order.

The application should attempt to provide reasonable behavior in extreme situations: actions that are self-referential should not be subsequently executed; actions that close the document or otherwise render the execution of the next action impossible should terminate the execution of the action tree. The application should ensure that a sequence of actions is interruptible by the user.

Note During its execution, an action should not modify its dictionary or the action tree in which it resides. The effect of such modification on subsequent execution of actions in the tree is undefined.

6.8.1 Action Trigger Points

PDF 1.2

In PDF 1.1, the presence of an **A** key or **Dest** key in an annotation or Outline entry denotes an action that is to be performed when the mouse button is released after clicking inside the annotation or Outline entry. PDF 1.2 provides a more general mechanism by defining other “trigger points” (events) and associating actions with each one by means of an “additional actions” dictionary, which is included in an annotation or Outline entry as the value of the **AA** key.



Note The term “mouse” is used to represent a generic pointing device available for use inside a viewing application. It is assumed that the pointing device has the following characteristics:

- A selection button that may be activated, held, and released. This is often the (left) mouse button.
- A notion of location, that is, an indication of where the device is pointing. In many systems, this is usually denoted by the cursor.
- A notion of focus, that is, which element in the document is currently interacting with the user. In many systems, this is denoted by a blinking caret, a focus rectangle, or a color change.

The viewing application and any accessibility programs must ensure that a mapping exists into this environment in order for these actions to be executed correctly.

Table 6.37 Additional Actions attributes

Key	Type	Semantics
E (Enter)	dictionary	(Optional. Defined only for Annotation objects.) The action that is executed when the cursor enters the activation area.
X (Exit)	dictionary	(Optional. Defined only for Annotation objects.) The action that is executed when the cursor exits the activation area.
D (Down)	dictionary	(Optional. Defined only for Annotation objects.) The action that is executed when the user depresses the mouse button inside the activation area.
U (Up)	dictionary	(Optional. Defined only for Annotation objects.) The action that is executed when the user releases the mouse button inside the activation area. For backwards compatibility, the A key in an annotation dictionary takes precedence over the U key in this dictionary.
O (Open)	dictionary	(Optional. Defined only for Page objects.) The action that is executed after the page has finished being drawn. This is independent of any OpenAction that may be defined in the Catalog (see Table 6.1 on page 67), and would be executed after such an action.
C (Close)	dictionary	(Optional. Defined only for Page objects.) The action that is executed when the page is no longer being displayed (e.g., when the user goes to the next page or follows a link). This action applies to the page being closed, and is executed before any other page is opened.
F (Format)	dictionary	(Optional. Defined only for Field objects.) The action to be executed prior to when an Acroform field has its appearance created based on its current value. The action, if it is a JavaScript action, has the opportunity to modify the value for formatting purposes. PDF 1.3
V (Validate)	dictionary	(Optional. Defined only for Field objects.) The action to be executed prior to when an Acroform field's value is changed and saved. The action, if it is a JavaScript action, has the opportunity to reject the change as invalid. PDF 1.3

K (Keystroke) dictionary *(Optional. Defined only for Field objects.)* The action to be executed whenever a user types a keystroke in a text or combo box, or modifies the selection in a list box. The action, if it is a JavaScript action, has the opportunity to reject or modify the change.

PDF 1.3

C (Calculate) dictionary *(Optional. Defined only for Field objects.)* The action to be executed in order to calculate an Acroform field's value based on other field value changes in the document. The action, if it is a JavaScript action, will calculate the new value of the field.

PDF 1.3

Implementation note In PDF 1.2, the keys **FP**, **PP**, **NP**, and **LP** were also defined as Additional Actions attributes. They were never implemented. In PDF 1.3, those keys are obsolete and should be ignored.

PDF 1.3



Note The effects of a particular action attached to one of the field trigger points is essentially limited only by the action itself and can be outside of the described scope of the trigger. For example, even though the format trigger is used to invoke actions that format field values prior to display, it is possible that an action attached at this trigger point could perform a calculation or make any other modification to the document.

Also note that these trigger points can be activated either through user interaction (as described) or programmatically such as the resolution of the **NeedsAppearances** key in the document, FDF import, or JavaScript actions. For example, modifying a field value (by the user) can trigger a cascade of calculations and further formatting and validation for other fields in the document.

Trigger points for mouse-events are defined with the following constraints:

- An Enter event cannot occur unless the mouse button is already up.
- An Up event cannot occur without an Enter event and a Down event occurring.
- An Exit cannot occur without an Enter occurring.
- In the case of overlapping or nested annotations, entering a second annotation's activation area causes an Exit of the first annotation's area to occur.

Note In PDF 1.2, trigger points (**AA** dictionaries) were inherited. In PDF 1.3, they no longer are.

Note The Enter and Exit triggers are determined by the annotation's handler and may not correspond to the annotation's bounding box. Annotation handlers can implement non-rectangular activation regions.

Note Triggers that are related to AcroForms (see [page 129](#)) are defined only for non-button types of Fields.

6.8.2 GoTo action

A GoTo action has the same effect as specifying a destination (with a **Dest** key) in a Link annotation, but it is less compact and is not compatible with PDF 1.0. Destinations are preferred over **GoTo** actions.

Table 6.38 GoTo action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(Required) Action type. Always GoTo .
D (Dest)	Destination	(Required) The destination, as described in Table 7.1 on page 184 .

Example 6.12 GoTo action

```

42 obj
<<
  /Type /Annot
  /Subtype /Link
  /Rect [71 717 190 734]
  /Border [16 16 1]
  /A <<
    /Type /Action
    /S /GoTo
    /D [3 0 R /FitR -4 399 199 533]
  >>
>>
endobj

```

Note This example has the same effect as the Link annotation shown in [Example 6.9 on page 90](#), which uses a destination (a **Dest** key).

6.8.3 GoToR action

The **GoToR** action is similar to the **GoTo** action. However, it includes an additional parameter, the **F** key, that specifies the PDF file that contains the action's destination.

Table 6.39 GoToR action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(Required) Action type. Always GoToR .
D (Dest)	Destination	(Required) An explicit destination, represented by an array, as described in Table 7.1 on page 184 , except that the destination page (the first element of the array) must be specified by a <i>page number</i> , not by an indirect reference to a Page object. The first page is 0.

If the destination is specified as an array, the first element of the array (the destination page) must be a *page number*, not an indirect reference to a Page object.

F (File) File specification		<i>(Required)</i> The file containing the destination.
NewWindow	boolean	<i>(Optional)</i> If the value is <i>true</i> , then the destination document will be opened in a new window; if it is <i>false</i> , then the destination document will be opened in the same window as the source document. If this attribute is omitted, the viewer will behave in accordance with the current user preference.

PDF 1.2



6.8.4 Launch action

The Launch action specifies an application to launch or a document to open. The action must specify the application or document as a file, using the **F** key.

PDF also allows platform-specific information to be included in the Launch dictionary where that information is needed for specific platform. The key **Win** is used for information related to Microsoft Windows launches; the key **Unix** is used for information related to UNIX system launches. If there is no platform specific key, then the **F** key is used.

Table 6.40 *Launch action attributes (in addition to those in Table 6.36)*

Key	Type	Semantics
S (Subtype)	name	<i>(Required)</i> Action type. Always Launch .
F (File) File specification		<i>(Required if there is no alternative key)</i> The file to use in performing the specified action. A viewer that encounters an action with no F key and for which it does not understand any of the alternative keys does nothing.
Win	dictionary	<i>(Optional)</i> Windows-specific launch parameters as described in Table 6.41 .
Unix	<i>undefined</i>	<i>(Optional)</i> Not yet defined.
Mac	<i>undefined</i>	<i>(Optional)</i> Not yet defined.
NewWindow	boolean	<i>(Optional)</i> If the destination is not a PDF document, then this key is ignored. If the value is <i>true</i> , the document is opened in a new window; if it is <i>false</i> , then the document is opened in the same window as the source document. If this attribute is omitted, the viewer behaves in accordance with the current user preference.



PDF 1.2

Table 6.41 *Windows-specific launch attributes*

Key	Type	Semantics
F (File)	string	<i>(Required)</i> The document or application to launch, specified as a DOS file name using standard DOS syntax of eight alphanumeric characters followed by a period, then up to three more alphanumeric characters (commonly referred to as DOS

“8.3” naming). If the string includes a backslash (\), the backslash must itself be preceded by a backslash. Note: this value must be a simple string. It is *not* a [File specification](#).

O (Operation)	string	<i>(Optional)</i> The operation to perform: (open) or (print) . (open) is the default. If the F key specifies an application, this key is ignored and the application is launched.
P (Parameters)	string	<i>(Optional)</i> The parameters passed to the application specified by the F key. If the F key specifies a document, this key should not be provided.
D (Directory)	string	<i>(Optional)</i> The default directory, specified using standard DOS syntax.

6.8.5 Thread action

When a viewer performs a Thread action, it goes to the specified thread and enters thread mode. The thread need not be in the current PDF file.

Table 6.42 *Thread action attributes (in addition to those in [Table 6.36](#))*

Key	Type	Semantics
S (Subtype)	name	<i>(Required)</i> Action type. Always Thread .
F (File) File specification		<i>(Optional)</i> If this key is omitted, then the thread is in the current file. Otherwise this key specifies an external file containing the thread.
D (Dest)		<i>(Required)</i> The desired thread destination. One of the following forms must be provided:
	dictionary	An indirect reference to a thread. This form of the destination requires that the thread be in the current file. (See Section 6.11, “Articles.”)
	number	The index of a thread. (The index of the first thread in a document is 0.)
	Text	The title of the thread. If more than one thread has the same title, the first thread in the document’s list of threads with that title is chosen.
B (Bead)		<i>(Optional)</i> The desired bead in the destination thread. One of the following forms may be provided:
	dictionary	An indirect reference to a Bead dictionary. This form of the bead requires that the thread be in the current file. See Table 6.55 on page 122 .
	number	A number that specifies the bead’s index in the thread. (The index of the first bead in a thread is 0.)

6.8.6 URI action

A Uniform Resource Identifier (URI) is a string that identifies a resource on the Internet, typically a file that is the destination of a hypertext link, although it can also “resolve” to a query or other entity. In PDF, activating a URI action causes the URI to be resolved.

PDF 1.1



Note The URI action is resolved by the Acrobat WebLink plug-in.

Table 6.43 URI action attributes (in addition to those in Table 6.36)

Key	Type	Semantics
S (Subtype)	name	(Required) Action type. Always URI .
URI	string	(Required) The Uniform Resource Identifier to resolve, encoded in 7-bit ASCII.
IsMap	boolean	(Optional) If this key is <i>true</i> , the mouse position should be tracked when link is activated.

In a URI string, any characters following a **#** define a *fragment identifier*.

```
<fragmentID>::= <handler>[=<data>]
```

handler indicates a service, typically provided by a plug-in, that makes use of the resolved URI; *data* (the characters following an equal sign) provides additional information for the handler. The only handler defined by PDF is designated by the string **nameddest**. The data is interpreted as the name of a destination; such a URI action is similar to a **GoToR** action that uses a named destination. If there is no equal sign, then the entire fragment identifier is interpreted as data for the **nameddest** handler; i.e., it is a named destination. For example, the following two examples are equivalent; the first one is preferred.

```
<< /Type /Action /S /URI
  /URI (http://www.adobe.com/techdocs/PDF1.3Man\
  ual.pdf#nameddest=Chap2.Section3)
>>

<< /Type /Action /S /URI
  /URI (http://www.adobe.com/techdocs/PDF1.3Man\
  ual.pdf#Chap2.Section3)
>>
```

The syntax for *handler* is similar to a PDF name object (see page 39) except that **=** is not permitted (because it separates the handler from the data). Furthermore, PDF names allow characters that are not allowed in URI strings. To use such a character in a fragment identifier, in either the handler or the data, write its two hex-digit character code, preceded by a percent sign. The name **X&Y**, for example, would be written as **X%26Y**.

A URI action's **IsMap** attribute indicates that when the action is performed, the (x, y) position of the mouse within the parent link annotation (relative to the upper left hand corner of the link rectangle) should be concatenated to the end of the URI, preceded by a question mark. Here is an example:

```
http://www.adobe.com/intro?100,200
```

Suppose the bounding rectangle in user space of the Link annotation (the value of the **Rect** key) is $[ll_x ll_y ur_x ur_y]$. Given the coordinates of the mouse position in device space, (x_d, y_d) , transform the mouse coordinates to user space, (x_u, y_u) . The final coordinates, (x, y) , are obtained in this way:

$$x = x_u - ll_x$$

$$y = y_u - ll_y$$

Because these coordinates can be fractional and the **IsMap** attribute requires integers, the final coordinates should be rounded to the nearest integer.

6.8.6.1 URI dictionary in the Catalog

In order to support URI action types, the Catalog of the PDF file may include a URI dictionary.

Table 6.44 *URI attributes*

Key	Type	Semantics
Base	string	<i>(Optional)</i> Base URI to resolve relative references. This element allows the URI of the document itself to be recorded in situations in which the document may be accessed out of context. URI actions within the document may be in a “partial” form relative to this base address. When the base address is not specified, the URI is assumed to be the one originally used to locate the document. For example, if a document has been moved but the documents pointed to by relative links within the document have not, the Base key could be used to override the true URI of the document to fix the relative links. This concept is parallel to the description of the body element <BASE> as described in Section 2.7.2 of the HTML specification [17].

6.8.7 Sound actions

PDF 1.2

A Sound action can be used to play a sound from with a PDF document.

Table 6.45 *Sound action attributes (in addition to those in Table 6.36)*

Key	Type	Semantics
S (Subtype)	name	<i>(Required)</i> Object subtype. Always Sound .
Sound	stream	<i>(Required)</i> A Sound object. See Section 6.15 on page 147 .
Synchronous	boolean	<i>(Optional. The default is false.)</i> If this key is <i>true</i> , the viewer does not permit further user interaction, other than cancelling the action, until the sound has been completed played.
Repeat	boolean	<i>(Optional. The default is false.)</i> If this key is <i>true</i> , the sound is repeated indefinitely. If this key is present, the Synchronous key is ignored.

Mix	boolean	<i>(Optional. The default is false.)</i> If this key is <i>true</i> , the viewer attempts to mix this sound with any other sound playing. If this key is <i>false</i> , any playing sound is stopped before starting this sound. This may be used to stop a Repeating sound.
Volume	number	<i>(Optional. The default is 1.)</i> The volume setting for the sound. This number must be between 0 and 1.

Note In PDF 1.2, the value of the **Sound** key was allowed to be a [File specification](#). That is no longer supported, but the same effect can be achieved by using an external stream.

6.8.8 Movie Player actions

A Movie Player action can be used to play a movie in a “floating window” or within the rectangle of a Movie Player annotation. A Movie Player action dictionary is identical to the Movie Activation dictionary described in [Table 6.18](#), with the following additional elements.

Table 6.46 *Movie Player attributes (in addition to those in [Table 6.36](#))*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>								
Type	name	Object type. Always Action .								
S (Subtype)	name	<i>(Required)</i> Object subtype. Always Movie .								
Operation	name	<i>(Optional)</i> The action command for the movie. Currently defined values are: <table> <tr> <td>Play</td> <td>Start playing the movie. The Mode key (see Table 6.18 on page 92) determines the type of play. This is the default value.</td> </tr> <tr> <td>Stop</td> <td>Stop playing the movie and exit the action.</td> </tr> <tr> <td>Pause</td> <td>Pause a playing movie.</td> </tr> <tr> <td>Resume</td> <td>Resume a paused movie.</td> </tr> </table> <p>If Play is selected on a paused movie, the movie is repositioned to the Start position (if specified) before playing.</p>	Play	Start playing the movie. The Mode key (see Table 6.18 on page 92) determines the type of play. This is the default value.	Stop	Stop playing the movie and exit the action.	Pause	Pause a playing movie.	Resume	Resume a paused movie.
Play	Start playing the movie. The Mode key (see Table 6.18 on page 92) determines the type of play. This is the default value.									
Stop	Stop playing the movie and exit the action.									
Pause	Pause a playing movie.									
Resume	Resume a paused movie.									
T (Title)	Text	The title of the Movie annotation to be played. The movie annotation must be contained in the destination page.								
Annot	dictionary	An indirect reference to a Movie annotation.								

Either **T** or **Annot** must be specified, but not both.

PDF 1.2



6.8.9 SetState action

PDF 1.2



Corresponding to the use of appearances for an annotation is the **SetState** action. This action allows for the setting of the value of the **AS** key in the appearance dictionary. This allows another annotation, page, or document-level event to change the state of one or more target annotations. Note that the effect of this action is temporary and does not permanently affect the document.

Table 6.47 *SetState action attributes (in addition to those in Table 6.36)*

Key	Type	Semantics
S (Subtype)	name	(Required) Must be SetState .
T (Target)	string	(Required) The string contains the fully qualified field name of the AcroForm field that is to change state as a result of this action. See page 131 .
AS (Appearance State)	name	(Optional) Indicates which appearance in the sub-dictionary of appearances for the annotation should be shown. This has the effect of setting the appearance state (AS) value in the annotation's dictionary.

6.8.10 Hide action

PDF 1.2



Corresponding to the *Hidden* flag, the Hide action allows the author to determine when a particular annotation is hidden or shown. For example, combined with the action triggers (see [page 108](#)), when the user rolls the mouse over an area of the page, an annotation can appear that describes the action that would occur if the user clicked on that spot on the page. The Hide action is equivalent to setting or clearing the Hidden flag for the annotation. Note that the effect of this action is temporary and does not affect the document.

Table 6.48 *Hide action attributes (in addition to those in Table 6.36)*

Key	Type	Semantics
S (Subtype)	name	(Required) Action subtype. Must be Hide .
T (Target)	string	(Required) The string contains the fully qualified field name of the AcroForm field that is to be hidden or shown. See page 131 .
H (Hide)	boolean	(Optional) If this boolean is <i>true</i> , the action hides the target annotation. If <i>false</i> , the action shows the target annotation. The default for this key is <i>true</i> .

6.8.11 Named actions**PDF 1.2**

PDF 1.2 defines several actions, shown in the following table, that a viewer is expected to support. Additional names may be defined in the future. Viewers may support additional named actions, but a document that uses such names is not portable. If a named action is inappropriate for a viewing platform, or if the viewer does not recognize the name, it should take no action.

Table 6.49 *Named Action Attributes (in addition to those in Table 6.36)*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
S (Subtype)	name	(Required) Action type. Always Named .
N (Name)	name	(Required) One of the names listed in Table 6.50.

Table 6.50 *Named Action List*

<i>Name</i>	<i>Action</i>
FirstPage	Go to the first page of the document.
PrevPage	Go to the previous page in the document.
NextPage	Go to the next page in the document.
LastPage	Go to the last page in the document.

6.8.12 NOP action**PDF 1.2**

The **NOP** action does nothing when it is executed. It exists primarily as a means of overriding a trigger point that would inherit behavior either from some other part of the document or from the viewer itself. For example, a page could specify the **NOP** action for the **NP** key of its Additional Actions dictionary to prevent anything from happening when the user attempted to go to the next page of the document.

Table 6.51 *NOP Action Attributes (in addition to those in Table 6.36)*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
S (Subtype)	name	(Required) Action type. Always NOP .

Note In PDF 1.3, the **NOP** action is no longer defined and should be ignored. In PDF 1.2, **NOP** was defined as a way to block the inheritance of the actions associated with a trigger point. In PDF 1.3, there is no inheritance for trigger points, and therefore no need for the **NOP** action.

PDF 1.3

6.9 Names dictionary

PDF 1.2

The Catalog of a document may contain a **Names** key. The value of this key is a dictionary. Each value in this dictionary is a [Name tree](#), which is a tree, similar to the Pages tree, where the leaf nodes contain pairs of strings (the “names”) and objects (the “values”). A name tree has the same purpose as a dictionary, mapping keys to values, but it does so in a different manner: the keys in a name tree are strings, not PDF name-objects.

Table 6.52 *Names dictionary in the Catalog*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Dests	Name tree	The root of a name tree for named destinations. Destination names are encoded with PDFDocEncoding . (See Section 7.3.2 on page 185.)
Pages	Name tree	The root of a Pages tree for AcroForms. (See Section 6.14.14 on page 146.)
Templates	Name tree	The root of a Templates tree for AcroForms. (See Section 6.14.14 on page 146.)
URLS	Name tree	<i>(Optional)</i> The root of a name tree that maps URLs to Web Capture Content Sets (or arrays of Web Capture Content Sets). (See Section 6.16.2 on page 152.)
IDS	Name tree	<i>(Optional)</i> The root of a name tree that maps IDs to Web Capture Content Sets (or arrays of Web Capture Content Sets). (See Section 6.16.2 on page 152.)
JavaScript	Name tree	<i>(Optional)</i> The root of a name tree for document-level JavaScript actions.
		Note that the names of the individual JavaScripts in the Names dictionary serve as a convenient means for organizing and packaging scripts. The names are arbitrary and specified by the user when authoring the document and may or may not have anything to do with the JavaScript name space. (See Section 6.14.13.4 on page 145.)
AP	Name tree	<i>(Optional)</i> The root of a name tree for named appearances. Appearances can be referenced from multiple annotations. In order to save space in PDF files, appearance attributes can be “named.” The values associated with the names are appearance dictionaries

PDF 1.2

PDF 1.3

PDF 1.3

PDF 1.3

PDF 1.3

PDF 1.3

PDF 1.3

The values in the leaf nodes of the **URLS** and **IDS** name trees contain indirect references to either a Web Capture Content Set object or an array of such objects. The purpose of the array is to allow a single name, such as a URL, to map to multiple Web Capture Content Sets. The Web Capture Content Sets need not, and usually don’t, have the same subtype. For example, it is legal to have the same URL map to multiple Web Capture Page Set objects, or to both Web Capture Page Set objects and Web Capture Image Set objects. See [Figure 6.13](#) for an example.

6.10 Info dictionary

A document's trailer may contain a reference to an Info dictionary that provides information about the document. This optional dictionary may contain one or more keys, whose values are strings. In most cases, the strings are actually [Text](#). These strings may be displayed and/or updated using a viewer of PDF files. For this reason, the information in these keys may not always accurately reflect the state of the PDF file.

Implementation note These strings are displayed in an Acrobat viewer's Document Info dialog.

Note Omit any key in the Info dictionary for which a value is not known, rather than including it with an empty string as its value.

Table 6.53 PDF Info dictionary attributes

Key	Type	Semantics	
Author	Text	(Optional) The name of the person who created the document.	
CreationDate	Date	(Optional) The date the document was created.	
ModDate	Date	(Optional) The date the document was last modified.	PDF 1.1
Creator	Text	(Optional) If the document was converted into a PDF document from another form, this is the name of the application that created the original document.	
Producer	Text	(Optional) The name of the application that converted the document from its native format to PDF.	
Title	Text	(Optional) The document's title.	PDF 1.1
Subject	Text	(Optional) The subject of the document.	PDF 1.1
Keywords	Text	(Optional) Keywords associated with the document.	PDF 1.1
Trapped	name	(Optional) This key is used to indicate whether the data in this PDF file has been modified to express trapping information. <i>Trapping</i> means the modification of the boundaries of colored objects within the PDF file to prevent the media color from accidentally showing through when color planes are slightly misregistered in a reproduction process. The defined names which this key can take are:	PDF 1.3
	True	The file has been modified to include trapping data and no further trapping is needed; that is, all necessary traps for all pages of the file have been created. (Note that this is the name /True , not the boolean true .)	
	False	No trapping modifications have been made to any page or object of the entire file. Any desired trapping must be done before final rendering of the file. (Note that this is the name /False , not the boolean false .)	
	Unknown	Either it is unknown whether the file has been trapped, or it is the case that some but not all necessary traps have been created.	

If the **Trapped** key is omitted, the default value of the key is **Unknown**. The **Trapped** key may be set either by some processor that created traps or, because trapping information may be generated at so many places in the workflow that created the PDF file, the amount of trapping may only be known to a human operator and it may be necessary for the trapping information to be manually entered via a user interface to the Info dictionary.

Info strings that are to be interpreted as dates must include the **D:** prefix (see on [Section 7.2 on page 183](#)). In particular, the 1.0 key **CreationDate** and the 1.1 key **ModDate** should use this format. All Info strings that represent dates should be displayed as a human-readable date. Other Info strings are uninterpreted.

Info keys and strings may be added to or changed by users or extensions, and some extensions may choose to permit searches on these keys. PDF 1.1 does not define short names for the keys in [Table 6.53](#), to make it easier to browse and edit Info dictionary entries. New names should be chosen with care so that they make sense to users.

Although private data can be stored in the Info dictionary, it is more appropriate to store it in the Catalog. This allows a user or program to alter entries in the Info dictionary with less chance of unforeseen side effects.

The following example shows an Info dictionary.

Example 6.13 *Info dictionary*

```
1 0 obj
<<
/Creator (Adobe Illustrator)
/CreationDate (D:19930204080603-08'00')
/Author (Werner Heisenberg)
/Producer
(Acrobat Network Distiller 1.0 for Macintosh)
>>
endobj
```

6.11 Articles

PDF 1.1

An *article thread* identifies related elements in a document, enabling a user to follow a flow of information that may span multiple columns or pages.

A PDF document may include one or more article threads. Each thread has a title and a list of thread elements, which are referred to as *beads*. A viewer may allow the user to select a particular thread and then navigate through it. As in newspaper or magazine articles, beads are not necessarily located on contiguous pages or in sequential page order. The viewer automatically maintains a comfortable zoom level for reading and moves from one bead to the next, rather than from one page to the next.



If a document includes any threads, they are stored in an array as the value of the **Threads** key in the Catalog object. Each thread and its beads are dictionaries. [Table 6.54](#) lists the attributes of a Thread dictionary, and [Table 6.55](#) lists the attributes of a Bead dictionary.

Table 6.54 *Thread attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
F (First)	dictionary	<i>(Required; must be an indirect reference)</i> Specifies the first bead in this thread.
I (Info)	dictionary	<i>(Optional)</i> Information about the thread. This dictionary should contain information similar to the document's Info dictionary and should use the same key names and data formats for entries that correspond to Info dictionary entries. Most entries in this dictionary will be of type Text .

Table 6.55 *Bead attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
T (Thread)	dictionary	<i>(Required for the first bead of a thread; must be an indirect reference)</i> The thread in which this bead is an element.
V (Prev)	dictionary	<i>(Required; must be indirect)</i> The previous bead of this thread; for the first bead in a thread, V specifies the last bead in the thread.
N (Next)	dictionary	<i>(Required; must be indirect)</i> The next bead of this thread; for the last bead in a thread, N specifies the first bead in the thread.
P (Page)	dictionary	<i>(Required; must be indirect)</i> The Page on which this bead appears.
R (Rect)	Rectangle	<i>(Required)</i> Rectangle specifying the location of this bead.

PDF 1.2

The following example shows a thread with three beads:

Example 6.14 *Thread*

```

22 0 obj <<
  /F 23 0 R
  /I << /Title (Man Bites Dog) >> >>
endobj

23 0 obj <<
  /T 22 0 R
  /V 25 0 R
  /N 24 0 R
  /P 8 0 R
  /R [158 247 318 905] >>
endobj

```

```

24 0 obj <<
  /T 22 0 R
  /V 23 0 R
  /N 25 0 R
  /P 8 0 R
  /R [322 246 486 904] >>
endobj

25 0 obj <<
  /T 22 0 R
  /V 25 0 R
  /N 23 0 R
  /P 10 0 R
  /R [157 254 319 903] >>
endobj

```

The Page object for each page on which beads appear should contain a **B** key, as described in [Section 6.4, “Page objects.”](#) The value of this key is an array of indirect references to each bead on the page, in drawing order.

6.12 File ID

PDF 1.1

A PDF file may contain a reference to another PDF file. Storing a file name, even in a platform-independent format, does not guarantee that the file can be found, even if it exists and its name has not been changed. Different server software applications often present different names for the same file. For example, servers running on DOS platforms must convert all file names to eight letters and a three-letter extension. Different servers use different strategies for converting long names to this format.

References to PDF files can be made more reliable by making the PDF file reference consist of two parts: (1) a normal operating system-based file reference and (2) a file ID. The file ID characterizes the file and is stored with the file. Placing a file ID with the file reference and in the file itself increases the chances that a file reference can be resolved correctly. Matching the ID in the reference with the ID in the file indicates whether the desired file was found.

A PDF file may have an **ID** key in its trailer. The value of this key is an array of two strings. The first element is a permanent ID, based on the contents of the file at the time the file was created. This ID does not change when the file is incrementally updated. The second element is a changing ID, based on the contents of the file at the time the file is incrementally updated. When a file is first written, the IDs are set to the same value. When resolving a file reference, if both IDs match, it is very likely that the correct file has been found. If only the first ID matches, then a different version of the correct file has been found.



To help insure the uniqueness of the file ID, it is recommended that the file ID be computed using a message digest algorithm such as MD5, as described in *RFC 1321: The MD5 Message-Digest Algorithm* [35]. It is recommended that the following information be passed to the message digest algorithm:

- the current time
- a string representation of the location of the file, usually a path name
- the document size in bytes
- the value of each entry in the document's Info dictionary.

Implementation note

Adobe applications pass this information to the MD5 message digest algorithm to calculate file IDs. Note that the calculation of the file IDs need not be reproducible. All that matters is that the file IDs are likely to be unique. For example, two implementations of this algorithm might use different formats for the current time. This will cause them to produce different file IDs for the same file created at the same time, but this does not affect the uniqueness of the ID.

6.13 Encryption dictionary

PDF 1.1

Documents can be protected via encryption, as described in [Section 5.18, "Encryption."](#) Every protected document must have an Encryption dictionary, which specifies the security handler to be used to authorize access to the document. The Encryption dictionary also contains whatever additional information the security handler chooses to store in it.

The value of the **Filter** key is the name of the security handler that encrypted the document.

The **V** key indicates the algorithm that is used to encrypt and decrypt document data, as described below. A security handler may require additional keys in the Encryption dictionary. The keys required by the built-in standard security handler are described below.

The **V** key is set and used by the PDF consumer (e.g., viewer) rather than by the security handler. A value of 1 corresponds with [Algorithm 5.3 on page 63](#). A value of 0 corresponds to an algorithm that is no longer supported and has been deprecated. Values greater than 1 are not defined for PDF 1.3; PDF files with values greater than 1 cannot be opened by PDF viewers that support PDF 1.3.

Standard encryption

The built-in encryption method provided by PDF allows for the following functionality. A document has two passwords: an owner password and a user password. The document also specifies operations that should be restricted even when the document is decrypted: printing; copying text and graphics out of the document; modifying the document; and adding or modifying text notes and AcroForm fields. When the correct user password is supplied, the document is

opened and decrypted but these operations are restricted; when the owner password is supplied, all operations are allowed. The owner password is required to change these passwords and restrictions.

A document is encrypted whenever a user or owner password or restrictions are supplied for the document. However, a user is prompted for a password on opening a document only if the document has a user password. It is possible to determine if a document has a user password by testing the empty string as the user password, as described in [Algorithm 6.10](#) below.

Note If the Owner and User passwords are the same, the document is opened with User permission. It is, therefore, impossible at this point to get Owner permission for the file.

Note [Table 6.56](#) lists the Encryption dictionary attributes of a document that is encrypted using the standard security handler. The data in this dictionary is used to determine if a candidate password string is the user password, the owner password, or neither. If it is the user password, the dictionary data is used to determine the operations that should be restricted. The calculation and use of data in this dictionary is described below.

Table 6.56 Standard security handler attributes

Key	Type	Semantics
F (Filter)	name	(Required in all encrypted documents) Name of security handler. The name of the built-in handler is Standard .
R (Revision)	number	(Required if F is Standard) Revision number of algorithm used to encode data in this dictionary. The current revision number for the Standard security handler is 2.
O (Owner)	string	(Required if F is Standard) Data that can be used to determine whether or not the owner password was entered. Contains an encrypted version of the padded user password. Step 1 of Algorithm 6.7 shows how to pad the user password.
U (User)	string	(Required if F is Standard) Data that can be used to determine whether or not the user password was entered. Contains an encrypted version of the fixed padding string shown in step 1 of Algorithm 6.7 .
P (Permissions)	integer	(Required) Collection of flags describing permissions granted to user who opens a file with the user password. See Table 6.57 .
V (version)	number	(Optional) Indicates the algorithm that is used to encrypt and decrypt document data. The value should be 1, indicating Algorithm 5.3 on page 63 . If the value is omitted, the default value is 0, but that indicates an algorithm that is undocumented and deprecated. It is strongly recommended that a value of 1 be specified.

The value of the **P** (Permissions) key is an unsigned 32-bit integer that contains a collection of access rights. These rights are enabled on opening the document with the user password if the corresponding bit is set in the integer. [Table 6.57](#) specifies the meanings of the bits, with bit 1 being the least significant.

Table 6.57 *Permission flags*

<i>Bit position</i>	<i>Semantics</i>
1–2	Reserved. Must be set to 0.
3	Enable printing of the document.
4	Enable changing the document other than by adding or changing text notes and AcroForm fields.
5	Enable copying of text and graphics from the document.
6	Enable adding and changing text notes and AcroForm fields.
7–32	Reserved. Each bit must be set to 1.

As described in [Section 5.18, “Encryption,”](#) one function of a security handler is to produce a five-byte key that is provided as input to the PDF encryption algorithm. Given a password string, the standard security handler computes an encryption key using [Algorithm 6.7](#):

Algorithm 6.7 *Computing an encryption key*

1. Truncate the password string to 32 bytes. (Treat an omitted password as an empty string.) If the string is less than 32 bytes long, pad it with bytes from the following string to make it exactly 32 bytes long:

```
<28 BF 4E 5E 4E 75 8A 41 64 00 4E 56 FF FA 01 08
 2E 2E 00 B6 D0 68 3E 80 2F 0C A9 FE 64 53 69 7A>
```

Pad the password string from the beginning of the pad string. That is, if the password string is n bytes long, concatenate bytes 1 to $32 - n$ of the pad string to the password string.

2. Provide the string as input to the MD5 hash function.
3. Input the value of the Encryption dictionary’s **O** (Owner) key (a 32-byte string) to the MD5 hash function. ([Algorithm 6.8](#) explains how the **O** string is computed.)
4. Treat the value of the **P** (Permissions) key as an unsigned four-byte integer, and provide these bytes to the MD5 hash function, low-order byte first.
5. Input the first element of the file’s **ID** to the MD5 hash function. (See [Section 6.12, “File ID.”](#))

6. The first five bytes of the MD5 output make up the encryption key.

The encryption key is used to encrypt strings and stream data, as described in [Section 5.18, “Encryption.”](#) The same key decrypts data as well.

To create an encrypted document, the standard security handler must compute the entries in the Encryption dictionary as well as the encryption key. While calculation of the value of the **F**ilter, **R** (Revision), and **P** (Permissions) keys are straightforward, calculation of the value of the **O** (Owner) and **U** (User) keys requires further explanation.

[Algorithm 6.8](#) shows how to calculate the value of the **O** key in the Encryption dictionary:

Algorithm 6.8 *Computing the O (Owner) key in the Encryption dictionary*

1. Truncate or pad the owner password string as needed, following the same process as in step 1 of computing an encryption key ([Algorithm 6.7](#)). If there is no owner password, the user password is used instead.
2. Provide the string as input to the MD5 hash function.
3. Create an RC4 key using the first five bytes of the MD5 output.
4. Truncate or pad the user password string as needed, following the same process as in step 1 of [Algorithm 6.7](#).
5. Encrypt the padded user password string using the RC4 algorithm with the key from step 3.
6. Store the encrypted string as the value of the **O** key in the Encryption dictionary.



[Algorithm 6.9](#) shows how to calculate the value of the **U** key in the Encryption dictionary:

Algorithm 6.9 *Computing the U (User) key in the Encryption dictionary*

1. Create an encryption key ([Algorithm 6.7](#)) based on the user password string.
2. Encrypt the 32-byte string specified in step 1 of [Algorithm 6.7](#) using the RC4 algorithm with the encryption key from the previous step.
3. Store the encrypted string as the value of the **U** key in the Encryption dictionary.

The standard security handler performs several other functions beyond calculating the key needed to encrypt or decrypt data, as described above. It determines if the candidate password string is the user password, the owner password, or neither. If it is the user password, it determines the operations that should be restricted.

Given a password, the standard security handler uses the contents of the Encryption dictionary to determine if a document should be opened and what permissions should be granted. If the password is the user password, the document is opened only with the permissions specified in the Encryption dictionary. If the password is the owner password (but not the same as the user password), all permissions are enabled.

To determine if a password is the correct user password, the security handler uses [Algorithm 6.10](#):

Algorithm 6.10 *Checking for the correct user password*

1. Use steps 1 through 3 of [Algorithm 6.7](#) to compute an encryption key from the password string.
2. Use RC4 with this encryption key to decrypt the value of the **U** key.
3. If the decrypted value of the **U** key is the fixed padding string shown in step 1 of [Algorithm 6.7](#), the password is the user password.

To determine if a password is the correct owner password, the security handler uses [Algorithm 6.11](#):

Algorithm 6.11 *Checking for the correct owner password*

1. Use steps 1 through 3 of [Algorithm 6.7](#) to compute an encryption key from the password string.
2. Use RC4, with this encryption key, to decrypt the value of the **O** key.
3. Use [Algorithm 6.7](#) to compute an encryption key from the decrypted value of the **O** key.
4. Use RC4, with this encryption key, to decrypt the value of the **U** key.
5. If the decrypted value of the **U** key is the fixed padding string shown in step 1 of [Algorithm 6.7](#), the password is the owner password.

Note Despite the specification of document permissions in a PDF file, PDF cannot enforce the restrictions specified. It is up to the implementors of PDF viewers to respect the intent of the document creator by limiting access to an encrypted PDF file according to the permissions and passwords contained in the file.

Document creators have two choices if the standard encryption provided by PDF is not sufficient. They can use an alternative, more secure, security handler, or they can encrypt whole PDF documents, bypassing PDF security entirely.

6.14 Acrobat Forms

PDF 1.2

Beginning with version 1.2, PDF defines a number of interactive document features. Prominent among these is the feature that allows a PDF document to represent a *form*. By form we mean the PDF equivalent of the familiar paper instrument and not the Form XObject defined on [page 255](#). The term *AcroForm* refers to the feature defined here; the other is referred to explicitly as a Form XObject.

An AcroForm consists of a collection of *fields*. A field should be thought of as a set of properties. The three most important properties of a field are its type, its name, and its value. Other properties are used to specify the appearance of a field. Fields can be organized into a hierarchy, and there are still other properties of the field that associate it with its parent and children.

6.14.1 AcroForm

A PDF file may contain at most one AcroForm, although that AcroForm may have an arbitrary number of fields that appear on any page of the PDF file. The properties of the AcroForm itself are encoded in a dictionary that is referenced in the Catalog dictionary of the PDF file using the key **AcroForm**.

Note Although each PDF file contains at most a single AcroForm, arbitrary sets of fields may be imported and exported from the PDF file; see [Section 6.14.13 on page 143](#).

The AcroForm dictionary has the following attributes:

Table 6.58 *AcroForm dictionary attributes*

Key	Type	Semantics
Fields	array	<i>(Required)</i> This array contains a reference to each root field in the PDF file. A root field is a field with no parent.
NeedAppearances	boolean	<i>(Optional)</i> AcroForms use annotations; in particular, AcroForms use the AP dictionary of an annotation to represent the appearance of a field's value. Form-authoring applications can omit the AP dictionary for fields that contain text (see Section 6.14.5 on page 133). When the PDF file is opened, if NeedAppearances is <i>true</i> , the viewer creates an AP dictionary and all of its subparts including the N dictionary and the appearance stream it refers to, for any field containing variable text that does not have one. The default value for this attribute is <i>false</i> .
CO (Calculation Order)	array	<i>(Required if calculation scripts are specified for fields in the document)</i> A JavaScript key (see page 145). An array of indirect pointers to field dictionaries. The order of fields within the calculation order array is the order in which actions attached to the fields' calculation triggers will be executed.

PDF 1.3

SigFlags	integer	<p><i>(Optional)</i> Signature flags. Collection of document-level flags defining various characteristics related to signature fields (see Section 6.14.12 on page 141). The default value is 0. The following flags are defined:</p> <p>bit 1 <i>Signatures-exist</i> flag. A value of 1 indicates that at least one of the fields in the document is of type Sig. This is intended for use by viewer interfaces to enable menu items or buttons appropriate for manipulating signatures. This bit is always available with the first page of a linearized file, while scanning all the fields potentially requires loading all the pages.</p> <p>bit 2 <i>Append-only</i> flag. A value of 1 indicates that the user is to be discouraged from <i>saving</i> (writing) the file in a way that alters the previous contents of the file, such as saving with the “optimize” option, because doing so will likely invalidate signatures contained in the document. Merely <i>updating</i> the file, which involves appending new information to the end of the previous of the version of the file, is fine. (See Section A.6 on page 435.) The intended use of this flag is for the viewer to present the user doing a save/optimize with an additional dialog box stating that signatures will be invalidated and asking whether to continue, with a default action of canceling the save.</p>
-----------------	---------	---

PDF 1.3

In addition to the attributes in [Table 6.58](#), the AcroForm dictionary can contain document-wide defaults for the **Q**, **DR**, **OH** and **DA** field dictionary attributes, described below.

In order to ensure efficient performance of field value calculation, it is necessary to calculate field values in a single pass through all of the calculation scripts in the document. Because JavaScript is a full programming language, it is difficult to determine dependencies between calculations and the proper calculation order based on a simple analysis of these scripts.

PDF 1.3

The order of fields within the Calculation Order array is the order in which actions attached to the fields' calculation triggers will be executed.

6.14.2 Fields

The properties that define each distinct field are stored in a dictionary. All such dictionaries must be indirect objects. A particular property of a field is found in one of three ways. In the simplest case, the property is found immediately in the dictionary, referenced by a predefined key. For example, the simple way to represent the type of a checkbox field is to place the key-value pair

```
<< /FT /Btn >>
```

in the field's dictionary, indicating that is a button.

Fields are arranged in a hierarchy, and many properties that can be stored directly in the dictionary can also be inherited from the field's parent. In general, an inherited property might be found in the grandparent or some arbitrary ancestor of the field.

One field property is neither stored directly in the dictionary nor inherited from an ancestor. This property, the *fully qualified field name*, must be derived from the **T** attribute of the field and its ancestors. A fully qualified field name is an ordered collection of the partial field names of a field and all of its ancestors. The partial field name is stored in the dictionary as the key-value pair

<< /**T** *string* >>

The fully qualified field name is formed by concatenating the partial field names, separated by periods, such as (**a.b.c.d**).

One way of producing the ordered collection is to store the partial field names in an array. Using this technique, two fully qualified field names are identical if the arrays are of equal size and the components of the array are pairwise identical.





It is possible for different fields to have the same fully qualified field name, albeit under a restricted set of circumstances. All fields with the same fully qualified field name must be descendants of a common ancestor with that same name. In practice, this means that none of these descendants can have a **T** attribute. If different fields have the same fully qualified field name, the fields should differ only in properties that specify appearance. In particular, whenever different fields have the same fully qualified field name, they must have the same type, value, and default value, as described below. Given a fully qualified field name, there is an effective way to find the common ancestor of all fields with that name.

6.14.3 Field dictionaries

All types of fields share the following attributes. Since the notion of required and optional fields is somewhat clouded by inheritance, in addition to the normal categories of *Required* and *Optional*, we use the terms *Required, inheritable* for attributes that *must* exist somewhere in the inheritance hierarchy, and *Optional, inheritable* for attributes that *may* exist somewhere in the inheritance hierarchy.

Table 6.59 *Attributes common to all types of fields*

Key	Type	Semantics
FT	name	(<i>Required, inheritable</i>) The type of field. Valid field types include: <ul style="list-style-type: none"> Btn Button field. This field type is used for a checkbox, radio button, or a push button. A checkbox is a field that toggles between two states, on and off. A radio button is a field made up of an arbitrary number of toggles, only one of which can be on at any given time. A push button is a simple interactive element that retains no value. See Section 6.14.6 on page 135. Tx Text field. A field whose value is text. Text may be single- or multi-line. See Section 6.14.11 on page 139. Ch Choice field. Like a radio button, this field type takes on one value from an arbitrary set. It is presented to the user as a pop-up list or a scrollable list, rather than as a set of mutually exclusive toggles. See Section 6.14.10 on page 138. Sig Signature field. See Section 6.14.12 on page 141.

T	Text	<p><i>(Optional)</i> The partial field name. The fully qualified field name is derived from the partial field name, as described above.</p> <p>Acrobat 4.0 does not support Unicode encoding of field names.</p>	 
V	<i>various</i>	<p><i>(Optional, inheritable)</i> The value. The format of the value is dependent on the type of the field as specified by the FT attribute. See the following sections for details on the field values for particular types.</p>	
DV	<i>various</i>	<p><i>(Optional, inheritable)</i> The default value. The value of the field reverts to this when a ResetForm action (see Section 6.14.13.2 on page 144). The format of DV is the same as V.</p>	
Ff	integer	<p><i>(Optional, inheritable)</i> Flags. Collection of flags defining various characteristics of the field. The default value is 0. The following flags apply to any field, regardless of field type:</p> <p>bit 1 <i>Read-only</i> flag, indicating that the user is not allowed to change the field value in the viewer. This is intended for use with computed fields or fields that are imported from a database.</p> <p>bit 2 <i>Required</i> flag, indicating that the this field must have a value at the time the field is exported by a Submit action; see Section 6.14.13.1 on page 143.</p> <p>bit 3 <i>No-export</i> flag, indicating that this field must not be exported by a Submit action; see Section 6.14.13.1 on page 143.</p>	
Kids	array	<p><i>(Optional)</i> The Kids and Parent attributes define the child/parent relationship between field dictionaries. This array contains references to child field dictionaries of this field. A child may be referenced from only one Kids array and thus has a unique parent.</p> <p>A field dictionary may use a Kids array to reference sub-fields (i.e. fields that share a common prefix with their fully qualified field). That shared prefix is the fully qualified field name of this field. For instance, you might create a field dictionary whose T key is (Address) with sub-field dictionaries whose T keys are (Street) and (City).</p> <p>A field dictionary may use a Kids array to reference child fields that differ from this field only in their appearance. Such child fields inherit FT, V, and DV attributes and may not specify a T attribute.</p>	
Parent	dictionary	<p><i>(Required if this field is referenced from another field's Kids array)</i> The parent field.</p>	
TU	Text	<p><i>(Optional)</i> User Name. The name to be used when generating error or status messages for a field.</p>	
TM	Text	<p><i>(Optional)</i> Mapping Name. The name to be used when exporting AcroForm field data from a document.</p>	

6.14.4 Widget annotations

AcroForms use annotations to represent the appearance of fields and to manage user interactions. These annotations have a subtype of **Widget**.

If a field has an appearance, that appearance must be represented with an **AP** dictionary (see [page 87](#)). Conversely, if an **AP** dictionary exists, the appearance it represents must be consistent with the value of the field.

As an optimization, the attributes of a field are merged into a single dictionary with the attributes of the Widget annotation that represents the appearance of the field and manages user interactions. Hence, a field may also be a bona fide annotation that is referenced in the **Annots** array of some Page dictionary. The keys for a field dictionary and an annotation dictionary are mutually exclusive, so there is no ambiguity in doing this. A field is also a Widget annotation if and only if it contains a **Subtype** attribute whose value is **Widget**. A Widget annotation is also a field if and only if it contains a **T** attribute, or if it contains a **Parent** attribute one of whose ancestors, as found by following the **Parent** references, contains a **T** attribute.

6.14.5 Fields comprising variable text

Both text fields and choice fields can contain variable text (i.e., text stored in field dictionary attributes that are used to generate the appearance of the field). The sections on text and choice fields, below, contain more information on exactly which attributes are used in each case. When there is a change in the attributes of a text or choice field that affect the appearance, the viewer must generate an appearance stream for the field, contained in the **AP** dictionary of the Widget annotation, conforming to these new attributes. A number of additional attributes are defined to specify the initial appearance of this text.

Table 6.60 *Attributes common to all types of fields containing variable text*

Key	Type	Semantics
DR	dictionary	<i>(Required, inheritable)</i> Default resources. This is a Resources dictionary (see page 195). At a minimum, this dictionary must contain a resource of type Font . The Font resource is, in turn, a dictionary which must contain at least one key-value pair specifying the resource name and resource dictionary of the default font for this field. All resources from the DR dictionary are copied to the appearance stream when it is created or updated.
DA	string	<i>(Required, inheritable)</i> Default appearance string. This string contains a sequence of valid page-content graphics or text-state operators that determine such appearance properties as text color, text size, etc.
H (Highlight)	name	<i>(Optional)</i> Indicates how click-feedback appears.
Q	integer	<i>(Optional, inheritable)</i> Quadding. This determines whether the input text is left-justified (0), centered (1) or right-justified (2). The default value is 0.

OH	integer	<i>(Optional)</i> Original height. This determines the height of the BBox attribute of the appearance stream. The default for this attribute is the height of the Rect attribute of the Widget annotation that contains the appearance stream.
-----------	---------	--

A viewer must generate the appearance stream according to the following procedures. There are two cases to consider. In the first case, the appearance stream is created where none existed previously. In the second, an existing appearance stream is updated to reflect a new value.

If the appearance stream is created, when the Widget annotation used to represent the appearance of the field contains no appearance stream, a viewer performs the following actions:

- Create the appearance dictionary, **AP**, in the Widget annotation. Create an empty stream. Set the normal face, **N**, of the appearance dictionary to the newly created stream. Set the Resources dictionary of the newly created stream to **DR**. Set the **BBox** attribute of the stream. The lower and left coordinates of the **BBox** are set to 0. The top coordinate is set to **OH**. The right coordinate is set to the width of the **Rect** of the Widget annotation divided by the ratio of **OH** to the height of the **Rect** of the annotation. All other stream attributes are set to their defaults.
- Insert **/Tx BMC q** into the stream data. The **BMC** operator labels the stream data as marked content with the tag **Tx** (see [page 353](#)). The **q** operator saves the current graphics state.
- If any graphics-state changes are required, such as clipping, emit them here.
- Insert **BT** to begin the text object.
- Insert the **DA** string into the stream data, subject to the following constraints:

Graphics state operators, if specified, must be operators that are legal within a text object (see [Figure 8.1 on page 320](#)). At a minimum, a **Tf** operator along with its two parameters, *fontname* and *size*, must be present. The *fontname* must match a resource name in the **Font** resource of the default resource dictionary, **DR**. If *size* is 0, then the size of the font is computed as a function of **OH** (or the height of the annotation's rect when **OH** is not defined), and that value is inserted instead of 0. (This is referred to as "auto-sizing.") If a **Tm** operator is present, the *x* and *y* translation components are replaced with positioning information determined by the viewer to be appropriate, based on the field value, the **Q** attribute, and any layout rules it employs.

Note The **DA** string should contain at most one **Tm** operator.

- If the **DA** string contains no **Tm** operator, emit **Tm** with appropriate *x* and *y* translation components, as described above.
- Emit any text string operators necessary to show the variable text along with any additional, necessary text positioning operators.

- Emit **ET Q EMC**. The **ET** operator ends the text object, the **Q** operator restores the previous graphics state, and the **EMC** operator indicates the end of the content that was marked with tag **Tx**.

If an existing appearance stream is updated to reflect a new value, when an appearance stream already exists, the viewer performs the following actions:

- Copy the resources from **DR** to the Resources dictionary of the appearance stream. Conflicts between resource names in **DR** and the existing appearance stream are resolved in favor of the appearance stream (i.e., no copy takes place).
- Search the appearance stream for the sequence **/Tx BMC** and the matching **EMC**, accounting for any nested **BMC/EMC** pairs in this matching process. Valid appearance streams should have exactly one such **BMC** operator with tag **Tx**.
- Replace the stream contents identified in the previous search by the stream contents generated in the first case beginning at the second step; i.e. emit **/Tx BMC q**, etc. If the original appearance stream contains no content marked with tag **Tx**, then append the new stream contents to the end of the original stream.

6.14.6 Button field

The **Btn** field type is used for checkboxes, radio buttons, and simple push buttons that retain no value. Simple push buttons are distinguished from checkboxes or radio buttons by the *push-button* flag. Radio buttons and checkboxes are distinguished with the *radio* flag.

Table 6.61 *Field flags (Ff) for Btn fields*

<i>Bit Position</i>	<i>Semantics</i>
16	<i>Radio</i> flag. If this flag is <i>true</i> (i.e., if this bit is 1), then the Btn field is a radio button; otherwise it is a checkbox. This flag is meaningful only if the <i>push-button</i> flag is <i>false</i> .
17	<i>Push-button</i> flag. This flag is used to indicate a button field that is “purely interactive”; i.e., one that responds to user input in some way but does not change state as a result of that input.

6.14.7 Push button

A push button is the simplest type of field. It has no value. It can assume any of the attributes common to all fields except the **V** and **DV** attributes. The **FT** attribute must be set to **Btn**, the **Ff** key is required, and the *push-button* flag must be set to 1.

6.14.8 Checkbox

A checkbox is a field that can assume two states, *on* and *off*. Each state can have a separate appearance created from arbitrary marking operations. These appearances are found in the **AP** dictionary of the field annotation; see [page 87](#). The appearance for the *off* state is optional but, if present, must be stored in the **AP** dictionary under the name **Off**. The recommended name for the *on* state is **Yes**, but this is not required. The **Ff** attribute of a checkbox field should have both the *push-button* and *radio* flags set to 0. In addition to the attributes common to all fields, a checkbox field can contain the following attributes:

Table 6.62 *Checkbox attributes*

Key	Type	Semantics
FT	name	(Required, inheritable) Field type. Must be Btn .
Ff	integer	(Optional, inheritable) Flags. The <i>push-button</i> and <i>radio</i> flags must both be set to 0.
V	name	(Optional, inheritable) Field value. This indicates which of the two states the checkbox is in. The value of a checkbox is the name that is used to identify its appearance in the AP dictionary. So, if the checkbox is in the <i>off</i> state, then the value is Off . If it is in the <i>on</i> state, then the value contains whatever name is used to identify this state in the AP dictionary. As mentioned above, it is recommended but not required that this name be Yes . The default value for this attribute is Off .

Example 6.15 *Simple checkbox field with Zapf Dingbats check character (X)*

```

1 0 obj
<< /FT /Btn
/T (Urgent)
/V /On
/AS /On
/AP << /N << /On 2 0 R >> >>
>>
endobj

2 0 obj
<< /Resources 20 0 R /Length 50 >>
stream
q 0 0 1 rg BT /ZaDb 12 Tf 0 0 Td (8) Tj ET Q
endstream
endobj

```

6.14.9 Radio button

A radio button is a field that can assume at most one of $n+1$ predefined states. The radio button is conceptually formed from n checkboxes, with the added semantics that selecting any one of the checkboxes deselects all the others. The PDF

representation of a radio button is $n+1$ fields, one for the radio button itself and one for each of the underlying checkboxes. The **Ff** attribute of a radio button has the *push-button* flag set to 0 and the *radio* flag set to 1.

Each child of the radio button is a checkbox that inherits its type and value from the parent radio button.

In addition to the attributes common to all fields, the radio button contains the following attributes:

Table 6.63 *Radio button attributes*

Key	Type	Semantics
FT	name	(Required, inheritable) Field type. Must be Btn .
Ff	integer	(Required, inheritable) Flags. The <i>push-button</i> flag should be 0. The <i>radio</i> flag should be 1. A radio button contains one additional flag: bit 15 “No toggle to Off” flag. If this bit is 1, then selecting the checkbox that is in the <i>on</i> state leaves it in that state, rather than toggling it to <i>off</i> ; in effect, the checkbox is “re-selected.” If this bit is 0, the checkbox that is currently in the <i>on</i> state can be set to <i>off</i> even if no other checkbox enters the <i>on</i> state.
Kids	array	(Required) An array of n checkboxes.
V	name	(Optional, inheritable) The value. V is the name of the key in the AP dictionary of whichever child checkbox is in the <i>on</i> state. The default value for this attribute is Off .

Example 6.16 *Radio button field with two buttons*

```

10 0 obj                                The radio button
<<
/FT /Btn
/Ff ...7                                ... radio flag = 1, push-button = 0 ...
/T (Credit Card)
/V /Master-Card
/Kids [11 0 R 12 0 R]
>>
endobj

11 0 obj                                checkbox 1
<<
/Parent 10 0 R
/AS /Master-Card
/AP
<</N << /Master-Card 8 0 R /Off 9 0 R >> >>
>>
endobj

```

```

12 0 obj                                checkbox 2
<<
/Parent 10 0 R
/AS /Off
/AP
<< /N << /Visa 8 0 R /Off 9 0 R >> >>
>>
endobj

8 0 obj                                An "On" stream
<< /Resources 20 0 R /Length 50 >>
stream
q 0 0 1 rg BT /ZaDb 12 Tf 0 0 Td (8) Tj ET Q
endstream
endobj

9 0 obj                                An "Off" stream
<< /Resources 20 0 R /Length 50 >>
stream
q 0 0 1 rg BT /ZaDb 12 Tf 0 0 Td (4) Tj ET Q
endstream
endobj

```

6.14.10 Choice

A choice is a field that contains n items, at most one of which may be selected as the field value. A choice may be presented to the user as a scrollable list within a rectangle on the page, or as a pop-up list triggered by user input. A pop-up list box may allow user input in addition to the predefined choice options, in which case it is referred to as a *combo box*. A choice can be used to represent the behavior of list boxes, combo boxes, or pop-up list boxes, depending on the setting of certain flags.

In addition to the attributes common to all fields, the field dictionary that represents a choice contains the following attributes:

Table 6.64 *Choice attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
FT	name	(<i>Required, inheritable</i>) Field type. Must be Ch .
Ff	integer	(<i>Optional, inheritable</i>) Flags. Collection of flags defining various characteristics of the list box field. The following flags apply specifically to choices: <ul style="list-style-type: none"> bit 18 <i>Pop-up</i> flag. Distinguishes between a list box (0) and a pop-up (1). bit 19 <i>Edit</i> flag. Indicates that this list box is a combo box and hence, the user may input a string value rather than selecting it from the pre-defined list.

bit 20	<i>Sort</i> flag. This flag is included for use by form-authoring tools that need to keep track of the user's preference regarding sort order of list items. If set, the authoring tool should sort the items in the Opt array in alphabetical order. This flag is <i>not</i> used by viewers. It is the responsibility of the authoring tool to put the items in the Opt array in the order desired. The viewer simply presents them in that order.	
Opt	array	<i>(Required, inheritable)</i> An array of n elements. Each element is either a Text representing one of the n potential values of the field, or an array containing a Text and a string: the text indicates a potential value, and the string is used to produce the appearance.
V	Text or name	<i>(Optional, inheritable)</i> Field value. Refers to the Opt array element that is currently selected. If that Opt array element is a text, then the value is that text. If it is an array, then the value is the text found in the first element of that array. The default value of this attribute is null .
TopIndex	integer	<i>(Optional, inheritable)</i> List boxes are scrollable. This integer indicates the index in the Opt array of the first visible element in the list box.



If the choice is presented to the user as a scrollable list, then the appearance contains the texts found in the **Opt** array, one per line, beginning with the text at the index indicated by **TopIndex**.

If the choice is presented to the user as a pop-up, then the appearance contains the string found in the **V** attribute.

Note When a pop-up choice is activated in the viewer by user input, the choices appear on the screen, typically immediately under the current value. The viewer may choose to present to the user this transient appearance of the field by generating an appropriate appearance stream, based on the **V** and **Opt** attributes.

Example 6.17 List box field

```
<<
/FT /Ch
/Ff ...
/T (Body Color)
/V (Blue)
/Opt [(Red) (My favorite color) (Blue)]
>>
```

6.14.11 Text field

A text field is a field whose value is an arbitrary, user-specified [Text](#). The text may be restricted to a single line or allowed to span multiple lines. The text is presented on the page in a single style; i.e. font, size, color, etc. The variable text used to generate the appearance of the field as described on [page 133](#) is found in the value of the text field.

In addition to the attributes common to all fields, the field dictionary that represents a text field contains the following attributes:

Table 6.65 *Text field attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
FT	name	<i>(Required, inheritable)</i> Field type. Must be Tx .
Ff	integer	<i>(Optional, inheritable)</i> Flags. Collection of flags defining various characteristics of the text field. Flags that apply to text fields include: <ul style="list-style-type: none"> bit 13 <i>Multi-line</i>. Distinguishes between a multi-line (1) and single-line (0) text field. bit 14 <i>Password</i>. Keyboard input is not displayed on the screen when this flag is 1. Rather, some benign feedback is given, such as displaying an asterisk for each input character. In order to preserve the confidentiality of password fields, the value of a text field with the password flag set should never be stored in the PDF file by a viewer.
V	Text	<i>(Optional, inheritable)</i> Field value.
MaxLen	integer	<i>(Optional, inheritable)</i> Maximum number of characters allowed in this text field.

Example 6.18 *Text field*

```

6 0 obj
<<
/FT /Tx
/Ff ... set multi-line flag ...
/T (Silly prose)
/DR 21 0 R
/DA (0 0 1 rg /Ti 12 Tf)
/V (The quick brown fox ate the lazy mouse)
/AP << /N 5 0 R >>
>>
endobj

5 0 obj
<< /Resources 21 0 R /Length 65 >>
stream
/Tx BMC BT 0 0 1 rg /Ti 12 Tf
1 0 0 1 100 100 Tm
0 0 Td (The quick brown fox ) Tj
0 -13 Td (ate the lazy mouse) Tj
ET EMC
endstream
endobj

```

6.14.12 Signature field**PDF 1.3**

A *Signature Field* is a field whose value is an electronic form of a signature. The signature may be purely mathematical, such as a public-private-key encrypted document digest, or it may be a biometric representation such as handwritten signature, fingerprint, or retinal scan. Filling in a Signature Field (signing it) involves updating at least the **V** key, and usually involves updating the **AP** key of the associated annotation. Exporting a Signature Field follows the rules for other fields and typically exports the **T**, **V**, and **AP** key-value pairs.

The digital signature public key specification is available online ([\[11\]](#)).

The field dictionary that represents a signature field contains the following attributes:

Table 6.66 *Signature field attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
FT	name	<i>(Required, inheritable)</i> Field type. Must be Sig .
T	string	<i>(Optional)</i> The partial field name.
Ff	integer	<i>(Optional, inheritable)</i> Flags.
Kids	array	<i>(Optional)</i> This array contains references to child field dictionaries of this field or to Widget annotation dictionaries that describe this field. As an optimization, the Widget annotation key-value pairs may be merged into the signature field dictionary.
Parent	dictionary	<i>(Required if this field is referenced from another field's Kids array)</i> The parent field.
V	dictionary	<i>(Optional, inheritable)</i> The value, a Signature Dictionary (see below).
DV	dictionary	<i>(Optional, inheritable)</i> The default value. The value of the field reverts to this when a ResetForm action occurs (see page 144). The format of DV is the same as V .

Signature Dictionary

The value of a Signature Field is a Signature Dictionary. It contains the following attributes:

Table 6.67 *Signature dictionary attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	<i>(Required)</i> Object type. Always Sig .

Filter	name	<i>(Required, inheritable)</i> Name of specific signing method, such as Adobe.PPKLite , Entrust.PPKEF , CICI.SignIt , and VeriSign.PPKVS . This name specifies which signature handler was used to fill in this signature field. Third-party handler writers are encouraged to register their name with Adobe; see Appendix F .
ByteRange	array	<i>(Required)</i> An array of pairs of integers, describing the exact byte range for the digest calculation: start byte-offset, length in bytes. Multiple discontinuous ranges of bytes may be used to describe a digest that does not include the signature token itself.
Contents	string	<i>(Required)</i> A (hex) string containing the encrypted signature token. A number of keys are optional, but all Adobe and 3rd-party signature handler writers are encouraged to use them in a standard way if they are used at all:
SubFilter	name	<i>(Optional)</i> Name of a specific submethod of the specified handler.
Name	Text	<i>(Optional)</i> String giving the name of the person or authority signing the document.
M	Date	<i>(Optional)</i> The time of signing. It is up to the specific signature handler whether this is a normal unverified computer time or a time generated in a verifiable way from a secure time server.
Location	Text	<i>(Optional)</i> String giving the CPU hostname or physical location of the signing.
Reason	Text	<i>(Optional)</i> String giving the reason for the signing. E.g., (I agree...) .

Finally, a Signature Dictionary may contain arbitrary handler-private key-value pairs. Third-party handler writers are encouraged to use a key-naming convention to avoid duplicate names. The suggested naming convention is to use the registered filter name followed by a period as a prefix of all private keys.

The **DV** key in a Signature Field may specify a signature dictionary for an unsigned field that is pre-loaded with default values for some keys. There may be private-key default values for multiple handlers. Signature handlers are free to use or ignore these values in constructing the **V** key for a signed field. A signature dictionary for a signed document should contain keys meaningful only to the handler that was actually used in the signing; all others should be stripped out by that handler.

Signature Annotation

Each Signature Field is described by one or more **Widget** annotations, just as other field types are described. This annotation should be included in the **Annots** array of the page containing the signature. The **Rect** (rectangle) specified by this annotation gives the position of the field on some page, and the **AP** (appearance) specified by this annotation gives the visible representation of the signature field. A rectangle of zero height and zero width is explicitly allowed for signatures (such as **PPK** ones) that are not intended to be visible.

Signature Appearance

For a signature field that has not been filled in, the normal appearance is usually blank. For a signature that has been filled in, the normal appearance is a stream that shows an appropriate representation of the unvalidated signature: text, strokes of a handwritten signature, a bitmap of a fingerprint, or whatever other representation is deemed appropriate by the signing method. When a document containing a signature field is displayed or printed, the normal appearance is used, just as it is for other field types. This behavior is independent of the PDF viewer having any signature-related plug-ins.

If a viewer does have a signature plug-in that matches the specific signing method used in a given signature field, and the plug-in is requested to validate the signature, it is expected to alter the appearance of the signature from *unvalidated* to either *valid* or *invalid*. It is suggested that the standard appearance of an unvalidated signature contain an overprinted yellow question mark, the appearance of an invalid signature contain an overprinted red X, and the appearance of a valid signature contain the logo of the validating method underprinted as a watermark.

6.14.13 AcroForm actions

AcroForms support four additional types of actions: **SubmitForm**, **ResetForm**, **ImportData**, and **JavaScript**.

6.14.13.1 SubmitForm Action

The **SubmitForm** action is used to send name-value pairs from the selected fields to the indicated URL, presumably that of a Web server where they will be processed, and from which a response sent back. The name that is sent is the fully qualified name of the field, and the value that is sent is the value of the **V** entry in the field dictionary (or the value of the **AP** entry if the field is of type push button and the data is sent using Forms Data Format (FDF)).

These name-value pairs can be sent using Forms Data Format (see [Appendix H](#)) or HTML Form format, as specified in RFC 1866, *Hypertext Markup Language - 2.0* (i.e. application/x-www-form-urlencoded).

Table 6.68 *SubmitForm* action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(<i>Required</i>) Action type. Always SubmitForm .
F	File specification	(<i>Required</i>) A URL file specification; see Section 7.4.4 on page 194 . This is the URL of the script at the Web server that will process the request.
Fields	array	(<i>Optional</i>) The name-field value pairs to send, or to exclude from sending. Which of these two interpretations to use is decided based on bit 1 of the Flags key below. Each entry in the Fields array is either an indirect object reference to the corresponding field dictionary in the AcroForm, or (in PDF 1.3) the string representing the fully qualified name of the field. The array may contain a mixture of these types. If the Fields key is not present, then <i>all</i> name-field value pairs in



the document are sent (except for those that have the “no-export” flag set), and bit 1 of **Flags** is ignored. The actual selection of name-value pairs to send can be further refined through bit 2 of the **Flags** key.

If the **Fields** array represents fields to include, and an entry in this array represents a non-terminal field, that is, a field that has descendants (e.g. **Kids**) and those descendants can have values, then all name-value pairs from terminal fields in the subtree underneath it get sent (except for those that have the “no-export” flag set).

If (and only if) the Submit Form action includes the Fields attribute, then any fields of type push button that end up in the set to transmit have their **AP** entry exported.

Flags

- integer *(Optional)* The binary value of the integer is interpreted as a collection of flags that define various characteristics of the action. The default value is 0.
- bit 1 *Include/exclude* flag. This determines how the **Fields** key is interpreted. If this bit is 0, then **Fields** represents the individual fields to send. If the bit is 1, then all name-value pairs in the AcroForm are sent, except for those in the **Fields** array (and those that are flagged as “no-export”).
 - bit 2 *“Include no-value fields”* flag. This can further restrict which name-value pairs get sent. If this bit is 0, then those fields (previously selected by the combination of the **Fields** array and bit 1 of **Flags**) that don’t have a value (i.e. they don’t have a **V** key) is *not* sent. If, on the other hand, this bit is 1, then for any fields that don’t have a **V** key only the field name is sent.
 - bit 3 *Export format* flag. If this bit is 0, then the data is sent using the Forms Data Format (FDF). Otherwise it is sent using the HTML Form format.
 - bit 4 *Get Method* flag. If this bit is 0, the data is submitted using a POST action; if it is 1, a GET action is used for the submission. (This bit has an effect only when bit 3 is 1. If bit 3 is 0, this bit must also be 0.)
 - bit 5 *Submit Coordinates* flag. If this bit is 1, the x and y coordinates of the mouse click that caused this action are sent as part of the AcroForm data. The x and y values are relative to the upper left hand corner of the widget contains the submit action. They are represented in the data as *name.x=xval&name.y=yval* where *name* is the value of the **TM** field if present, otherwise the field name. The special case of **TM** equal to a single space character results in *x=xval&y=yval* being submitted (no name, no dot). (This bit has an effect only when bit 3 is 1. If bit 3 is 0, this bit must also be 0.)
-

6.14.13.2 ResetForm Action

This action is used to reset the indicated fields to their default values (if defined); that is, the value of the **V** key is set to the value of the **DV** key in the field dictionary. If no default value is defined for a field, then the **V** key is removed. Resetting a field that can have no **V** key (i.e., a push button) has no effect.

Table 6.69 *ResetForm* action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(<i>Required</i>) Action type. Always ResetForm .
Fields	array	(<i>Optional</i>) The fields to reset, or to exclude from resetting. Which of these two interpretations to use is decided based on bit 1 of the Flags key below. Each entry in the Fields array is either an indirect object reference to the corresponding field dictionary in the AcroForm, or (in PDF 1.3) the string representing the fully qualified name of the field. The array may contain a mixture of these types. If the Fields key is not present, then <i>all</i> fields in the document are reset, and bit 1 of Flags is ignored. If the Fields array represents fields to include, and an entry in this array represents a non-terminal field, that is, a field that has descendants (e.g. Kids) and those descendants can have values, then all fields in the subtree underneath it get reset.
Flags	integer	(<i>Optional</i>) The binary value of the integer is interpreted as a collection of flags that define various characteristics of the action. The default is 0. bit 1 <i>Include/exclude</i> flag. This determines how the Fields key is interpreted. If this bit is 0, then Fields represents the individual fields to reset. If the bit is 1, then all fields in the AcroForm are reset <i>except</i> for those in the Fields array.

PDF 1.3

6.14.13.3 ImportData Action

Import data in Forms Data Format from a specified file into the AcroForm.

Table 6.70 *ImportData* action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(<i>Required</i>) Action type. Always ImportData .
F (File) File specification		(<i>Required</i>) File specification of the FDF file to import.

**6.14.13.4 JavaScript Action**

The JavaScript action causes a script to be compiled and executed by the JavaScript interpreter. Depending on the nature of the script, this can cause various AcroForm fields in the document to update their values or change their graphical properties.

Table 6.71 *JavaScript* action attributes (in addition to those in [Table 6.36](#))

Key	Type	Semantics
S (Subtype)	name	(<i>Required</i>) Action subtype. Always JavaScript .

PDF 1.3

JS (JavaScript)

string or stream *(Required)* The string or the contents of the stream contains the JavaScript to be executed.

One of the desirable properties of a scripting language that supports function definitions is the re-use of code in the form of function calls. In some instances, a Acrobat author will want to define a JavaScript that is very specific to a single field. For the most part, however, field scripts will make parameterized calls into subroutines that will do the bulk of the processing.

A PDF document allows for function definitions that are particular to a document and used by field scripts to be stored at the document level in the Names tree (see [Section 6.9, “Names dictionary”](#)). All of the scripts specified in the name tree are executed when the document first opens thus defining functions for use by other scripts in the document.

Note *PDFDoc or Unicode encodings (identified by a prefix pair of <FE FF> bytes) are used to encode the contents of the string or stream.*

Implementation note *These encodings are translated to a platform specific encoding prior to interpretation by the JavaScript engine as the current version of JavaScript (1.2) is not Unicode compatible.*

Implementation note *See the Netscape JavaScript Reference Manual [33] and the Adobe Acrobat Forms JavaScript Object Specification [10] for details on the contents and effects of these types of scripts.*

6.14.14 Named Pages**PDF 1.3**

In PDF 1.3, it is possible to name a page using the **Names** data structures described in [Section 6.9 on page 119](#). Naming a page allows it to be referenced in two different ways.

- During an FDF import operation, a referenced page can be added to the document into which the FDF is being imported, either as a button appearance, or as a page.
- During the execution of a JavaScript, the named page can be added as a regular page to the current document.

Pages can be named in one of two ways: If a named page is to be visible in the viewer, the page is left in the Pages data structure (described in [Section 6.3 on page 71](#)), and a reference to the page is placed in the appropriate leaf node of the Pages tree of the **Names** dictionary. If the page is to be made invisible to the viewer, it is instead removed from the Pages tree and put in the Templates structure. Additionally, the **Type** field must be changed to **Template**, the **Parent** key must be removed, and the **B** key must be removed.

Regardless of how the page is named, when it is copied into a document either by using a JavaScript or by importing an FDF, the new copy is not itself named.

6.15 Sounds

PDF 1.2

A PDF document may contain Sound objects, and it may refer to external sound files.

6.15.1 Sound objects

A Sound is represented as a stream. In addition to the standard keys in any stream dictionary, a Sound dictionary is defined with the following keys:

Table 6.72 *Sound attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	<i>(Required)</i> Stream type. Always Sound .
R (Rate)	number	<i>(Required)</i> Sample rate, in samples per second.
C (Channels)	integer	<i>(Optional)</i> Number of sound channels in the data. The default is 1.
B (Bits)	integer	<i>(Optional)</i> Number of bits per sample value (per channel). The default is 8.
E (Encoding)	name	<p><i>(Optional)</i> The format of the sample data. Allowable values are:</p> <p>Raw unspecified or unsigned values, 0 to 2B (Bits) - 1. Signed two's-complement values. muLaw μ-law encoded samples. ALaw A-law encoded samples. The default is Raw.</p>
CO (Compression)	name	<i>(Optional)</i> The type of sound compression (not file compression) used on the data, if any. If this key is omitted, then no sound compression has been used, and the data contains sampled waveforms, to be played at R (Rate) samples per second, per channel.
CP (Compression Parameters)	<i>various</i>	<i>(Optional)</i> Optional parameters, specific to the Compression used.

Dual-channel uncompressed sound is saved in an interleaved format, in channel order (1, 2). For 2-channel stereophonic sounds, channel 1 is the left channel, and channel 2 is the right channel. Acrobat does not support more than two channels of sound.

Samples are stored in the stream with the most significant bits first (big-endian order for samples larger than 8 bits). Samples that are not a multiple of 8 bits are to be packed into the byte stream, starting at the most significant end of the byte and proceeding to the least significant end. If a sample extends across a byte boundary, the most significant bits are placed in the first byte, followed by less significant bits in the subsequent bytes.

To maximize portability of PDF documents containing embedded sounds, Adobe recommends that PDF applications and plug-ins support at least the following formats (assuming the platform has sufficient hardware and OS support to play sounds at all):

B	8 or 16 bits per channel
C	1 or 2 channels
E	Raw, Signed, or muLaw encoding
R	8000, 11025, or 22050 samples per second

If the encoding is **muLaw**, **B** must be 8, **C** must be 1, and **R** must be 8000. If the encoding is **Raw** or **Signed**, **R** must be 11025 or 22050. Sound players should be prepared to convert between formats, downsample rates, and combine channels as necessary to render sound on the target platform.

6.15.2 External sounds

If a Sound stream is external (that is, if it includes an **F** key), the specified file must be a *self-describing* sound file, containing all the information necessary to render the sound. No additional information need be present in the PDF file.

Note *AIFF, AIFF-C (Macintosh), RIFF (.wav) and snd (.au) files all contain sufficient information for playing.*

6.16 Web Capture

PDF 1.3

Web Capture (sometimes referred to as AcroSpider) is an Acrobat plug-in that creates PDF files from Internet-based or locally resident HTML, PDF, GIF, JPEG, and ASCII text files. The plug-in provides several useful features, including the ability to automatically print multiple pages from a website and the ability to locally save and preserve the visual appearance of information from the web. The plug-in also allows the user to update and modify PDF files that it created, to retrieve more material or delete unwanted material.

AcroSpider performs the following operations, among others:

- It finds all pages or Image XObjects in a PDF file that were generated from a given URL.
- It finds all pages or Image XObjects in a PDF file that match a given digital ID (MD5 hash).
- Given a page in the PDF file, it find out its creation information, such as the URL (if any) from which the page was generated.
- It lists all URLs that were used to generate pages in the PDF file.

Web Capture records the information it needs in the PDF file, in two data structures. The first, known as the *Web Capture Content Database*, keeps track of the material retrieved and where it came from, enabling Web Capture to avoid

needlessly downloading material that is already in the file. The second, known as the *Web Capture Info Dictionary*, holds information about the state of the PDF file, and how it was created. The following sections provide a detailed overview of these structures. To understand the need for these data structures, it is important to know how Web Capture works. (The discussion below is centered on HTML and GIF files, although Web Capture handles other types as well.)

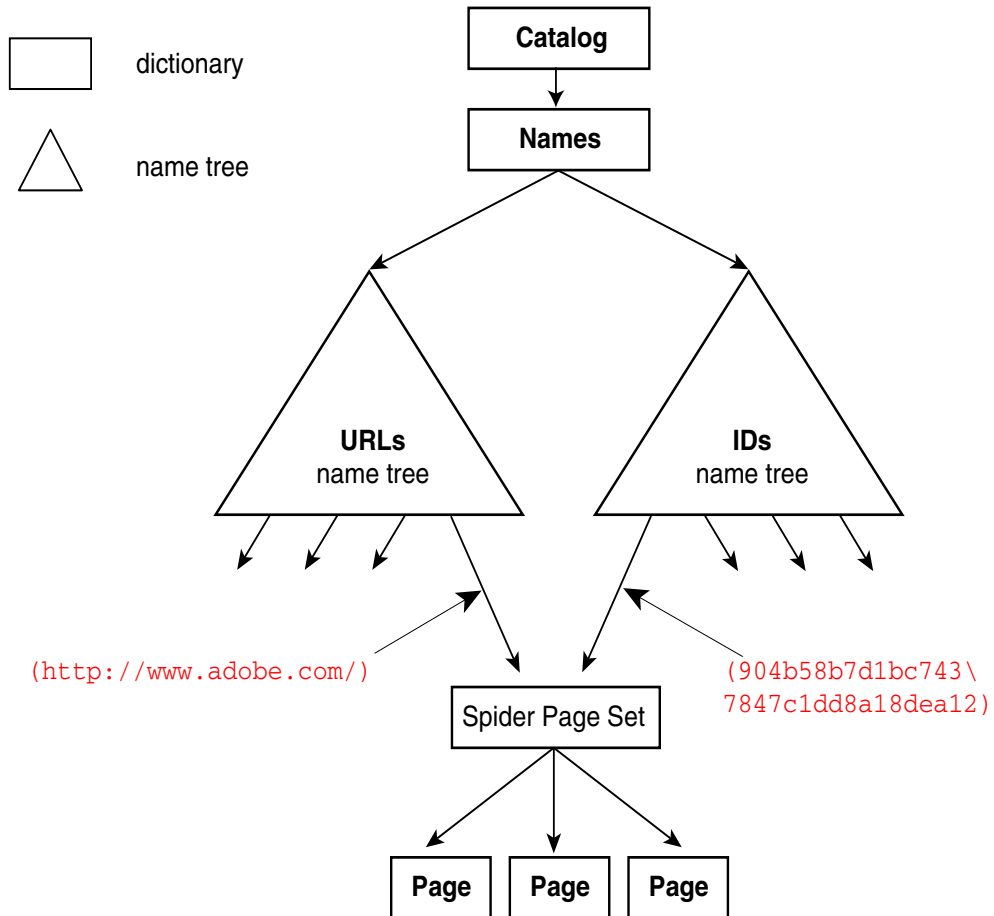
6.16.1 Web Capture Content Database

Web Capture retrieves HTML files from URLs and converts them into PDF. Since HTML pages do not have a fixed size, the conversion of HTML pages to PDF frequently results in a one-to-many mapping: a single HTML page may result in several letter-sized PDF pages. Additionally, a single Web Capture-generated PDF file may contain the contents of multiple URLs. Web Capture keeps track of all of the URLs used to create the resulting PDF file, by means of the Web Capture Content Database. This database maps URLs and digital IDs to PDF objects such as pages and Image XObjects. This allows Web Capture to perform on-the-fly optimizations, such as storing only one copy of an image that is referenced by multiple HTML pages. By looking up digital IDs in the database, Web Capture determines whether newly downloaded content is identical to content already in the PDF that was retrieved from a different URL.

[Figure 6.12](#) depicts a sample portion of a Web Capture Content Database. There are two Web Capture-related name trees in a PDF document's Catalog (see [page 119](#)); one contains URLs, and the other contains digital IDs. Each name in these trees is an indirect reference to one or more *Web Capture Content Sets*.

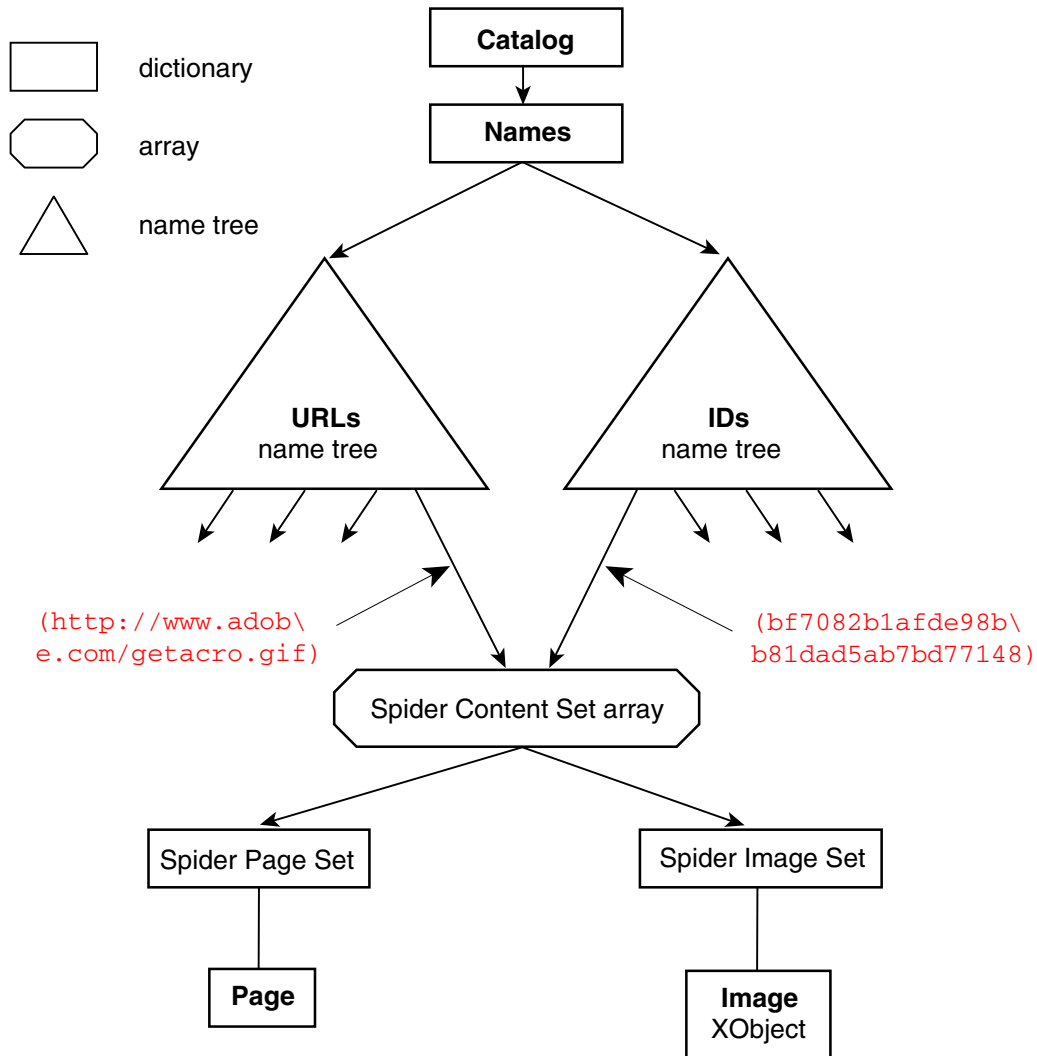
A Web Capture Content Set is a dictionary that holds information about a group of PDF objects generated from the same source data. This information includes indirect references to the objects. There are two subtypes of Web Capture Content Set: *Web Capture Page Set* and *Web Capture Image Set*. When Web Capture converts an HTML file into PDF pages, for example, it creates a Web Capture Page Set that holds information about all of the new pages. Similarly, when it converts a GIF image into one or more Image XObjects, it create a Web Capture Image Set to hold information about the new XObjects.

In the example in [Figure 6.12](#), an HTML file has been retrieved from the Web. It has been converted into three pages in the PDF file. A Page Set has been created for the new pages. It has indirect references to the pages themselves. The source URL, the string (`http://www.adobe.com/`), has been added to the **URLS** name tree and its tree entry points to the new Page Set. The digital ID calculated from the HTML source, the string (`904b58b7...`), has been added to the **IDS** name tree in a similar fashion.

Figure 6.12 *Simple Web Capture file structure*

A more complicated situation is depicted by [Figure 6.13](#). Here, a GIF has been retrieved from a URL, `(http://www.adobe.com/getacro.gif)`, independently of any HTML file, and converted into a single page in a PDF file. As in [Figure 6.12](#), a Web Capture Page Set has been created to hold information about the new page. However, since the retrieval also resulted in a new Image XObject, a new Web Capture Image Set has also been created. In this case, the same entries in the name trees must point to both a Web Capture Image Set and a Web Capture Page Set. Thus, a Web Capture Content Set *array* is created to hold these pointers, and the name tree entries point to this array.

Figure 6.13 *Complex Web Capture file structure*



URL strings

The entries in a name tree are stored in sorted order, to facilitate fast lookup. This requires that the URLs used as keys in the **URLS** name tree be stored in a canonical form. URLs being looked up in this name tree must also be transformed to the canonical form before searching.

Here is the algorithm to canonicalize a URL (relevant for HTTP, FTP, and file URLs). These steps refer to terminology from *RFC 1738: Uniform Resource Locators* [18] and *RFC 1808: Relative Uniform Resource Locators* [23].

- If the URL in question is relative, make it absolute.
- If the URL contains a number sign character (#), strip the leftmost # and any characters after it.

- Transform the scheme section of the URL into lowercase ASCII.
- If the URL contains a host section, transform the host section into lowercase ASCII.
- If the scheme is **file**, and the host is **localhost** (case insensitive), strip the host section.
- If the URL contains a port section, and the port is the default port for the given protocol (80 for HTTP, 21 for FTP), strip the port section.
- If the URL's path section contains “. . .” or “. ” subsequences, transform the path as described in section 4 of RFC 1808 [23].

Note Concerning the encoding of characters in a URL which are considered unsafe: no number of encoding or decoding passes on a URL will ever cause it to reach a steady state, because the percent character (%) is both an unsafe character and the escape character for encoded characters. In general, it is not possible to tell the difference between a URL with non-encoded characters and a URL with encoded characters. For example, it is impossible to decide whether the sequence **%00** represents a single encoded NUL character, or a non-encoded sequence of three characters. Empirically, URLs embedded in HTML files have unsafe characters encoded with one pass of encoding, and web servers perform one decoding pass on received paths (though CGI scripts are free to make their own decisions). Thus, canonical URLs are assumed to have one and only one pass of encoding. If the initial encoding state of a URL is known, then it can be safely transformed into a URL which has undergone only one pass of encoding.

6.16.2 Web Capture Content Sets

A Web Capture Content Set is a dictionary that holds information about a set of PDF objects generated from the same source data. This information includes an indirect reference to each of the objects in the set. A Content Set can be used to hold information that is common to all of the objects as well as information that describes the set of objects itself. Every Content Set must contain an **S** (Subtype) key. The subtype determines what type of objects are grouped by the Content Set (e.g., page objects, images). It is legal for a subtype to be defined that groups heterogeneous object types. Unless a subtype definition imposes stricter constraints, it is legal for the same object to be a member of multiple Content Sets at once.

Table 6.73 *Web Capture Content Set attributes (common to all Web Capture Content Sets)*

Key	Type	Semantics
Type	name	(Optional) Object type. Always SpiderContentSet .
S (Subtype)	name	(Required) Web Capture Content Set subtype.
O (Objects)	array	(Required) List of indirect references to the objects grouped by this Content Set. The order of objects in the array is undefined, in general, but subtypes of Web Capture Content Set may define an order.

ID	string	<i>(Required)</i> An ID for the Content Set.
TS (TimeStamp)	Date	<i>(Optional)</i> The date at which the Content Set was created.
CT (Content Type)	string	<i>(Optional)</i> A Content Type that describes the entity from which the objects grouped by the Content Set were created.

If a Content Set is located by traversing the **URLS** name tree, the ID key allows its related entry in the **IDS** name tree to be found. Every Content Set must be mapped into the **IDS** name tree. The ID may also be used as a unique (for practical use) identifier for the Content Set. The ID calculation method is specified on a per-subtype basis.

The optional **CT** (Content Type) key can be used to store origin information about the objects mapped by the Content Set. The Content Type string should conform to the Content Type specification described in *RFC 2045: Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies* [26]. For example, for a Web Capture Content Set (Web Capture Page Set) that maps a group of Page objects created from an HTML file, the value of the **CT** key would be the string (`text/html`).

6.16.2.1 Web Capture Page Set (subtype of Web Capture Content Set)

A *Web Capture Page Set* is a Web Capture Content Set subtype for groups of PDF page objects that are generated from the same source data, such as an HTML file. A Page Set can optionally store information that applies to the entire set of pages, such as a title. A page object may not be a member of more than one Page Set.

Table 6.74 *Web Capture Page Set attributes (in addition to those in Table 6.73)*

Key	Type	Semantics
S (Subtype)	name	<i>(Required)</i> Web Capture Content Set type. Always SPS (Spider Page Set).
SI (Source Info)	dictionary or array	<i>(Required)</i> One or more Source Info objects. See Section 6.16.2.2, “Source Info Object.”
T (Title)	Text	<i>(Optional)</i> A string intended to be displayed as the Page Set’s title.
TID (Text ID)	string	<i>(Optional)</i> An ID generated from the text of the Page Set.

For each Web Capture Page Set it is necessary to store information related to the location of the source data from which the pages were generated. This information is stored in a *Source Info* object (See [Section 6.16.2.2, “Source Info Object”](#) below). The **SI** (Source Info) key allows one or more Source Info objects to be associated with a Page Set. If only one such object is present, the value of the **SI** key may be the object itself, rather than a single-element array.

Multiple Source Info objects may be associated with a Page Set. Each holds information about a location from which the same source-data bytes were retrieved. It is not unusual for the exact same source bytes to be referenced by two or more completely unrelated URLs. If Web Capture detects such a condition (via ID comparison), only one copy of the PDF pages created from this source-data will be placed in the file, and only one Page Set will be created for the pages. Information specific to each source will be recorded in a Source Info object.

The optional **TID** (Text ID) key may be used to store an ID generated from the text of the pages grouped by the Page Set. This ID may be used, for example, to determine whether the text of a document has changed. For calculation info, see [page 160](#). A Text ID may not be appropriate for some Page Sets (such as ones with no text), and should be omitted in these cases.

For a Page Set, the order of the indirect references to page objects in the Objects array is defined to be the same as the order in which the page objects were initially added to the file.

6.16.2.2 Source Info Object

A Source Info object contains information specific to a single location from which the source data for a Web Capture Content Set was retrieved. It is possible for a Content Set to have multiple Source Info objects associated with it.

Table 6.75 *Source Info attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
AU (Aliased URLs object) string or dictionary		<i>(Required)</i> A URL (string) or an Aliased URLs object (dictionary) for this Source Info object.
TS (TimeStamp)	Date	<i>(Optional)</i> The date at which the Content Set's data was last found to be up to date with the source referenced by this Source Info object.
E (Expires)	Date	<i>(Optional)</i> The date at which the Content Set's data should be considered out of date with the source referenced by this Source Info object.
S (Submitted)	integer	<i>(Optional)</i> Indicates that the location that this Source Info object describes was accessed via a form submission, and if so, the type of form submission. 0 Not accessed via a form submission. This is the default value. 1 Accessed via an <i>HTTP GET</i> submission. 2 Accessed via an <i>HTTP POST</i> submission.
C (Command)	dictionary	<i>(Optional)</i> Web Capture Command object whose execution caused the retrieval of the source-data referenced by this Source Info object.

The value of the **AU** (Aliased URLs) key is either a single URL, or an Aliased URLs object (see [Section 6.16.2.3](#)) representing all of the known URLs that map (through redirection) to the location described by this Source Info object. If there is only one such URL, an Aliased URLs object is unnecessary; a string storing the URL is sufficient. For file size efficiency, it is recommended that the entire Aliased URLs object (excluding the URL strings) be represented as a direct object, as its internal structure should never be shared or externally referenced.

The **TS** (TimeStamp) key allows each source location associated with a Content Set to have its own TimeStamp. This is necessary because the **TS** (TimeStamp) in the Content Set itself merely refers to the creation date of the Content Set object. A hypothetical “Update Web Capture Content Set” command might change the Source Info TimeStamp to the current time if it found the source data not to have changed since the last time the TimeStamp was set.

The **E** (Expires) key allows each source-data location associated with a Content Set to have an expiration date. If the current date is later than the expiration date, the data mapped by the Content Set should be considered out of sync with the original source.

The **S** (Submitted) key should only be present in Source Info objects which are referenced by Web Capture Page Sets. The flag indicates whether or not the source retrieved from the location described by this Source Info object was retrieved as the result of a form submission.

The **C** (Command) key should be present only in Source Info objects that are referenced by Web Capture Page Sets. If present, the value of the key is an indirect reference to the Web Capture Command object (see [Section 6.16.4](#)) whose execution caused the retrieval of the source-data referenced by this Source Info object, thus causing this Source Info object to be created. This allows us to take a PDF page and find out which Web Capture Command generated it.

6.16.2.3 Aliased URLs Object

When a URL is accessed by a user-agent via HTTP, a response header is frequently returned that indicates that a different URL should be accessed to reach the desired data. The agent will then access the new URL. This redirection process can continue indefinitely (though in practice agents should place an upper limit on the number of URLs that will be accessed). It is not uncommon to find multiple URLs that eventually lead to the same final URL through one or more redirections. An *Aliased URLs Object* represents the set of chains of URLs that lead to a common URL.

Table 6.76 *Aliased URLs Object attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
URL	string	<i>(Required)</i> The common (final) URL at which all chains in this object terminate.
C (Chains)	array	<i>(Optional)</i> An array that holds one or more arrays of strings (URLs).

The **C** (Chains) key should be omitted if the Aliased URLs Object contains only one URL. If the **C** key is present, its value is an array of arrays. Each nested array represents a chain of URLs. The URLs are stored as strings. The order of the URLs in an array should be the same as the order of the URLs in a chain. The common URL (the last URL in a chain) may be omitted, since it is already stored in the **URL** key.

Implementation note An implementation should use an indirect object reference to a shared URL string for all URLs in an Aliased URLs Object. These strings may then be shared with other structures, and among the chains of the Aliased URLs Object.

6.16.2.4 Web Capture Image Set (subtype of Web Capture Content Set)

A *Web Capture Image Set* is a Web Capture Content Set subtype for groups of Image XObjects that were generated from the same source (multiple frames of an animated GIF, for example). In the usual case, only a single Image XObject will be generated. An Image XObject may not be a member of multiple Web Capture Image Sets.

Table 6.77 *Web Capture Image Set attributes (in addition to those in Table 6.73)*

Key	Type	Semantics
S (Subtype)	name	(Required) Web Capture Content Set type. Always SIS (Spider Image Set).
R (Refcounts)	integer or array	(Required) Reference count for each Image XObject grouped by the Web Capture Image Set.
SI (Source Info)	dictionary or array	(Required) One or more Source Info objects.

The **R** (Refcounts) key associates a reference count (an integer) with the indirect reference to an Image XObject in the same position in the Objects array. That is, it is indexed in parallel with the Objects array. If there is only one Image XObject mapped by the Web Capture Image Set, the value of the **R** key may be the integer itself, rather than a single-element array.

When Web Capture creates a page that references an Image XObject that is mapped by a Web Capture Image Set, the XObject's reference count is incremented (once per page). When a page that references the XObject is copied inside the same file, the reference count must be incremented as well. When a page is deleted, the reference count for each Image XObject to which it referred is decremented. When a reference count reaches zero, it is assumed that there are no remaining pages that reference the XObject, and that the indirect reference to it may be removed from the Web Capture Image Set.

Implementation note In Acrobat 4, if this reference is not removed, the XObject will never be garbage-collected during a Save operation. Thus, the Image Set's reference to the XObject may be considered a weak one that is relevant only for caching purposes: when the last strong reference goes away, so does the weak one.

The **SI** (Source Info) key allows one or more Source Info objects to be associated with a Web Capture Image Set. If only one such object is present, the value of the **SI** key may be the object itself, rather than a single-element array. Each Source Info object stores information related to the location of the source from which the images were generated.

Note Unlike the *Refcounts* array, the *Source Info* array is not indexed in parallel with the *Objects* array, or any other array.

Multiple Source Info objects may be associated with a Web Capture Image Set. Each holds information about a location from which the same source data was retrieved. It is not uncommon for the exact same data to be referenced by two or more unrelated URLs. Web Capture can detect this condition (via ID comparison), and place only one copy of the Image XObjects generated from the data into the file. In this case, only one Web Capture Image Set will be created, but it will have multiple Source Info objects associated with it.

6.16.3 Web Capture Info Dictionary

A *Web Capture Info dictionary* may be added to a document's Catalog. This provides a place for Web Capture to store information regarding the state of the file.

Table 6.78 *Web Capture Info attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
V (Version)	number	<i>(Required)</i> Version number of the Web Capture structures. For PDF 1.3, this version number is 1.0.
C (Commands)	array	<i>(Optional)</i> An array of indirect references to one or more Web Capture Command objects.

As Web Capture builds a PDF file in response to user commands, it can save the information necessary to replay these commands as Web Capture Command objects (see the next section). The **C** key groups one or more of these objects. If the **C** key is defined, its value is an array that holds indirect references to Command objects. The Commands in the array appear in the order in which they were created.

6.16.4 Web Capture Command

A *Web Capture Command* is a dictionary that encapsulates the information that Web Capture uses to execute a command that adds pages to a PDF file.

Table 6.79 *Web Capture Command attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
URL	string	<i>(Required)</i> Initial URL for this command.

L (Levels)	integer	<i>(Optional)</i> Number of levels to retrieve from the initial URL.
F (Flags)	integer	<i>(Optional)</i> Collection of flags. The default value is 0.
	bit 1	<i>Same-Site</i> flag, indicating that the traversal should not include pages that are not on the same host.
	bit 2	<i>Same-Path</i> flag, indicating that the traversal should not include pages that are not found under the same path (see below).
	bit 3	<i>Submit</i> flag, indicating that the Command represents a form submission (see below).
P (Post Data)	string or stream	<i>(Optional)</i> Data that was posted to the URL.
CT (Content Type)	string	<i>(Optional)</i> A Content Type that describes the data that was posted to the URL.
H (Headers)	string	<i>(Optional)</i> Additional HTTP request headers sent to the URL.
S (Settings)	dictionary	<i>(Optional)</i> A Web Capture Command Settings object.

The value of the **URL** key is the initial URL to access for the command. The **L** (Levels) key specifies the number of levels of pages to retrieve from the initial URL. If the **L** key is omitted, its value is assumed to be 1, which specifies the retrieval of the initial URL only.

The *Same-Path* flag requires further explanation. (The following description borrows terminology from RFC 1808 [23].) Each time a URL is considered for traversal, it is compared to the initial URL. First, the scheme and net_loc components of the two URLs must match (case insensitive). Finally, the URL in question is considered to be in the same path as the initial URL if the path component of the initial URL is identical (case sensitive) to the path component of the URL in question, up to and including the last forward slash character in the initial URL. For example,

`http://www.adobe.com/foo/bar/bleh.html`

is considered to be in the same directory as

`http://www.adobe.com/foo/initial.html`

because **`foo/`** is present at the start of the path components for both URLs.

If the *Submit* flag is set, then the Web Capture Command represents a form submission. If the **P** (Post Data) key is not present, then the submitted data is encoded in the URL (an *HTTP GET* query). If the **P** key is present, then the Web Capture Command represents an *HTTP POST* form submission. In this case, the value of the *Submit* flag is ignored. If the Post Data is small enough, a string may be used to represent it. For large amounts of data, a stream is recommended, as it can offer compression.

The **CT** (Content Type) key is relevant only in a POST situation. It describes the Content Type (format described in RFC 2045 [26]) of the Post Data stored under the **P** key. If the **CT** key is omitted, the Content Type is assumed to be **application/x-www-form-urlencoded**.

The **H** (Headers) key specifies additional HTTP request headers that were sent in the request for the URL. Each header line in the string is terminated with a carriage return and a newline. Example:

```
(Referer: http://foobar.com\015\012From: foo@bar.com\015\012)
```

The HTTP request header format is specified in *RFC 2068: Hypertext Transfer Protocol – HTTP/1.1* [24].

The **S** (Settings) key references a Web Capture Command Settings dictionary (see the next section). This object holds settings specific to the conversion engines. If this key is omitted, default values are assumed. It is recommended that the Command Settings objects be shared by any Command objects that use the same settings.

6.16.4.1 Web Capture Command Settings

A *Web Capture Command Settings* dictionary holds conversion-engine-specific settings that are needed in the conversion process. Web Capture Command Settings objects are normally referenced by Web Capture Command objects.

It must be possible to deeply copy a Web Capture Command Settings dictionary without explicit knowledge of the settings it may contain. To facilitate this operation, a restriction is placed on the directed graph of PDF objects rooted by the Web Capture Command Settings dictionary. The graph must not contain indirect object references to objects that are not solely contained by the graph. Put another way, if the PDF file's graph is traversed as a tree starting from the Catalog, each object that has a Web Capture Command Settings dictionary as an ancestor must not appear on any path in which it does not have the same Web Capture Command Settings as an ancestor.

Table 6.80 *Web Capture Command Settings attributes*

Key	Type	Semantics
G (Global Settings)	dictionary	(Optional) Global conversion engine settings, relevant to all conversions. If this key is omitted, default settings will be used.
C (Converter Settings)	dictionary	(Optional) Settings for specific Content Type conversion engines (e.g., HTML). Each key in the dictionary is the internal name of a conversion engine. The value associated with the key is a dictionary containing the settings information associated with that conversion engine. If the settings for a particular conversion engine are not found in the dictionary, default settings will be used.

The internal name of a conversion engine should be a name object of the following form (case sensitive):

```
<converterName> :=
    /<company>:<product>:<version>:<contentType>
```

company is the name (or abbreviation) of the company that created the conversion engine. *product* is the name of the conversion engine; this field may be left blank, but the trailing colon character is still required. *version* is the version of the conversion engine. *contentType* is an identifier for the content type that the settings are associated with. This is required because some converters may handle multiple content-types.

Example: **/ADBE:H2PDF:1.0:HTML**

6.16.4.2 Page and Image XObject Attributes

A Web Capture Content Set (Page Set or Image Set) points to all the objects (pages or images) that are derived from its source. It is referred to as the *Parent Content Set* of those objects. An object has at most one such parent. There is no direct pointer in the opposite direction, from a “child” object to its Parent Content Set. Such a pointer might present problems for an application that traces all pointers from an object to determine what resources, for example, the object depends on. Instead, a child object contains an **ID** string, and the document’s **IDS** name tree maps that string to the Parent Content Set, or to an array containing the Parent. (The Parent is that element of the array whose Objects list contains the child object.)

In the course of creating PDF pages from HTML files, Web Capture frequently scales the page contents downward to fit on fixed-sized pages. The **PZ** (Preferred Zoom) key specifies a scaling factor by which the page can be scaled to “undo” the scaling and view the page at its original size. That is, when the page is viewed at the preferred zoom, one unit in default user space will correspond to one pre-scaled source pixel.

6.16.4.3 Implementation details

Web Capture Page Set ID

The ID for a Web Capture Page Set is calculated as follows. First, the source-data for the Web Capture Page Set is fed into the MD5 algorithm described in *RFC 1321: The MD5 Message-Digest Algorithm* [35]. For an HTML file, the source-data is the HTML source itself, not any auxiliary data such as referenced images. Next, the strings representing the IDs of any auxiliary data files are fed into the MD5 algorithm. It is important that the IDs for the auxiliary data files be fed into the MD5 algorithm in a deterministic order. This order is defined to be the same as their order of use in the source-data. An auxiliary data file that is referenced more than once must have its ID fed only once to the MD5 algorithm, when it is first referenced by the source-data.

The calculation process combines the IDs of the auxiliary data with the ID of the source-data to create a composite ID that identifies the visual appearance of the pages in the Page Set. For two HTML source files that are identical, but for which the referenced images contain different data (they may have been generated by a script, or they may be pointed to by relative URLs), the IDs will not be the same. When Web Capture generates a Page Set from a PDF file, the ID key will be generated solely from the bytes of the PDF file; there is no auxiliary data.

Web Capture Image Set ID

Web Capture calculates the ID for a Web Capture Image Set by feeding the source-data for the original image into the MD5 algorithm. For example, the ID for an Image Set created from a GIF image is calculated from the bytes of the GIF itself.

Web Capture Page Set Text ID

The **TID** (Text ID) for a Web Capture Page Set is calculated by feeding all of the rendered text present in the source-data into the MD5 algorithm. For example, for an HTML file, only the text between markup tags will be considered for the Text ID. No images are considered in this calculation.

Unique Name Generation

Note In this section, the term “name” refers to a string, not a name object.

In generating PDF pages and a Web Capture Page Set, Web Capture converts certain items such as hypertext links into named destinations and AcroForms fields. When updating the file, Web Capture needs to locate all the destinations and fields that were constructed for the Page Set. The names for these destinations and fields are derived from their original names, but they must be unique, in order to avoid conflicts among identical original names used in constructing destinations and fields for different Page Sets.

The unique name is formed by concatenating an encoded form of the Page Set’s **ID** string to the original name. The **ID** string must be encoded to remove characters that have special meaning in destinations and fields. For example, period (.) is used as the field separator in an AcroForms field name, so it must not appear in the **ID** portion of the unique name. Period is encoded as **#p**, NUL (ASCII 0) is encoded as **#0**, and **#** is encoded as **##**.

Implementation note See [page 465](#) for information about implementation limits in Web Capture.

6.17 Logical Structure in PDF

PDF 1.3

PDF 1.3 provides a mechanism by which structural information, such as the organization of a document into chapters and sections, can be represented in a PDF file. This mechanism is extensible, allowing PDF producers to choose which structural information is included and how it is represented, yet it enables PDF consumers to navigate a PDF file without knowing the producer’s conventions.

PDF logical structure shares basic features with standard document markup languages such as HTML, SGML, and XML. A PDF document's logical structure is expressed as a hierarchy of *structural elements*, each represented by a PDF dictionary. PDF structural elements are analogous to elements in SGML (or XML or HTML). Like their SGML counterparts, PDF structural elements can have content and attributes. The content of structural elements can be references to other structural elements, references to snippets of document content, or both. In PDF, however, rendered document content takes over the role occupied by text in SGML, XML, and HTML.

A PDF document's logical structure is stored apart from its visible content, with pointers from the logical structure to the physical representation and *vice versa*. This separation allows the ordering and nesting of logical elements to be entirely independent of the order and location of graphic objects on the document's pages.

6.17.1 Representing logical structure in PDF

The logical structure facilities of PDF comprise several kinds of object:

- The *structure tree root* (page 163) and *structural elements* (page 164) make up the tree that organizes the content of the PDF document. The structure tree root serves not only as the attachment point for the entire structure hierarchy, but also as a repository for information about the hierarchy as a whole.
- *Attribute objects* (page 169) allow structural elements to be annotated with application-specific information; they also allow such information to be associated with many elements by means of the *class* mechanism (page 168).
- The *role map* (page 177) and the *class map* (page 178) provide document-wide information that assists in the interpretation of structural elements. The *ID tree* (page 163) provides a document-wide index to structural elements by optional unique identifiers.
- *Marked Content references* (page 170) and *object references* (page 175) serve an auxiliary role in allowing structural elements to refer to a document's graphic content or to other information stored as a PDF dictionary or stream.
- The *parent tree* (page 163) is another auxiliary structure that forms part of the mechanism by which marked portions of graphic content refer to their containing structural elements (page 170).

In addition to these constructs, which exist only to support logical structure, there are several kinds of items that support logical structure but are located in other constructs within a PDF document:

- The document's **Catalog** dictionary has a key referring to the root of the document's structure tree (page 67).
- Marked content within streams (page 351) has a special mechanism allowing portions of graphic content to refer to their containing structural elements (page 170).

- Certain PDF objects such as Pages and Form Xobjects contain graphic content that can be referred to from a structural element. Such objects participate in the special backpointer mechanisms.

6.17.2 Structure tree root

The tree representing a document's logical structure is accessed from the Catalog object using the **StructTreeRoot** key. The structure tree root is the starting point for access to all of the tree's elements. [Table 6.81](#) gives the keys that a structure tree root may contain.

In one respect, the structure tree root is like a structural element: it may have kids. But because of its special role as the central point of access to a document's logical structure, the structure tree contains keys not present in structural elements, and it may only have structural elements as children.

Table 6.81 *Structure tree root attributes*

Key	Type	Semantics
Type	name	(Required) Object type. Always StructTreeRoot .
K (Kids)	array or dictionary	(Optional) The immediate child or children of the structure tree root in the structure hierarchy. This may be a single dictionary representing a structure element, or an array of such dictionaries.
RoleMap	dictionary	(Optional) The role map for logical structure in the document. The role map (Section 6.17.10) is a dictionary whose keys and values are element types that are roughly equivalent to each other.
ClassMap	dictionary	(Optional) The class map for logical structure in the document. The ClassMap (Section 6.17.12) is a dictionary whose keys are class names and whose values are one or more attributes objects (Section 6.17.4).
ParentTree Number tree		(Required if any element points to <i>Marked Content</i> or a <i>PDF Object</i>) A tree of indirect references to elements. This tree is part of the mechanism by which marked content within elements points back to its containing element. The value of each tree element can be either a reference to an element or an array of such references. See Section 6.17.5 and Section 6.17.6 .
ParentTreeNextKey	integer	(Optional) An integer greater than any key in the ParentTree . Maintaining this value allows new unique StructParents indices to be created easily.
IDTree	Name tree	(Required if any element has an <i>ID field</i>) A name tree that maps a string to a structural element having that string as the value of its ID field.

Example 6.19 *A structure tree root having one kid*

```
104 0 obj
```

```

<<
/Type /StructTreeRoot
/IDTree 105 0 R
/K 106 0 R
/ParentTree 107 0 R
/ParentTreeNextKey 3
>>
endobj

```

Example 6.20 *A structure tree root having multiple kids*

```

104 0 obj
<<
/Type /StructTreeRoot
/IDTree 105 0 R
/K [106 0 R 108 0 R]
/ParentTree 107 0 R
/ParentTreeNextKey 3
>>
endobj

```

6.17.3 Structural elements

Structural elements are represented by PDF dictionaries. The dictionary key-value pairs common to all structural elements are given in [Table 6.82](#).

Every structural element must have an *element type*, given by a name associated with its **S** key. The element type represents a classification of information within a document: typically, element types represent divisions of content such as sections, paragraphs, and figures, but element types may in fact be any name consistent with the PDF Naming Conventions described in [Appendix F](#).

There is a set of proposed standard element types, defined in *Standard Element Types for Logical Structure in PDF* [12], which facilitates data interchange among PDF applications. Applications may use whatever element types are convenient. Where an application's element types correspond in meaning to standard element types, however, a role map ([Section 6.17.10](#)) should be provided to indicate the relationship.

Attributes of structural element are represented by *attribute objects*, which are PDF dictionaries or PDF streams. A structural element can directly reference one or more attribute objects via the **A** (Attribute) key. An attribute object can be referenced by more than one structural element.

Another means of associating attribute objects with structural elements is by class. Unlike its usage in object-oriented programming, the use of the term *class* here does not refer to an inheritance mechanism, but rather to a way of specifying sets of elements that share certain properties. Structural elements have an optional key **C** (Class) key, whose value is the name of the class or classes that the structural element belongs to. These class names can be used as keys into the structure tree root's class map ([Section 6.17.12](#)) in order to find the associated attribute.

Attributes created by different applications are differentiated by their **O** (Owner) key. Elements have an **R** (Revision) key that enables Attribute Object owners to know if other applications have modified either the associated element or the element's contents, since both Attribute Object reference mechanisms (direct and class) specify the element's **R** key at the time of association.

Table 6.82 *Structural element attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always StructElem .
S (Subtype)	name	<i>(Required)</i> Structural element type (e.g., Paragraph , Footnote). Element subtypes must follow the PDF Name Registry guidelines in Appendix F .
P (Parent)	dictionary	<i>(Required)</i> A reference to the immediate ancestor of the structural element.
T (Title)	Text	<i>(Optional)</i> A title for the structural element. It is intended that the title describe the specific structural element, such as “Chapter 1”, rather than a generic element type, such as “Chapter”.
ID	string	<i>(Optional)</i> An identifying string by which the element can be accessed. This string must be unique among all the elements in a document's structure tree. The structure tree root's ID Tree (page 163) indexes structural elements according to their ID.
K (Kids)	<i>special</i>	<i>(Optional)</i> The contents of the structural element. These contents may be either another structural element or PDF graphic content. PDF graphic content within a stream is represented either as an MCR object (Section 6.17.7) or as an MCID (Section 6.17.5); PDF graphic content represented by an entire PDF dictionary or stream is referenced via an OBJR object (Section 6.17.8).
		The actual value associated with the K key may be a single dictionary (in the case where there is one child of a kind that is represented as a dictionary), a single integer (in the case of one marked content child when the Pg key is present); or it may be an array whose elements may each be a dictionary or an integer.
Pg (Page)	dictionary	<i>(Optional, required under conditions noted below)</i> A page containing some or all of the structural element's kids as given under the K attribute.
R (Revision)	integer	<i>(Optional, required as noted below)</i> Revision count.
A (Attributes)	<i>various</i>	<i>(Optional)</i> Attribute information associated with the structural element. The value can be a dictionary, a stream, or an array each of whose elements is a stream or integer, as explained below.
C (Class)	<i>special</i>	<i>(Optional)</i> A single name specifying a class, or an array of alternating classes (names) and revisions (integers) associated with this element.
Alt	Text	<i>(Optional)</i> An alternate representation of the structural element's contents.

The **K** (Kids) key identifies all the contents of an element. An element's contents can be portions of PDF page descriptions, PDF objects, or other structural elements.

6.17.3.1 Page description as structural element content

PDF page descriptions are stored in streams. To enable a specific sequence of page description operators to be referenced from outside the stream, the appropriate page descriptions must be enclosed in one or more marked content containers (see [Section 8.10.3](#)), abbreviated in this document as MCCs. An MCC containing the **MCID** key in its property list dictionary can be indirectly referenced from outside of the PDF stream by specifying both a reference to the containing PDF stream object and the specific MCID value. In the general case, a reference to an MCC is packaged up into a marked content reference (MCR) object, described in [Section 6.17.7](#).

The following example shows how a structural element refers to page content via an MCID.

Example 6.21 Structural element kid reference via MCID

```

1 0 obj                               Structural element object
<</Type /StructElem
/S /P                                 Subtype
/P 5 0 R                               Parent
/Pg 2 0 R                               Page
                                        Kid references given as integers will be
                                        interpreted as being in the content of this page.
/K 0                                     There is only one kid, so an array is not necessary.
>>
endobj

2 0 obj                               Page object
<</Type /Page
/Contents 3 0 R
...
>>
endobj

3 0 obj                               Page Contents stream
<</Length 3587>>
stream
...
/P <</MCID 0>> BDC                    Marked Content, with ID = 0.
(Here is some text) TJ ...
EMC
...
endobj

```

6.17.3.2 PDF objects as structural element content

A structural element can logically contain an entire PDF object, such as an annotation or a Form XObject. In this case, the reference to that PDF object is achieved via an object reference object (OBJR), described in [Section 6.17.8](#).

6.17.3.3 Structural elements as structural element content

Structural elements can also contain other structural elements, even if the containing element also contains PDF page descriptions or other PDF objects. The contained structural elements are stored in the containing structural element's **K** array. The following example shows a structural element containing other structural elements and page contents.

Example 6.22 Structural element with mixed content

```

1 0 obj                                     Parent element object
<</Type /StructElem
/S MixedContainer
/K                                           Three kids:
    [2 0 R                                   a structural element (object 2 0),
      0                                       marked content on a page (MCID 0),
      3 0 R]                               and another structural element (object 3 0).
/Pg 4 0 R
>>
endobj

```

Child elements

```

2 0 obj <</Type /StructElem ... >> endobj
3 0 obj <</Type /StructElem ... >> endobj

```

6.17.3.4 Optimizing page references from structural elements

To reduce file size in the case in which all or most of a structural element's kids are on the same page, the **Pg** key can be supplied. If **Pg** is present in a structural element, then two optimizations are possible:

1. MCR objects referencing MCCs and OBJR objects referencing other objects on that element's page need not have their own **Pg** key.
1. PDF page descriptions located in the specified page's contents stream can be identified in **K** simply by their MCID, which is an integer.

This latter form will be the most common case in a typical structured document.

6.17.3.5 Structural element attributes

Applications can attach attributes to any element by creating an attribute object (see [Section 6.17.11](#)) to hold them, and referencing the attribute object from the structural element, either directly via the **A** (Attributes) key, or indirectly via the **C** (Class) key. Applications can keep track of changes made by other clients to a structural element or its contents via the **R** (Revision) key.

R is a non-negative integer; it must be incremented by each successive application that modifies the structural element or its hierarchical contents. If the **R** key is absent, the value is assumed to be zero. The **R** key is used if and only if an application *A* modifies a structural element that had attributes attached to it by another application *B*, unless that attribute object is removed by application *A*.

R only counts changes to the structural element itself and its hierarchical contents: changing the attributes one has attached to a structural element does not change the revision count. See [Section 6.17.13](#) for a complete discussion of structural element modification rules.

Note A Revision number is independent of an indirect object's generation number.

The **A** key associates attribute objects with a structural element. In the general case, the associated value is an array of alternating attribute object references and revision numbers: an attribute object is associated with the element at the revision number immediately following it in the array. Thus, if the element is subsequently modified and its revision number incremented, clients accessing the attribute object will be alerted that the element has changed. To enable a single attribute object to be associated with many elements (each with different revision numbers), the revision number is not placed in the attribute object, but in the array that associates the attribute object with the structural element.

To save file space, if a attribute object is associated with an element whose revision number is zero, the revision number can be omitted from the **A** array. Furthermore, if the array contains only one element (which must be an Attribute Object reference), the value of **A** is simply a reference to the Attribute Object. The use of **A** in elements that are subsequently modified is further discussed in [Section 6.17.11](#).

6.17.3.6 Structural element class

The **C** (Class) key associates attribute objects with a structural element, but it does so indirectly through the *class map* of the root of the structure tree.

In the general case, the value associated with **C** is an array of alternating class names and revision numbers. The names are used as the keys in the class map to obtain the actual attribute object or objects. As with **A**, revision numbers can be omitted from the array if they are zero, and an array with only one element can be represented just as the name the class. Class names must follow the PDF Name Registry guidelines (see [Appendix F](#)).

C (Class) can also be used by itself, without being associated with any attribute object. In this case the class map will not contain a key for the class name. Even if not associated with an attribute object, class names attached to a structural element with a non-zero revision number must be followed by the revision number.

This class mechanism has nothing to do with object-oriented programming. Class is not an inheritance mechanism: it is a convenient and compact information association mechanism.

Use of **C** is further discussed in [Section 6.17.12](#).

6.17.3.7 Structural element alternative text

Alt is provided as a standard place for applications to locate a human-readable text representation of the structural element's contents. For instance, it can be used to provide text that takes the place of the structural element's contents when those contents cannot be displayed by a particular viewing application.

6.17.4 Attribute Objects

Any PDF dictionary or stream may serve as an attribute object as long as it contains the following key:

Table 6.83 *Attribute object attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
O (Owner)	name	<i>(Required)</i> The name of the client responsible for the attribute data, such as an application that processes PDF. The name must follow the PDF Name Registry guidelines in Appendix F .

6.17.5 Marked Content Containers (MCC)

Marked content containers referenced by structural elements must have dictionaries that contain the following key:

Table 6.84 *Marked content property list attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
MCID	integer	<i>(Required)</i> A non-negative integer. This number uniquely identifies the particular marked content container among all those in a stream (or array of streams forming the contents of a page).

The **MCID** value is used to identify which MCC is being referenced by an MCR object or structural element. Since MCIDs serve as indices into an array in the parent tree of the structure tree root, assigning MCIDs as small as possible will conserve space in the array.

Note As documented in [Section 8.10.3](#), marked content delimiters associate a named tag with marked content. This tag name is not used by the structure tree or its reference mechanisms. However, it is suggested for readability that the tag be the same as the referencing structural element's **S** value.

6.17.6 Parent pointers from streams containing marked content

Marked content is stored in streams. Streams cannot contain PDF object references, so the marked content cannot refer directly to its parent structural element. A different mechanism, the **StructParents** key, is used for this purpose. It enables applications to find the containing structural element from the marked content.

Table 6.85 *Marked content container dictionary attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
StructParents	integer	<i>(Required)</i> Index into the StructTreeRoot's ParentTree , resulting in a value that is an array of pointers to elements.

If the marked content is in the contents stream of a Page, then the **StructParents** key appears in the Page dictionary; for Form and Image XObjects, the **StructParents** key appears in the XObject's dictionary.

Access to the containing structural element from a marked content container is accomplished indirectly in two steps:

1. The containing stream's **StructParents** value is used as an index into the number tree stored under the **ParentTree** in the structure tree root. This yields an array of references to structural elements.
2. The specific marked content's **MCID** value is used as an index into that array, yielding a reference to the structural element containing the MCC as one of its kids, either directly or via an MCR object.

See [Example 6.23 on page 171](#).

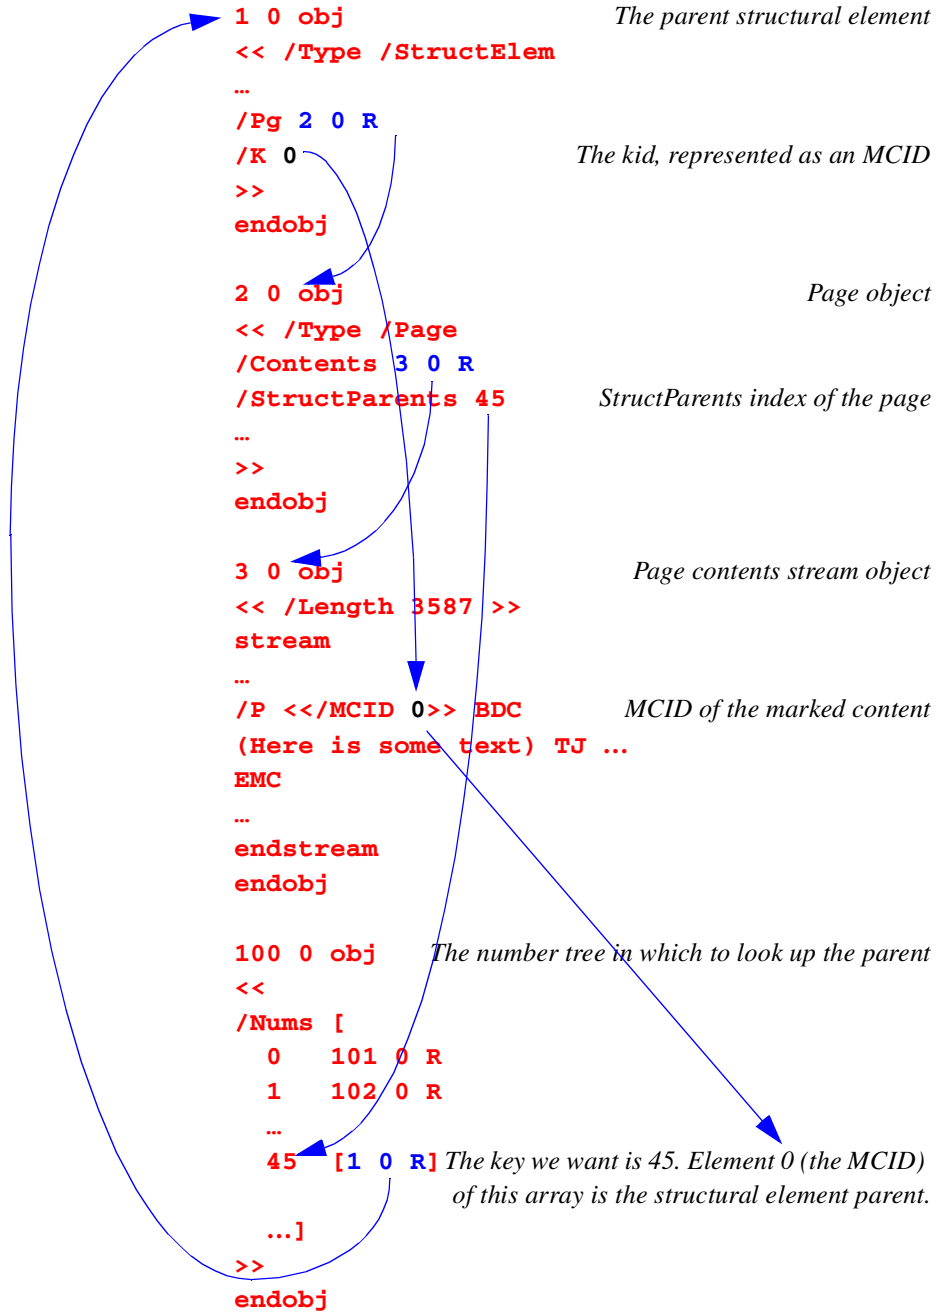
6.17.7 Marked Content Reference (MCR)

A marked content reference (MCR) is a PDF dictionary whose key-value pairs indirectly reference a specific MCC ([Section 6.17.5](#)). MCRs are typically referenced via a structural element's **K** key value.

Table 6.86 *Marked content reference attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	<i>(Required)</i> Always MCR .

Example 6.23 *Accessing containing structural element from marked content*



Pg (Page)	dictionary	<i>(Optional, Required under conditions described below)</i> The page on which the MCC's contents are rendered.
Stm (Stream)	stream	<i>(Required if value is not the page's contents)</i> The stream containing the referenced MCC.
StmOwn (StreamOwner)	<i>any</i>	<i>(Optional)</i> The PDF object owning the stream. (See the text below.)
MCID	integer	<i>(Required)</i> An integer identifying the specific MCC within the specified stream.

Pg (Page) specifies the page on which the PDF page descriptions inside the referenced MCC are rendered. **Pg** is required unless the structural element referencing the MCR object has a **Pg** key that matches the page that would be specified in the MCR.

Stm (Stream) specifies the PDF stream containing the MCC. If this is a stream that forms all or part of the contents of a page, then the **Stm** key can be omitted. Two examples of streams that are not the contents of a page are annotations and Form XObjects.

If the referenced stream is just a part of a larger PDF object, then that larger object can be referenced in **StmOwn** (StreamOwner). For example, an MCC in the Normal stream of an annotation appearance can reference the annotation in **StmOwn** and the actual stream that contains the MCC in **Stm**.

MCID is an integer that identifies which MCC in the stream is being referenced. It must be unique within the stream. For the stream or streams which make up a page's **Contents** these integers must be unique within the whole page, whether it has one content stream or several. [Section 6.17.6](#) describes the requirements placed on streams that contain MCC's referenced from a structural element's **K** key.

Example 6.24 *Referencing page content via an MCR*

```

1 0 obj                                The containing structural element
<<
/S /P
/K <</Type /MCR
  /Pg 3 0 R    This key is required in the MCR because
                there is no Pg key in the structural element.
  /MCID 0>>
/P 2 0 R
>>
endobj

3 0 obj                                The page containing the content
<<
/Type /Page
/Contents 4 0 R
...
>>
endobj

4 0 obj                                The content stream of the page
<</Length 1729>>
stream
...
/P <</MCID 0>> BDC
(Here is some text) TJ ...
EMC
...
endstream
endobj

```

Example 6.25 *Referencing the content of a Form XObject via an MCR*

```

1 0 obj <<                               The containing structural element
  /Type /StructElem
  /S /P
  /K <</Type /MCR /Pg 4 0 R /Stm 3 0 R /MCID 0>>
  /P 2 0 R
>> endobj

3 0 obj << The Form XObject stream containing the content
  /Type /XObject
  /Subtype /Form
  /Length ...
>> stream
...
/P <</MCID 0>> BDC
(Here is some text) TJ ...
EMC
...
endstream
endobj

4 0 obj <<
  /Type /Page
  /Resources << /XObject << /Fm3 3 0 R >> >>
  /Contents 5 0 R
  ...
>> endobj

5 0 obj <<
  /Length ...
>> stream
...
/Fm3 Do
...
endstream
endobj

```

6.17.8 Object reference (OBJR)

There are PDF objects, such as Form XObjects, that specify content that appears on a page without being part of that page. An object reference (OBJR) is a means of incorporating such a PDF object as a whole into the content of a structural element in conjunction with the Page on which it is drawn.

Table 6.87 *Object reference attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	(<i>Required</i>) Always OBJR .
Pg (Page)	dictionary	(<i>Optional, required under conditions described below</i>) The page on which the object is rendered.
Obj (Object)	any	(<i>Required</i>) The referenced object.

Pg (Page) specifies the page that the object is drawn on. **Pg** is required unless the structural element referencing the OBJR object has a **Pg** key that matches the page that would be specified in the OBJR.

If the referenced object is rendered on multiple pages, then each rendering requires a separate OBJR object. If it is drawn multiple times on the same page, then only one OBJR object is needed to identify all of them. Referenced objects may be Form or Image XObjects, or any other PDF object that is considered part of a structural element; for instance, an annotation object could be included as the content of a structural element by means of an OBJR.

Note OBJRs are for referencing an entire PDF object as part of the content of a structural element. Reference to graphic content within a PDF object is accomplished via MCR objects.

Note If an application needs to distinguish between multiple Form XObject renditions on the same page, then the application should access the rendition via a marked content container enclosing the specific “**Do** formName” portion of the stream rather than via an OBJR.

6.17.9 Parent pointers from PDF objects to structural elements

The structural element parent of an object referred to via an OBJR is stored by a mechanism similar to that used for such references from graphic content stored in a stream.

In the case of parent pointers from PDF objects, the object itself must contain a **StructParent** key, as specified in the following table.

Table 6.88 *PDF object structure pointer attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
StructParent	integer	(Required) Index in the structure tree root's ParentTree , mapping to a value that is a reference to the structural element containing the object.

Unlike the mechanism for structural element parent pointers from stream content, there is only one parent per object having the index into the parent tree. For this reason, the entry in the parent tree for a PDF object is an indirect reference to the parent structural element itself, rather than an array of such references.

Example 6.26 *Parent pointer from Form XObject*

```

3 0 obj                                     Form XObject
<<
  /Type /XObject
  /Subtype /Form
  ...
  /Length 1729
  /StructParent 45                           Index into parent tree
>>
stream
...
endstream
endobj

100 0 obj   The number tree in which to look up the parent
<<
  /Nums [
    0 101 0 R
    1 102 0 R
    ...
    45 1 0 R   The key we want is 45. This entry is
                the structural element parent itself.
  ...]
>>
endobj

1 0 obj   The parent structural element object
<< /Type /StructElem
...
  /K <</Type /OBJR /Obj 3 0 R /Pg 2 0 R>>
>>
endobj

```


6.17.10 The role map

The role map is a dictionary referenced from the structure tree root, whose keys and value are the names of structural element subtypes considered roughly equivalent. For instance, an application can create an element whose **S** value is **TextWithinAPath**. By including in the role map a key of **TextWithinAPath** whose value is **P** (which is a standard structural subtype, as proposed in [12]), applications that don't know about **TextWithinAPath** can treat the element as a **P** (paragraph) for their purposes.

When non-standard structural element types are used, it is strongly recommended that the role map be created to map them to the closest standard structural element types. This enables a structure-creating application to employ structural types of its own definition while providing clues to other applications as to what role this new structural element plays.

The association of two structural element types in the role map merely denotes approximate equivalence of meaning. It is not required that either the keys or the values in the role map be from any particular set of structural element types.

Structural element types entered as values in the role map can also be present as keys mapped to some other structural element type. Circular chains of association are explicitly permitted. Applications making use of the role map should iterate their lookups until they find either a structural element type they recognize, or else they encounter a structural element type already encountered. Circular mapping enables a single role map to define a bidirectional mapping.

6.17.11 Attaching attributes

An application can attach attributes to any structural element, even if the application did not create that structural element, or if other applications have already attached their own attributes. This can be achieved with the **A** (Attributes) and/or **C** (Class) element keys.

A (Attributes) element key

If the structural element does not yet have an **A** value, then the client simply creates an attribute object and makes that the value of **A** (including the structural element revision number if it is not zero, in which case **A**'s value must be an array).

If the structural element has an **A** value that is not an array, the client must create an array, move the existing attribute object reference into that array, then add the reference to its own attribute object (and revision number if not zero) in the array. Other than the association of attributes objects being followed by revision numbers at time of attachment, the array is unordered, but it is considered good policy to add new attribute objects to the end so that the structural element originator will normally have added the first entry in the array as its attribute object.

Similarly, if a client's attribute object is removed from the array and there remains only one attribute object, the array should be destroyed and **A** should again point to the only attribute object attached to the element.

6.17.12 Element class key and class map

If many structural elements share the same set of attribute values, the shared values can be put into an attribute object (Section 6.17.4) that is associated with a name, called the *class name*. This attribute object is indirectly attached to a structural element by adding that class name and element revision number at time of attachment to the array of names associated with the structural element's **C** (Class) key. The association of class name to attribute objects is maintained in the **ClassMap** dictionary attached to the structure tree root.

The class map is a dictionary whose keys are class names, and whose values are references to attribute objects. The value of any key can be an individual attribute object reference, or an array of such references.

6.17.13 Element Modification Rules

If a client wishes to modify a structural element or its contents and the structural element has attributes attached to it by other clients, the modifying client must do exactly one of the following:

1. Increment the **R** (Revision) key of the structural element. If the **R** key is not present, one must be added to the structural element, with a value of 1.
2. Remove all unknown attribute objects from **A** (Attributes) and **C** (Class) arrays associated with the structural element.

Which an application chooses (they are mutually exclusive) should be decided on the degree of structural element modification. Incrementing the revision count is appropriate for small changes, since it is possible for the other applications to reconstruct. For large changes, which are likely to destroy the state of the structure on which the originating application's attribute information was based, the outdated attribute information should just be removed.

Note Applications creating attribute objects should be designed keeping in mind that other applications may delete some or all of them at any time.

6.17.14 Extended example of logical structure

Here is an example showing portions of a PDF file that has a simple document structure. The Structure Tree Root (object #3) contains a MyChapter (#12) and a MyP (#15). The MyChapter, titled *Chapter 1*, contains an H1 (#13) and a MyP (#14).

Example 6.27 *Extended example*

```

1 0 obj <<                                     Catalog dictionary
  /Type /Catalog
  /Pages 2 0 R
  /StructTreeRoot 3 0 R
>> endobj

```

```

2 0 obj <<                                     Page tree
/Type /Pages
/Kids [4 1 R 5 0 R]
/Count 2
>> endobj

```

```

3 0 obj <<                                     Structure tree root
/Type StructTreeRoot
/K [12 0 R 15 0 R]   Two kids, both structural elements
/RoleMap
  << /MyP /P /H1 /H /MyChapter /Chapter >>
/ParentTree 100 0 R
/ParentTreeSize 2
/IDTree 95 0 R
>> endobj

```

```

100 0 obj <<                                   Parent tree objects (a number tree)
/Nums [
  0 101 0 R
  1 102 0 R ]
>> endobj

```

```

101 0 obj                                     Structure parents for page 1
[13 0 R 14 0 R]
endobj

```

```

102 0 obj                                     Structure parents for page 2
[14 0 R 15 0 R 15 0 R]
endobj

```

```

95 0 obj <<                                   End of parent tree objects
/Kids [96 0 R]                                   ID tree object (a name tree)
>> endobj

```

```

96 0 obj <<
/Limits [(Chap1) (Sec1.3)]
/Names [
(Chap1) 12 0 R
(Sec1.1) 13 0 R
(Sec1.2) 14 0 R
(Sec1.3) 15 0 R ]
>> endobj                                     End of ID tree objects

```

```

4 1 obj <<                                     First page object
/Type /Page
/MediaBox [0 0 612 792]
/Parent 2 0 R
/Contents 8 0 R

```

```

/StructParents 0
/Resources
  <<
    /Font <</F1 6 0 R /F12 7 0 R>>
    /ProcSet [/PDF /Text]
  >>
>> endobj

5 0 obj <<                                     Second page object
/Type /Page
/MediaBox [0 0 612 792]
/Parent 2 0 R
/Contents 10 0 R
/StructParents 1
/Resources
  <<
    /Font <</F1 6 0 R /F12 7 0 R>>
    /ProcSet [/PDF /Text]
  >>
>> endobj

8 0 obj <<                                     Page 1's Contents stream
/Length 842
>> stream
1 1 1 rg
0 0 612 792 re f
BT
/H1 <</MCID 0>> BDC
0 0 0 rg
/F1 1 Tf
30 0 0 30 18 732 Tm
(This is a first level heading. he\
llo world: )Tj
1.1333 TL T* (goodbye universe.)Tj
EMC
/P <</MCID 1>> BDC
/F12 1 Tf
14 0 0 14 18 660.8 Tm
(This is the first paragraph, which spans p\
ages. It has four fairly short and concise s\
entences. This is the next to last ) Tj
EMC
ET
endstream
endobj

10 0 obj <<                                     Page 2's Contents stream
/Length 1709

```

```

>> stream
1 1 1 rg
0 0 612 792 re f
BT
/P <</MCID 0> BDC
0 0 0 rg
/F1 1 Tf
30 0 0 30 18 732 Tm
(sentence. This is the very last sentence o\
f the first paragraph. ) Tj
EMC
/P <</MCID 1>> BDC
/F12 1 Tf
14 0 0 14 18 570.8 Tm
(This is the second paragraph. It has four \
fairly short and concise sentences. This is \
the next to last )Tj
EMC
0 0 m
42 0 1
0 42 1
-42 0 1
0 -42 1
/P <</MCID 2>> BDC
1.1429 TL T*
(sentence. Here is the last sentence of the \
second para. ) Tj
EMC
ET
endstream
endobj

```

```

12 0 obj      Structural element representing the chapter
<<
/Type /StructElem
/S /MyChapter
/ID (Chap1)
/T (Chapter 1)
/P 3 0 R      Parent is the Structure Tree Root.
/K [13 0 R 14 0 R] This Chapter contains 2 elements.
>>
endobj

```

```

13 0 obj      Structural element representing a section
<<
/Type /StructElem
/S /H1
/ID (Sec1.1)

```

```

/P 12 0 R      Parent is the Chapter structural element.
/Pg 4 1 R
/K 0          Reference to the MCC on the first page
>> endobj

```

```

14 0 obj << Structural element representing a paragraph
/Type /StructElem
/S /P
/ID (Sec1.2)
/P 12 0 R      Parent is the Chapter structural element
/Pg 4 1 R      One of the pages containing this element
               There are two kids, each on a different page:
/K [1
   <</Type /MCR /Pg 5 0 R /MCID 0>>
 ]
>> endobj

```

```

15 0 obj <<Structural element representing another paragraph
/S /P
/ID (Sec1.3)
/P 3 0 R      Parent is the structure tree root
/Pg 5 0 R
/K [1 2]      Two kids on second page in this paragraph
>> endobj

```

CHAPTER 7

Common Data Structures

Chapter 4 describes all the basic object types in PDF. In this chapter, we describe data structures that are built from these basic types but which occur so often in PDF files that it is useful to regard them as types in their own right.

7.1 Rectangle

Rectangles are used to describe locations on the page and bounding boxes for several objects in PDF, such as fonts. A rectangle is represented as an array of four numbers, `[ll_x ll_y ur_x ur_y]`, specifying the lower left x , lower left y , upper right x , and upper right y coordinates of the rectangle, in that order.

Note Many operations permit a rectangle to be described by any two diagonally opposite corners, not necessarily lower-left and upper-right. PDF consumers should be prepared to “normalize” rectangles in situations where specific corners are required.

7.2 Date

PDF defines a standard date format. The PDF date format closely follows the format defined by the international standard ASN.1 (Abstract Syntax Notation One, defined in CCITT X.208 or ISO/IEC 8824). A date is a string of the form:

`(D:YYYYMMDDHHmmSSOHH'mm')`

where

- `YYYY` is the year.
- `MM` is the month (`01–12`).
- `DD` is the day (`01–31`).
- `HH` is the hour (`00–23`).
- `mm` are the minutes (`00–59`).
- `SS` are the seconds (`00–59`).

PDF 1.1



- *O* is one of the characters **+**, **-**, or **Z**, indicating the relation of local time to GMT. **+** indicates that local time is later than GMT, **-** indicates that local time is earlier than GMT, and **Z** indicates that local time is GMT.
- *HH* is the absolute value of the offset from GMT in hours. The quote (**'**) is part of the syntax.
- *mm* is the absolute value of the offset from GMT in minutes. The quote (**'**) is part of the syntax.

Example 7.1 *Date string*

December 23, 1998, 7:52 PM, US Pacific time is represented by:

D:199812231952-08'00'

The **D:** prefix permits arbitrary keys to be recognized as dates. However, it is not required. Trailing fields other than the year are also optional. The default value for day and month is 1; all other numerical fields default to 0. If no GMT information is specified, the relationship of the specified time to GMT is considered unknown. Whether the time zone is known or not, the rest of the date should be specified in local time.

7.3 Destination



An annotation or Outline entry may specify a *destination*, which consists of a page, the location of the display window on that page, and the zoom factor to use when displaying that page.

A destination may be represented explicitly as an array, or implicitly through a name; see [page 185](#).

7.3.1 Explicit destinations

[Table 7.1](#) shows the allowed forms of the destination. In the table, *top*, *left*, *right*, and *bottom* are numbers specified in the default user space coordinate system. *page* is an indirect reference to the destination Page object, except in the case of the **GoToR** action, where it is a page number. The page's *bounding box* is the smallest rectangle enclosing all objects on the page. No side of the bounding box is permitted to be outside the page's crop box. If it is, that side of the bounding box is defined by the corresponding side of the crop box.

Table 7.1 *Destination specification*

<i>Destination</i>	<i>Semantics</i>
[page /XYZ left top zoom]	

If *left*, *top*, or *zoom* is null, the current value of that parameter is retained. For example, a destination of **[4 0 R /XYZ null null null]** specifies the page object with an object ID of 4, retaining the same *top*, *left*, and *zoom* as the current page. A zoom of 0 has the same meaning as a zoom of null.

[<i>page</i> /Fit]	Fit the page to the window.	
[<i>page</i> /FitH <i>top</i>]	Fit the width of the page to the window. <i>top</i> specifies the <i>y</i> -coordinate of the top edge of the window.	
[<i>page</i> /FitV <i>left</i>]	Fit the height of the page to the window. <i>left</i> specifies the <i>x</i> -coordinate of the left edge of the window.	
[<i>page</i> /FitR <i>left bottom right top</i>]	Fit the rectangle specified by <i>left bottom right top</i> in the window. If the height (<i>top</i> – <i>bottom</i>) and width (<i>right</i> – <i>left</i>) imply different zoom factors, the numerically smaller zoom factor is used, to ensure that the specified rectangle fits in the window.	
[<i>page</i> /FitB]	Fit the page’s bounding box to the window.	PDF 1.1
[<i>page</i> /FitBH <i>top</i>]	Fit the width of the page’s bounding box to the window. <i>top</i> specifies the <i>y</i> -coordinate of the top edge of the window.	PDF 1.1
[<i>page</i> /FitBV <i>left</i>]	Fit the height of the page’s bounding box to the window. <i>left</i> specifies the <i>x</i> -coordinate of the left edge of the window.	PDF 1.1

7.3.2 Named destinations

A destination may be represented implicitly, using a *string* or a *name*. Both of these cases are referred to as *named destinations*. These are especially useful when the destination is in another PDF file. For example, one file may contain a link to the first page of Chapter 6 in another file. If the link uses a name such as

/Chap6.begin

rather than an explicit location, such as a certain rectangle on page 42, then the place where Chapter 6 starts can change without invalidating the link.

The mapping from *strings* to destinations is a feature of PDF 1.2. The Catalog of any document may contain a **Names** key whose value is a dictionary (see [Section 6.9, “Names dictionary”](#)). Each value in this dictionary is a [Name tree](#), which maps strings to indirect objects. The **Dests** entry in this dictionary is a name tree where the indirect objects are destinations.

Note The keys in the **Dests** name tree may be treated as [Text](#) for display purposes.

The mapping from *names* to destinations is a feature of PDF 1.1, and it is supported in later versions for backwards compatibility. The Catalog of any document may contain a **Dests** key. Each key in this dictionary is a name, and the corresponding value is either a destination or a dictionary. If it is a dictionary, it must have a **D** key whose value is a destination. (The dictionary enables named destinations to have additional attributes.)

Note There are several performance advantages to using strings instead of names for named destinations. See [page 389](#).

If an annotation or Outline entry that contains a named destination also includes a file specification (**F** key), then the destination is in the specified file. Otherwise, if there is no **F** key, the destination is in the current file.

7.4 File specification

PDF 1.2

A file specification, together with a file system, describes the location of a file. A *simple* file specification is one that does not specify the file system to be used, and a *full* file specification includes information related to one or more file systems. A file specification is represented using a string or a dictionary. Only simple file specifications may be represented as strings. Both full file specifications and simple file specifications may be represented as a dictionary. The simple file specification has a standard format for representing the name of the referenced file; this format is independent of operating system naming conventions.

Although the files specified by the file specification are normally external to the PDF file, copies of these external files may be embedded within the PDF file, to allow transmission of a single PDF file that includes some or all of the files referenced by that PDF file. Embedding a file does not change the presumption that the file is external to the PDF file. Therefore, it may be necessary to copy the embedded files back into a local file system for correct processing of that PDF file.

PDF 1.3

The correct processing of a PDF file is dependent on the external references in that PDF file being correct. To enable the maintenance of file specifications, it should be possible to find all the file specifications that might need to be changed for any reason. Therefore, the following guidelines are strongly recommended:

- All file specifications should use the dictionary form, with a **Type** key whose value is the name **Filespec**.
- All references to file specifications should be indirect.

When external files are embedded within the PDF, the file specification for these files *must* use the indirectly referenced dictionary form of the file specification and *must* include a **Type** key with the value **Filespec**.

7.4.1 File specification strings

The standard format for a simple file specification divides the string into component strings separated by the slash character (**/**). The slash is used as a generic component-separator that is mapped to the appropriate separator when generating a system-dependent file name. The component string may be empty, and if the component string contains one or more slashes (e.g., “in/out”) then each such literal slash must be preceded by a backslash:

(in\\out)

Note that the backslash must itself be preceded by a backslash to indicate it is being used as a character in the string and not as the escape character. The backslashes are removed in defining the components; they are needed only to

distinguish the component values from the component separators. The component strings are stored as octets and are passed to the operating system without interpretation or conversion of any sort.

A simple file specification that begins with a slash is an *absolute* file specification. Within an absolute file specification, the last component is the file name, and the preceding components are the context. The file name may be empty in some file specifications; for example, URL specifications can specify directories instead of files. A file specification that begins with a component (i.e., one that does not begin with a slash) is a *relative* file specification. A relative file specification is relative to the file specification of the document that contains the relative file specification.

In the case of a URL file system, the rules of RFC 1808, *Relative Uniform Resource Locators* [23], are used to compute an absolute URL from the document's file specification and a relative file specification. Prior to this process, the relative file reference is converted into a relative URL by using the escape mechanism of RFC 1738, *Uniform Resource Locators* [18], to represent any octets that would be either "unsafe" according to RFC 1738 or not representable in 7-bit US ASCII. In addition, such URL-based relative file references are limited to being paths as defined in RFC 1808; the scheme, network location/login, fragment identifier, query information, and parameters are not allowed.

In the case of other file systems, an absolute file specification is created from a relative file specification and the file specification of the document that contains the relative file specification by removing the file name component of the document's file specification and appending the relative file specification.

The special component `..` allows condensing of a file specification. Proceeding from left to right, whenever a component that is not `..` is followed by `..`, that component and the `..` are eliminated from the file specification, and the process is begun again. This allows relative file specifications that are relative to an initial segment of an absolute file specification.

The conversion of a file specification into a system-dependent file name is specified for each file system. For the Macintosh, the components are separated by colons (`:`). For UNIX, the components are separated by slashes, and an initial slash, if present, is preserved. For DOS, the initial component is either a physical or logical drive identifier or a network resource name as returned by the Microsoft Windows function `WNetGetConnection` and is followed by a colon. A network resource name is constructed from the first two components of the file specification; the first component is the server name and the second component is the share name (volume name). All the components are then separated by backslashes. It is possible to specify an absolute DOS path without a drive by making the first component empty. (Empty components are ignored by other platforms.)

[Table 7.2](#) provides examples of file specifications on various platforms.

Multi-byte strings in file specifications

PDF 1.2

Table 7.2 Examples of file specifications

System	System-dependent path	Written as...
Macintosh	Mac HD:PDFDocs:spec.pdf	(/Mac HD/PDFDocs/spec.pdf)
DOS	\pdfdocs\spec.pdf (no drive)	(//pdfdocs/spec.pdf)
DOS	r:\pdfdocs\spec.pdf	(/r/pdfdocs/spec.pdf)
DOS	pclib/eng:\pdfdocs\spec.pdf	(/pclib/eng/pdfdocs/spec.pdf)
UNIX	/user/fred/pdfdocs/spec.pdf	(/user/fred/pdfdocs/spec.pdf)
UNIX	pdfdocs/spec.pdf (relative)	(pdfdocs/spec.pdf)

A string used to specify a file name may contain multi-byte character codes. Since the slash character <2F> is used as a delimiter for components of a file name, and the backslash character <5C> is used as an escape character, any occurrence of either of these bytes in a multi-byte character must be preceded by the ASCII byte for the backslash character.

For example, a file name that contains the double-byte character code <89 5C> must be written as <89 5C 5C>. When the viewer processes a file name with this sequence of bytes, it replaces the sequence with the original double-byte code.

The strings used to specify a file name are interpreted in the platform encoding where the document is being viewed. Where system portability is required, it is recommended that file names consisting only of ASCII characters be used; see [Section 7.4.5, “Safe path names.”](#)

7.4.2 File specification dictionaries

A file specification can be either a string, formatted as described above, or a dictionary. The dictionary form of the file specification provides for platform-specific file specifications and allows extension of the form of file specifications. A dictionary that contains a platform-specific file system key or a file system key (**FS**) is a *full* file specification. This provides alternate ways to locate a file.

A viewer should use the appropriate platform-specific key (**Mac**, **DOS**, or **UNIX**). If it does not find the appropriate platform-specific key and there is no file system value (**FS**), it should treat the value of the file specification key (**F**) as a simple file specification. The keys need not specify the same file, allowing a single file specification to describe appropriate but different files for different platforms.

[Table 7.3](#) describes the file specification dictionary attributes.

Table 7.3 *File specification attributes*

Key	Type	Semantics
Type	name	<i>(Required if an EF or RF key is present; recommended always)</i> Object type. Always Filespec .
FS (FileSystem)	name	<i>(Optional)</i> The name of the “file system” to be used to interpret this file specification. A viewer or plug-in can register a file system. A <i>file system</i> interprets file specifications, opens files, and provides the usual input and output operations. If a file specification includes a file system, all other keys are interpreted by this file system. Note that this key is independent of the F , Mac , DOS , and Unix keys.
F (File)	string	<i>(Required if the Mac, DOS, and Unix keys are all omitted)</i> A file specification using the string format described earlier in this section, or (if the file system is URL) a URL, described in Section 7.4.4 on page 194 . A viewer that encounters an action involving a file specification with no F key and that does not understand any of the alternative keys need not do anything.
Mac	string	<i>(Optional)</i> A string that specifies a Macintosh file name using the string format described above.
DOS	string	<i>(Optional)</i> A string that specifies a DOS file name using the string format described above.
Unix	string	<i>(Optional)</i> A string that specifies a UNIX file name using the string format described above.
ID	array	<i>(Optional)</i> An array of two strings. The ID is a file ID as described in Section 6.12, “File ID . This allows a viewer to find the intended file more often, and it allows viewers to warn a user if the file has changed since the link was made.
V (Volatile)	boolean	If V is <i>true</i> , this indicates that the document referenced by the file specification changes frequently with time. An implementation can use this value to determine whether it is safe to use a cached copy of a document. For example, a Movie annotation could reference a URL to a live camera; if V is <i>true</i> , then the implementation could determine that it should reacquire the Movie each time it is played. The default value is <i>false</i> .
EF (Embedded files)	dictionary	<i>(Optional)</i> The keys in the EF dictionary are a subset of the filename keys in the file specification dictionary: F , Mac , DOS , and Unix . The value of each such key is an Embedded File stream (see Section 7.4.3 on page 192), containing the corresponding file. See Example 7.2 on page 191 .
RF (Related Files)	dictionary	<i>(Optional)</i> A dictionary which has the same key structure as the EF dictionary. Each key in the RF dictionary must also be in the EF dictionary. Each value is a Related Files array (see Section 7.4.2.1 on page 190) that identifies file that are “related” to the corresponding file in the EF dictionary. See Example 7.2 on page 191 .

PDF 1.3**PDF 1.2****PDF 1.3****PDF 1.3**

The string values of the **DOS**, **Mac** and **Unix** keys should not be modified by the implementation.

If a file specification dictionary contains an **EF** key or **RF** key, then it must be indirectly referenced, and the **Type** key is required.

7.4.2.1 Related Files array

PDF 1.3

A Related Files array is used when an external entity consists of more than one file (all of which must be embedded), such as the set of five files that make up a DCS 1.0 entity, but where only one of those files is mentioned in a file specification.

The array has $2 \times N$ elements which are paired, in the form

[string stream string stream ...]

The first element of the pair is a string and the second element is an Embedded File stream. The string holds the name of the related file, and the Embedded File holds the contents of the related file.

Example 7.2 A file specification dictionary with **EF** and **RF** keys

```

10 0 obj                                A file specification dictionary
<< /Type /Filespec
  /Mac (Sunset.eps)                    The name of the Mac file
  /Unix (Sunset.eps)
  /DOS (SUNSET.EPS)
  /EF << /Mac 19 0 R                  The embedded Mac file
          /Unix 13 0 R
          /DOS 14 0 R >>
  /RF << /Mac 11 0 R >>               Files related to the Mac file
>>
endobj

11 0 obj                                The files related to the Mac file Sunset.eps
[ (Sunset.eps) 19 0 R                 This include Sunset.eps itself.
  (Sunset.C) 20 0 R
  (Sunset.M) 21 0 R
  (Sunset.Y) 22 0 R
  (Sunset.K) 23 0 R ]
endobj

19 0 obj                                An Embedded File stream
<< /Type /EmbeddedFile
  /Length ... /Filter ... >>
stream
...
endstream
endobj

20 0 obj                                Another Embedded File stream
<< /Type /EmbeddedFile
  /Length ... /Filter ... >>
stream
...
endstream
endobj

```

In this example, object 19 is the embedded file stream containing the Mac file **Sunset.eps**, object 13 is the embedded file stream containing the Unix file **Sunset.eps**, and object 14 is the embedded file stream containing the DOS file **SUNSET.EPS**. The **RF** key specifies a set of embedded files, object 11, that are related to the Mac file; this is a DCS 1.0 set. The first two are shown here.

7.4.3 Embedded File streams

In some cases, it is desirable to construct a single PDF file which contains all the information needed to process that file. For example, when there are OPI dictionaries in the PDF file that refer to externally stored high-resolution image data, it may be desirable to embed the high-resolution data in the PDF file for transmission or archiving of the PDF file. Similarly, when a Movie PLayer annotation refers to an external, platform-specific movie file, it may be useful to embed the movie file in the PDF file. An *Embedded File stream* provides for this kind of embedding of external file data within a PDF file. (These embedded files are included simply for convenience and need not be directly processed by any PDF consumer.)

The stream dictionary for an Embedded File contains the standard keys for any stream (e.g., **Length** and **Filter**), plus the additional keys shown in [Table 7.4](#).

Table 7.4 *Additional attributes for EmbeddedFile streams*

Key	Type	Semantics
Type	name	Object type. Always EmbeddedFile .
Subtype	name	(<i>Optional</i>) A name that indicates the subtype of the embedded file. The subtype must be a first-class name as defined in Appendix F . The names without a registered prefix must conform to the MIME Media type names defined in RFC 2046 [27], with the provision that characters not allowed in names must use the two-character hex code format described in Section 4.5 , “Names.”
Params	dictionary	(<i>Optional</i>) File-specific information. See the next section.

7.4.3.1 Embedded File Params dictionary

The **Params** dictionary in an Embedded File stream is used to hold information, other than that explicitly identified in the stream’s dictionary, about an Embedded File. Use of the **Params** dictionary will avoid name conflicts between the keys defined in the stream’s dictionary and keys introduced for a given type of Embedded File. For example, the **Params** dictionary may be used to include keys which would capture information contained in the file system file entry for the file that is being embedded in the EmbeddedFile stream.

The keys defined within the **Params** dictionary are shown in [Table 7.5](#).

Table 7.5 *Attributes for Params dictionary*

Key	Type	Semantics
Mac	dictionary	(<i>Optional</i>) See Section 7.4.3.2 , “Macintosh-specific file information.”
ModDate	Date	(<i>Optional</i>) Date when the file was last modified.
CreationDate	Date	(<i>Optional</i>) Date when the file was created.

Size	integer	<i>(Optional)</i> Size of the Embedded File, in bytes.
Checksum	string	<i>(Optional)</i> 16-byte string that is the checksum of the bytes of the uncompressed embedded file. This checksum is calculated by applying the standard MD5 message digest algorithm over the bytes of the stream.

7.4.3.2 Macintosh-specific file information

The Macintosh-Specific File Information dictionary may be specified as the value of the **Mac** key in the Params dictionary of an Embedded File stream dictionary.

Table 7.6 *Attributes of a Macintosh Specific File Information dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	string	<i>(Optional)</i> Macintosh file type
Creator	string	<i>(Optional)</i> Macintosh file creator.
ResFork	stream	<i>(Optional)</i> The binary contents of the resource fork associated with the Macintosh file that was embedded.

7.4.3.3 Finding all the file specifications in a PDF file

File specifications are references to external files. As such, the processing of a PDF file which contains those references is dependent on the referenced files being consistent with the PDF file. Therefore, it is desirable to enable a PDF file management application to easily determine what files are referenced from the PDF file. This facilitates the following sorts of operations:

PDF 1.3

- Update the relevant file specifications if one or more referenced files get renamed.
- Determine the complete collection of files that need to be copied to a mirror site.
- When making new links to external files, discover existing file specifications referring to the same files and share them.
- Find the file specifications associated with embedded files to be packed or unpacked.

It is not, in general, possible to find all file specifications in a PDF file, because there is no way to determine whether a given string is a file specification string. It is possible, however, to find all the file specification *dictionaries*, if they are indirectly referenced and have a **Type** key with the value **Filespec**. It suffices to traverse the entire cross-reference table (see [page 57](#)) and find all the dictionaries with a **Type** key whose value is **Filespec**.

Note If a PDF file uses a direct object for a file specification dictionary, the PDF file is still valid. However, PDF file-management applications may not be able to locate such file specifications, since they are neither self-typed nor necessarily reachable via any standard path of object references.

7.4.3.4 Finding all the embedded files within a PDF file

Files are embedded either directly using the **EF** dictionary in a File Specification dictionary, or indirectly using the Related Files arrays that are the values in the **RF** dictionary. If the file is embedded using the **EF** dictionary, the name of the file is obtained from the value of the corresponding key in the File Specification dictionary. For example, the **EF** dictionary in [Example 7.2 on page 191](#) has a **DOS** key; the name of the embedded DOS file is the value of the **DOS** key in the File Specification dictionary, **SUNSET.EPS**. If the file is embedded using a Related Files array, the name of the file is the string that precedes the Embedded File stream in this array.

7.4.3.5 Determining whether a file is embedded

A given external file may be referenced from more than one file specification. Therefore, when embedding a file with a given name, it is necessary to check for other occurrences of that name as the value of the same key in other file specifications. This involves finding all the embeddable file specifications and, for each matching key, checking if (1) the string that is the value of that key matches the name of the file being embedded and (2) that a value has not already been embedded for that file specification. (If there is already a corresponding key in the **EF** dictionary, then a file has already been embedded for that use of the file name.) There is no requirement that the files associated with a given file name be unique. The same file name, such as **readme.txt**, may have been associated with different embedded files in distinct file specifications.

7.4.4 URL

When the **FS** key in a file-specification dictionary has the value **URL**, the value of the **F** key in that dictionary is not a file specification string; instead, it is a URL formatted as specified in RFC 1738 and must follow the character-encoding requirements of that RFC. Because 7-bit US ASCII is a strict subset of the **PDFDocEncoding**, this value may also be considered to be in the **PDFDocEncoding**.

Note Protocols most expected to be seen in PDF are “http” and “ftp”.

Example 7.3 URLs

```

/Movie relative URL
<</F (AbbeyRoad.mov) >>

/Movie absolute URL
<</F (/Movies/Beatles/AbbeyRoad.mov) >>

/Movie relative URL

```

```

<</F <</FS /URL /F (AbbeyRoad.mov) >> >>

/Movie                                absolute URL
<</F
  <<
    /FS /URL
    /F (ftp://oranda/ftp/Movies/AbbeyRoad.mov)
  >>
>>

```

7.4.5 Safe path names

Care must be taken to use *safe* path names when creating collections of documents that are to be used on various file systems. A safe path name is one that can be used to locate files on the most common file systems. For maximum compatibility, only a subset of the US ASCII character set should be used: the uppercase alphabetic (**A-Z**) characters, the numeric characters (**0-9**), and the underscore (**_**). The period (**.**) has special meaning as part of a relative path specifier in DOS and Windows file names, and as the first character in a Macintosh pathname. When used in file names, the period should be used only to separate a base file name from a file extension.

Some systems are case-insensitive, so names within a directory should be distinguishable if lowercase letters are changed to uppercase, or vice versa. On DOS and Windows 3.1 systems and on some CD-ROM file systems, file names are limited to eight characters plus a three-character extension. File system software typically converts long names to short names by retaining the first six or seven characters and the first three characters after the last period, if any. The seventh or eighth characters are converted to other values unrelated to the original value. Therefore, safe file names are distinguishable from the first six characters.

7.5 Resources dictionaries

The marking operations for drawing a page are stored in a stream that is the value of the **Contents** key in the Page object's dictionary. In addition to pages, however, there are other objects in PDF that also include streams of marking operations: forms, patterns, and the procedures that draw characters in Type 3 fonts. These are all referred to as *content streams*. [Chapter 8](#) describes the operations used in all content streams, including those in pages.

The marking operations use various types of basic objects such as numbers and strings; these are represented as direct objects in the contents stream. Other objects, such as fonts, that are represented by streams or dictionaries may also be needed by the marking operations, but no indirect objects of any sort, including streams, may appear in a contents stream. Neither may dictionaries, with the exception of in-line property lists (see [Section 7.19 on page 313](#)). Instead, such objects are represented by names, and they are called *named resources*.

Each contents stream includes a list of the named resources it uses. This resource list is stored as a dictionary that is the value of the stream dictionary's **Resources** key, and it serves two functions: it enumerates the named resources in the contents stream, and it establishes the mapping from the names to the objects used by the marking operations.

For example, if a certain font were needed by a text operator, it might be referred to in the contents stream by the name **/F42**. The Resources dictionary would include the mapping from the name **/F42** to the actual font object.

PDF defines several types of named resources. These include:

- ProcSet ([page 198](#))
- Font ([page 198](#))
- Color space ([page 230](#))
- XObject ([page 246](#))
- Extended graphics state ([page 272](#))
- Pattern ([page 287](#))
- Property list ([page 313](#))
- Shading dictionary ([page 294](#))

Other types of objects may be used during marking operations, but they are not referred to directly and are therefore not named. These include:

- Font Encoding ([page 213](#))
- Font descriptor ([page 222](#))
- Halftone ([page 275](#))
- Function ([page 262](#))
- CMap ([page 215](#))

Each key in the Resources dictionary is the name of a *resource type*; each value is a dictionary or an array. If it is a dictionary, its keys are the *resource names*, and its values are indirect references to the PDF objects specifying those resources. If it is an array, it contains a list of names. Only ProcSet resources are represented as arrays in the Resources dictionary; all other resource lists are represented as dictionaries.

The following table shows the attributes of a Resources dictionary.

Table 7.7 *Resources dictionary*

Key	Type	Semantics
ColorSpace	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is either the name of a device-dependent color space or an array describing a color space, as described in Section 7.12 on page 230 .
XObject	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is an XObject. See Section 7.13 on page 246 .
ExtGState	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is an Extended Graphics State dictionary. See Section 7.15 on page 272 .
Font	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is a font. See Section 7.7 on page 198 .
Pattern	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is a pattern. See Section 7.17 on page 287 .
ProcSet	array	<i>(Optional)</i> An array of predefined names. See Section 7.6 on page 198 .
Properties	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is a property list. See Section 7.19 on page 313 .
Shading	dictionary	<i>(Optional)</i> A dictionary in which each key is a resource name, and each value is a shading dictionary. See Section 7.18 on page 294 .

PDF 1.3

[Example 7.4](#) shows a Resources dictionary containing ProcSets, Fonts, and XObjects. The ProcSets are specified by an array, as described in the following section. The Fonts are specified with a dictionary; it contains four names, **/F5**, **/F6**, **/F7**, and **/F8**, and these are associated with objects 6, 8, 10, and 12, respectively, which are fonts. Likewise, the XObject dictionary contains the names **/Im1** and **/Im2**, associated with objects 13 and 15, respectively, which are XObjects.

Example 7.4 *Resources dictionary*

```

<<
  /ProcSet [/PDF /ImageB]
  /Font
    << /F5 6 0 R /F6 8 0 R /F7 10 0 R /F8 12 0 R >>
  /XObject
    << /Im1 13 0 R /Im2 15 0 R >>
>>

```

Note The phrase “current Resources dictionary” refers to the Resources dictionary of the current contents stream. Since Page objects can inherit their attributes (see [page 72](#)) the current Resources dictionary for a Page may be stored in some ancestor of the Page object.

7.6 ProcSets

The types of instructions that may be used in a PDF contents stream are grouped into independent sets of related instructions. Each of these sets, called a ProcSet, may or may not be used in a particular stream. ProcSets contain implementations of the PDF operators and are used only when a page or other stream is printed. The Resources dictionary for each contents stream must contain a **ProcSet** key whose value is an array consisting of the names of the ProcSets used in that contents stream. Each of the entries in the array must be one of the predefined names shown in [Table 7.8](#). For an example of a **ProcSet** key, see the Resources dictionary in [Example 7.4](#).



Table 7.8 *Predefined ProcSets*

<i>ProcSet Name</i>	<i>Required if the page has any...</i>
/PDF	marks in the contents stream whatsoever
/Text	text
/ImageB	grayscale images or image masks
/ImageC	color images
/ImageI	indexed images (also called color-table images)

7.7 Fonts

A font is represented in PDF as a dictionary specifying the type of font, its PostScript font name, its encoding, and information that can be used to provide a substitute when the font is not available. PDF defines the following types of fonts:

- Type 1, including subsets and Multiple Master “snapshots”
- Type 3
- TrueType, including subsets
- Type 0

PDF 1.2

Any font may include a **Name** attribute; this is required only in PDF 1.0. The name is used as the operand of the **Tf** operator when selecting the font. If **Name** is supplied, it should match the name used in the Font dictionary within the current Resources dictionary.



All fonts specify the *width* of the characters in the font. The width of a character refers to the horizontal offset between the origin of the character and the origin of the next character when writing in horizontal mode, as shown in [Figure 7.1](#). In horizontal mode, the vertical offset is always 0.

Note In a PDF document, the character widths and metrics for Type 1, Multiple Master Type 1, True Type, and Type 0 font resources are measured in units in which 1000 units corresponds to 1 unit in text space.

Figure 7.1 Character metrics




Figure 7.2 shows the relationship between fonts, encodings, CMaps, and descriptors. Three types of fonts are shown: Type 1, TrueType, and Type 0. Note that the descriptors use different keys to refer to an embedded font file. (For illustration purposes only, the figure shows nested direct objects instead of indirect objects.)

The term *glyph* is used to refer to an element of a font, particularly in a multi-byte font, where the term *character*, which is the more common term, might be confused with *byte*.

Table 7.9 Attributes common to Type 1, Type 3, True Type and Type 0 fonts

Key	Type	Semantics
Type	name	Object type. The value must be Font .
Subtype	name	(Required) Type of font.
BaseFont	name	(Required) The “PostScript name” of the font, as described on page 201. There are special cases for True Type fonts, Multiple Master fonts, and font subsets.
Encoding	various	A map from character codes to glyph procedures.
ToUnicode	stream	(Optional) A mechanism for mapping character codes to Unicode values. See Section 7.10.1 on page 219.

PDF 1.3

7.7.1 Type 1 fonts

Type 1 fonts, described in detail in *Adobe Type 1 Font Format* [5], are special-purpose PostScript language programs used for defining fonts. As compared to Type 3 fonts, Type 1 fonts can be defined more compactly, make use of a special procedure for drawing the characters that results in higher quality output at small

Figure 7.2 *Fonts, encodings, CMaps, and FontDescriptors*

```

<< /Type      /Font
   /Subtype   /Type1
   /Encoding  << /Type /Encoding
                 /Differences [ ... ]
                 /BaseEncoding ... >>
   /FirstChar ...
   /LastChar ...
   /FontDescriptor << /Type /FontDescriptor
                     /Flags ...
                     ...
                     /FontFile stream >>
... >>

<< /Type      /Font
   /Subtype   /TrueType
   /Encoding  << /Type /Encoding
                 /Differences [ ... ]
                 /BaseEncoding ... >>
   /FirstChar ...
   /LastChar ...
   /FontDescriptor << /Type /FontDescriptor
                     /Flags ...
                     ...
                     /FontFile2 stream >>
... >>

<< /Type      /Font
   /Subtype   /Type0
   /Encoding  << /Type /CMap
                 /CIDSystemInfo ...
                 ... >>
   /DescendantFonts [ << /Type /Font
                       /Subtype /CIDFontType0
                       /CIDSystemInfo ...
                       /FontDescriptor << /Type /FontDescriptor
                                           /Flags ...
                                           ...
                                           /FontFile3 stream >>
                       /BaseFont ... >>
                       ... ]
... >>

```


sizes and low resolution, and have a built-in mechanism for specifying hints, which are data that indicate basic features of the character shapes not directly expressible by the basic PostScript language operators.

[Table 7.10](#) shows the attributes of Type 1 fonts.

Table 7.10 *Type 1 font attributes (in addition to those in [Table 7.9](#))*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	name	<i>(Required)</i> Type of font. Always Type1 .
BaseFont	name	<i>(Required)</i> The name of the base font, described below.
FirstChar	integer	<i>(Required except for base 14 Type 1 fonts listed in Table 7.11)</i> Specifies the first character code defined in the font's Widths array.
LastChar	integer	<i>(Required except for base 14 Type 1 fonts)</i> Specifies the last character code defined in the font's Widths array.
Widths	array	<i>(Required except for base 14 Type 1 fonts; indirect reference preferred)</i> An array of (LastChar - FirstChar + 1) widths. For character codes outside the range FirstChar to LastChar , the value of MissingWidth from the font's descriptor is used (see page 222). The units in which character widths are measured depend on the type of font.
Encoding	dictionary	<i>(Optional)</i> An Encoding dictionary that specifies the font's character encoding. If this key is not present, the font's built-in encoding is used. Appendix C describes the predefined encodings (MacRomanEncoding , MacExpertEncoding , and WinAnsiEncoding).
	<i>Note</i>	<i>To display correctly, fonts that do not use a predefined encoding must be embedded in the PDF file.</i>
FontDescriptor	dictionary	<i>(Required except for base 14 fonts; must be indirect reference)</i> A font descriptor describing the font's metrics other than its character widths.

BaseFont is the “PostScript name” of the font, that is, the name used by the **definefont** and **findfont** operators in a PostScript language program. This is usually the value of the **FontName** key in the PostScript language font dictionary of the font. For more information, see Section 5.2 of the *PostScript Language Reference Manual, Third Edition* [1].

[Example 7.5](#) shows the font dictionary for the Adobe Garamond Semibold font. The font has an encoding (object 25), although neither the encoding nor the font descriptor (object 7) is shown in the example.

Example 7.5 *Type 1 font and character widths array*

```
14 0 obj <<
  /Type /Font
```

```

/Subtype /Type1
/BaseFont /AGaramond-Semibold
/Encoding 25 0 R
/FontDescriptor 7 0 R
/FirstChar 0
/LastChar 255
/Widths 21 0 R
>> endobj

21 0 obj
[255 255 255 255 255 255 255 255 255 255 255 255 255
255 255 255 255 255 255 255 255 255 255 255 255 255
255 255 255 255 255 255 255 280 438 510 510 868 834
248 320 320 420 510 255 320 255 347 510 510 510 510
510 510 510 510 510 510 255 255 510 510 510 330 781
627 627 694 784 580 533 743 812 354 354 684 560 921
780 792 588 792 656 504 682 744 650 968 648 590 638
320 329 320 510 500 380 420 510 400 513 409 301 464
522 268 259 484 258 798 533 492 516 503 349 346 321
520 434 684 439 448 390 320 255 320 510 255 627 627
694 580 780 792 744 420 420 420 420 420 420 402 409
409 409 409 268 268 268 268 533 492 492 492 492 492
520 520 520 520 486 400 510 510 506 398 520 555 800
800 1044 360 380 549 846 792 713 510 549 549 510 522
494 713 823 549 274 354 387 768 615 496 330 280 510
549 510 549 612 421 421 1000 255 627 627 792 1016
730 500 1000 438 438 248 248 510 494 448 590 100 510
256 256 539 539 486 255 248 438 1174 627 580 627 580
580 354 354 354 354 792 792 790 792 744 744 744 268
380 380 380 380 380 380 380 380 380 380]
endobj

```

7.7.2 The base 14 Type 1 fonts

[Table 7.11](#) lists the PostScript names of fourteen Type 1 fonts known as the “base 14 fonts.” These fonts include members of the Courier, Helvetica*, and Times* families, along with Symbol and ITC Zapf Dingbats. These fonts, or their font metrics and suitable substitution fonts, are guaranteed to be available to the Acrobat program. For more information on font metrics, see *Adobe Font Metrics File Format Specification* [13].

Table 7.11 *Base 14 fonts*

Courier
Courier-Bold
Courier-BoldOblique
Courier-Oblique
Helvetica
Helvetica-Bold
Helvetica-BoldOblique
Helvetica-Oblique
Times-Roman
Times-Bold
Times-Italic
Times-BoldItalic
Symbol
ZapfDingbats

7.7.3 TrueType fonts

The TrueType font format was developed by Apple Computer. A TrueType font, shown in [Table 7.12](#), has the same keys as a Type 1 font. The **Subtype** and **BaseFont** keys have different values than those in a Type 1 font.

Note Character widths in TrueType fonts are measured in units in which 1,000 units corresponds to 1 unit in text space.

Table 7.12 *TrueType font attributes (in addition to those in [Table 7.9](#))*

Key	Type	Semantics
Subtype	name	<i>(Required)</i> Type of font. Always TrueType .
BaseFont	name	<i>(Required)</i> A name containing a “style” specifying the base TrueType font, described below.
FirstChar	integer	<i>(Required)</i> The same as for Type 1.
LastChar	integer	<i>(Required)</i> The same as for Type 1.
Widths	array	<i>(Required; indirect reference preferred)</i> The same as for Type 1.
Encoding	dictionary	<i>(Optional)</i> The same as for Type 1.
FontDescriptor	dictionary	<i>(Required; must be indirect reference)</i> The same as for Type 1.

For TrueType fonts, the **BaseFont** key in the font dictionary is a name that may contain a *style*. If the font is a bold, italic, or bold italic font for which no PostScript name is available, the **BaseFont** key contains the base name of the

font with any spaces removed, followed by a comma, followed by a style. The style is one of *Italic*, *Bold*, or *BoldItalic*. For example, the italic variant of the New York font has a **BaseFont** that is written as

```
/NewYork,Italic
```

Example 7.6 *TrueType font*

```
17 0 obj <<
  /Type /Font
  /Subtype /TrueType
  /BaseFont /NewYork,Bold
  /FirstChar 0
  /LastChar 255
  /Widths 23 0 R
  /Encoding /MacRomanEncoding
  /FontDescriptor 7 0 R
>> endobj
```



```
23 0 obj
[0 333 333 333 333 333 333 333 0 333 333 333 333 333
333 333 333 333 333 333 333 333 333 333 333 333 333
333 333 0 333 333 333 303 500 666 666 882 848 303
446 446 507 666 303 378 303
```

... omitted data ...

```
303 530 1280 757 605 757 605 605 355 355 355 355 803
803 790 803 780 780 780 340 636 636 636 636 636 636
636 636 636 636]
endobj
```

7.7.4 Font Subsets

PDF 1.1

PDF permits documents to include *subsets* of Type 1 and TrueType fonts. The font and the descriptor that describe a font subset are slightly different from those of ordinary fonts. These differences allow an application to recognize font subsets and to merge documents containing different subsets of the same font.

The value of the font's **BaseFont** key and the font descriptor's **FontName** key use the following format:

pseudoUniqueTag+PostScriptName

pseudoUniqueTag consists of exactly six uppercase alphabetic characters. *PostScriptName* must be the name of the complete Type 1 font. A plus sign separates *pseudoUniqueTag* and *PostScriptName*. For example, **EODIA+Poetica** is the name of a subset of the Poetica[®] font. The purpose of the tag is to identify the subset. Different subsets should have different tags.

Note Both the **BaseFont** and the **FontName** must use this format for subset fonts.

7.7.5 Multiple master Type 1 fonts

The multiple master font format is an extension of the Type 1 font format that allows the generation of a wide variety of typeface styles from a single font. This is accomplished through the presence of various design dimensions in the font. Examples of design dimensions are weight (light to extra-bold) and width (condensed to expanded). Coordinates along these design dimensions (such as the degree of boldness) are specified by numbers.

To specify the appearance of the font, numeric values must be supplied for each design dimension of the multiple master font. A completely specified multiple master font is referred to as a *snapshot* of the multiple master font.

The note *Type 1 Font Format Supplement* [6] describes multiple master fonts. A snapshot of a multiple master font, shown in [Table 7.13](#), has the same keys as an ordinary Type 1 font.

Note Character widths in multiple master Type 1 fonts are measured in units in which 1,000 units corresponds to 1 unit in text space.

Table 7.13 Multiple master Type 1 font attributes (in addition to those in [Table 7.9](#))

Key	Type	Semantics
Subtype	name	(Required) Type of font. Always MMType1 .
BaseFont	name	(Required) Specifies the PostScript name of the snapshot. If the name contains spaces (such as “MinionMM 366 465 11”), these spaces are replaced with underscores (e.g., /MinionMM_366_465_11).
FirstChar	integer	(Required) The same as for Type 1.
LastChar	integer	(Required) The same as for Type 1.
Widths	array	(Required; indirect reference preferred) The same as for Type 1.
FontDescriptor	dictionary	(Required; must be indirect reference) The same as for Type 1.
Encoding	dictionary	(Optional) The same as for Type 1.

Example 7.7 Multiple master font and character widths array

```
7 0 obj <<
  /Type /Font
  /Subtype /MMType1
  /BaseFont /MinionMM_366_465_11
  /FirstChar 32
  /LastChar 255
  /Widths 19 0 R
```

```

/Encoding 5 0 R
/FontDescriptor 6 0 R
>>
endobj

19 0 obj
[187 235 317 430 427 717 607 168 326 326 421
619 219 317 219 282 427 427 427 427 427 427
427 427 427 427 219 219 619 619
... omitted data ...
301 301 301 569 569 0 569 607 607 607 239 400
400 400 400 253 400 400 400 400 400]
endobj

```

7.7.6 Type 3 fonts

Type 3 fonts, also known as user-defined fonts, are described in Section 5.7 of the *PostScript Language Reference Manual, Third Edition* [1]. PDF provides a variant of Type 3 fonts in which characters are defined by streams of PDF page-marking operators. These streams, known as CharProcs, are associated with the character names. As with any font, the character names are accessed via an encoding vector.

PDF Type 3 fonts differ from the other fonts supported by PDF. A Type 3 font defines the font itself, while the other font dictionaries simply contain information *about* the font.

Type 3 fonts are more flexible than Type 1 fonts because the character-drawing streams may contain arbitrary PDF marking operators. However, Type 3 fonts have no mechanism for improving output at small sizes or low resolutions, and no built-in mechanism for hinting. Table 7.14 shows the attributes specific to Type 3 fonts.

Table 7.14 *Type 3 font attributes (in addition to those in Table 7.9)*

Key	Type	Semantics
Subtype	name	(Required) Type of font. Always Type3 .
FirstChar	integer	(Required) The same as for Type 1.
LastChar	integer	(Required) The same as for Type 1.
Widths	array	(Required; indirect reference preferred) The same as for Type 1.
CharProcs	dictionary	(Required) Each key in this dictionary is a character name; the value associated with that key is a stream object that draws the character. Any operator that can be used in a PDF page description can be used in this stream. However, the stream must include as its first operator either d0 (d zero) or d1 (d one). ¹

1. These are equivalent to the PostScript language **setcharwidth** and **setcachedevice** operators.

FontBBox	Rectangle	<i>(Required)</i> The font's bounding box. The coordinates are measured in character space. The bounding box is the smallest rectangle enclosing the shape that would result if all characters in the font were placed with their origins coincident, and then painted. FontBBox is identical to the PostScript Type 3 font FontBBox.
FontMatrix	array	<i>(Required)</i> Specifies the transformation from character space to text space. FontMatrix is identical to the PostScript Type 3 font FontMatrix.
Resources	dictionary	<i>(Optional, but strongly recommended)</i> A list of the named resources such as fonts and images required by the CharProcs in this font (see page 195). If this dictionary is not present, then the resources used by the CharProcs are specified in the Resources dictionary of the Page on which this font is used.
Encoding	dictionary	<i>(Required)</i> The same as for Type 1.

PDF 1.2



When a PDF viewer needs to find the CharProc for a glyph in a Type 3 font, it must first look in the **Encoding** array to find the glyph name corresponding to the character code, then look up the CharProc for that glyph name in the **CharProcs** dictionary. For example, if the **Encoding** array says `[0 /h ...]`, then the **CharProcs** dictionary must contain an entry for the key `/h`.

Example 7.8 Type 3 font

```

6 0 obj <<
  /Type /Font
  /Subtype /Type3
  /Name /T36
  /CharProcs 1928 0 R
  /FontBBox [-3 -241 875 856]
  /FontMatrix [.001 0 0 .001 0 0]
  /FirstChar 3
  /LastChar 101
  /Widths 7 0 R
  /Encoding 1927 0 R
>> endobj

7 0 obj
[55 0 0 589 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
640 40 0 40 0 40 40 0 0 0 0 0 0 0 0 60 0 58 61 54 52
603 0 29 0 0 853 73 60 62 504 0 659 44 58 60 60 0 0
603 0 0 0 0 0 0 0 0 0 0 35 0 35]
endobj

```

7.7.7 Type 0 Fonts

PDF 1.2

Type 0 fonts support single-byte or multi-byte encodings and can refer to one or more *descendant* fonts; see [Figure 7.2 on page 200](#). These fonts are analogous to the Type 0 or composite fonts supported by PostScript interpreters. However PDF

Type 0 fonts only support character encodings defined by a CMap, which maps character codes to glyphs in the descendant fonts. The details of a CMap are defined in [Section 7.10 on page 215](#).

The following table shows the attributes of a Type 0 font.

Table 7.15 *Type 0 font attributes (in addition to those in [Table 7.9](#))*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	name	<i>(Required)</i> Type of font. Always Type0 .
BaseFont	name	<i>(Required)</i> The PostScript name of the font, or a name with a style specifying the TrueType font.
DescendantFonts	array	<i>(Required)</i> An array of one or more fonts. These fonts are selected by the font number defined in the CMap.
Encoding	name or stream	<i>(Required)</i> The name of a predefined CMap, or a stream containing a CMap.

Implementation note

*Acrobat 3.0 and Acrobat 4.0 support only one entry in the **DescendantFonts** array.*

Type 0 fonts can also be used to embed fonts where the characters are encoded using an unknown character encoding. An example of this is when characters are encoded using glyphIndex values or CID values from an unknown character collection. GlyphIndex and CID values are two-byte values that are used to directly index glyph procedures in **CIDFontType2** fonts (TrueType, see [Section 7.8.2 on page 212](#)) or **CIDFontType0** fonts (see [Section 7.8.1 on page 211](#)).

The following example shows a Type 0 font which refers to a single CIDFont. The CMap used is one of the predefined CMaps and is referenced by name.

Example 7.9 *Type 0 font referring to a single CIDFont*

```
14 0 obj <<
  /Type /Font
  /Subtype /Type0
  /BaseFont /HeiseiMin-W5-90ms-RKSJ-H
  /Encoding /90ms-RKSJ-H
  /DescendantFonts [15 0 R]
>> endobj
```


7.7.8 Vertical Writing

Vertical writing is specified by the **WMode** entry in a CMap. When this value is 1, the Type 0 font that uses the CMap will position text for vertical writing.

In vertical writing, the character position is described by a *position vector* from the origin used for horizontal writing (origin 0) to the origin used for vertical writing (origin 1). See the *PostScript Language Reference Manual, Third Edition* [1], section 5.4, for a detailed explanation of vertical writing.

Figure 7.3 illustrates the metrics for horizontal and vertical writing modes.

Figure 7.3 Horizontal and vertical writing metrics



The left diagram illustrates the character metrics associated with writing mode 0, horizontal writing. The coordinates ll and ur specify the bounding box of the character relative to origin 0. $w0$ is the displacement vector that specifies how the current point is changed after the character is shown in writing mode 0. Its x component (horizontal displacement or *character width*) may be different from the width of the bounding box. Its y component (vertical displacement) is always 0.

The center diagram illustrates writing mode 1, vertical writing. $w1$ is the displacement vector for writing mode 1; its x component is always 0.

In the right diagram, v is the *position vector* that defines the position of origin 1 relative to origin 0.

In vertical writing, the default position vector and displacement vector are specified by the **DW2** attribute, which is an array of two values: the y component of the position vector, and the y component of the displacement vector. The x value of the position vector is always half the character width. The x component of the displacement vector is always 0. For example, if the **DW2** entry is:

```
/DW2 [880 -1000]
```

the position vector and displacement vectors are:

$$v = (hw \div 2, 880)$$

$$w = (0, -1000)$$

where hw is the character width. Note that a negative value for the y component will place the origin of the next character *below* the current character, because in a standard coordinate system, the positive direction of the y -axis points upward.

Glyphs whose vertical metrics differ from the defaults must be included in the **W2** array. In this array, the position vector and displacement vector are defined by three numbers: the x and y values of the position vector, and the y component of the displacement vector. There are two formats that can be used to define these vectors, as shown below:

```
/W2 [
c [W1,y V1,x V1,y W2,y V2,x V2,y ...]
cfirst clast W1,y V1,x V1,y
]
```

In the first entry, c gives a starting CID and is followed by an array. The array contains sets of three numbers: the y value for the displacement vector, followed by the x and y values for the position vector. The sets of three numbers defines the vertical metrics for consecutive character codes starting with CID value c . The second format defines a range of CIDs from c_{first} to c_{last} . These ranges use the same format as defined for the **W** array. The range is followed by three numbers that define the vertical metrics for all CIDs in the range.

Following is an example of a **W2** entry.

```
/W2
[
120 [-1000 250 772]
7080 8032 -1000 500 900
]
```

This entry defines the writing mode 1 (vertical) displacement vector for character CID 120 as (0,-1000) and the position vector as (250, 772). The second list in the array defines the displacement vector to be (0, -1000) and the position vector of (500, 900), for CIDs in the range 7080 to 8032.

7.8 CIDFonts

PDF 1.2

A CIDFont is a font-like object designed to contain a large number of glyph procedures. Instead of being accessed by a name, each glyph procedure is accessed by an integer known as a *character identifier* or *CID*.

Although its **Type** is **Font**, a CIDFont is not actually a font; it does not have an **Encoding** attribute, and it cannot be used as the operand to the **Tf** operator. It is used only as a descendant of a Type 0 font. It is the CMap in the Type 0 font that defines the encoding that maps character codes to CIDs in the CIDFont.

There are two types of CIDFonts. A Type 0 CIDFont contains glyph procedures based on Adobe's Type 1 font format; a Type 2 CIDFont contains glyph procedures based on the TrueType font format.

7.8.1 CIDFontType 0

A Type 0 CIDFont uses Adobe Type 1 charstrings for glyph procedures and the CIDFont file format.

Table 7.16 *CIDFontType 0 attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always Font .
Subtype	name	<i>(Required)</i> Always CIDFontType0 .
BaseFont	name	<i>(Required)</i> The PostScript name of the CIDFont.
CIDSystemInfo		
	dictionary	<i>(Required)</i> A dictionary containing the Registry, Ordering, and Supplement attributes that define the character collection of the CIDFont. See Table 7.17 .
FontDescriptor		
	dictionary	<i>(Required; must be an indirect reference)</i> A font descriptor describing the CIDfont's default metrics other than its character widths.
DW	integer	<i>(Optional)</i> Default width for glyphs in the CIDFont.
W	array	<i>(Optional)</i> The W array consists of a set of lists that define the widths for the glyphs in the CIDfont. Each list can specify individual widths for consecutive CIDs, or one width for a range of CIDs. (See the section below for more details on this format.)
DW2	array	<i>(Optional; applies only to CIDFonts that are used for vertical writing)</i> The default metrics for writing mode 1 (see Section 7.7.8, "Vertical Writing," for a description of writing mode 1 and examples of the DW2 key). This entry is an array of two numbers: the y component of the position vector and the y component of the displacement vector for writing mode 1. The x component of the position vector is always half the width of the character. The x component of the displacement vector is always 0. The default value is [880 -1000] .
W2	array	<i>(Optional; applies only to CIDFonts that are used for vertical writing)</i> This array defines the metrics for vertical writing (see Section 7.7.8, "Vertical Writing," for a description of vertical writing and examples of the W2 key). Its format is similar to the format of the W array. It defines the x and y components of the position vector, and the y component of the displacement vector for individual CID values. The x component of the displacement vector is always 0. The contents of this array is described in detail below.

Table 7.17 *Entries in a CIDSystemInfo dictionary*

<i>Key</i>	<i>Type</i>	<i>value</i>
Registry	string	<i>(Required)</i> A string identifying an issuer of character collections—for example, Adobe . For information about assigning a registry identifier, consult the Adobe Developer Relations Web site (see the Bibliography) or send e-mail to the Adobe Unique ID Coordinator at fontdev-person@adobe.com .
Ordering	string	<i>(Required)</i> A string that uniquely names a character collection issued by a specific registry—for example, Japan1 .
Supplement	integer	<i>(Required)</i> The <i>supplement number</i> of the character collection. An original character collection has a supplement number of 0. Whenever additional CIDs are assigned in a character collection, the supplement number is increased. Supplements do not alter the ordering of existing CIDs in the character collection. This value is not used in determining compatibility between character collections.

7.8.2 CIDFontType 2

A Type 2 CIDFont uses TrueType glyph procedures.

Table 7.18 *CIDFontType 2 attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always Font .
Subtype	name	<i>(Required)</i> Always CIDFontType2 .
BaseFont	name	<i>(Required)</i> A name with a style specifying the TrueType font.
CIDToGIDMap	stream or name	<i>(Optional)</i> If the value of this key is a stream, the bytes in the stream contain the mapping from CID to glyphindex (“GID”). The glyphindex for a particular CID value <i>c</i> is a 2-byte value stored in bytes <i>2c</i> and <i>2c+1</i> ; the first byte is the high-order byte. If the value of this key is a name, it must be Identity , indicating that the mapping between CIDs and glyphindices is the identity mapping. The default is the name Identity .
CIDSystemInfo	dictionary	<i>(Required)</i> The same as for CIDFontType 0. See Table 7.17 .
FontDescriptor	dictionary	<i>(Required; must be an indirect reference)</i> The same as for CIDFontType 0.
DW	integer	<i>(Optional)</i> Default width for glyphs in the CIDFont.
W	array	<i>(Optional)</i> The same as for CIDFontType 0.
DW2	array	<i>(Optional; applies only to CIDFonts that are used for vertical writing)</i> The same as for CIDFontType 0.

W2 array (Optional; applies only to CIDFonts that are used for vertical writing) The same as for CIDFontType 0.

7.8.3 Character widths in CIDFonts

In any font, the *width* of a character (glyph) refers to the horizontal displacement between the origin of the character and the origin of the next character when writing in horizontal mode. In this mode, the vertical displacement between origins is always 0.

Widths for a CIDFont are defined using the **DW** and the **W** attributes. The **DW** entry defines the default width. This entry is particularly useful for Chinese, Japanese, and Korean fonts, where many of the characters have the same width. The **W** attribute allows the definition of widths for single CIDs or a range of CIDs. The entry is an array of lists in either of the following two formats:

$$C \ [W_1 \ W_2 \ \dots \ W_n]$$

$$C_{\text{first}} \ C_{\text{last}} \ W$$

In the first list, C is an integer specifying a starting CID value. It is followed by an array of n numbers which specify the widths for n consecutive CIDs, starting with C . The second list defines the same width, w , for all the CIDs in the range C_{first} to C_{last} . If the **DW** key is present, the **W** attribute need not specify the width of a CID that uses the default width.

Following is an example of a **W** attribute.

```
/W
[
120 [400 325 500]
7080 8032 1000
]
```

In this example, the CIDs 120, 121, and 122 have widths 400, 325, and 500 units respectively. The second list in this example specifies that CIDs in the range 7080 to 8032 have a width of 1,000 units.

See [Section 7.7.8, “Vertical Writing](#) for a description of the **W2** and **DW2** keys.

7.9 Font encodings

An encoding describes a font’s character encoding, the mapping between numeric character codes and character names. These character names are keys in the font dictionary and are used to retrieve the code which draws the character. Thus, the font encoding provides the link which associates numeric character codes with the glyphs drawn when those codes are encountered in text. For details on the mapping from character codes to glyphs, see Section 5.11.5 in *PostScript Language Reference Manual, Third Edition* [1].

An encoding is represented as a dictionary whose contents are shown in [Table 7.19](#).

Table 7.19 *Font encoding attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always Encoding .
BaseEncoding	name	<i>(Optional)</i> Specifies the encoding from which the new encoding differs. This key is not present if the encoding is based on the base font's encoding. Otherwise it must be one of the predefined encodings MacRomanEncoding , MacExpertEncoding , or WinAnsiEncoding , described in Appendix C .
Differences	array	<i>(Optional)</i> Describes the differences from the base encoding.

The value of the **Differences** key is an array of character codes and glyph names organized as follows:

```
code1 name1,1 name1,2 ...
code2 name2,1 name2,2 ...
...
coden namen,1 namen,2 ...
```

Each code is the first index in a sequence of characters to be changed. The first glyph name after the code becomes the name corresponding to that code. Subsequent names replace consecutive code indexes until the next code appears in the array or the array ends.

For example, in the encoding in [Example 7.10](#), the glyph **quotesingle** (') is associated with character code 39. **Adieresis** (Ä) is associated with code 128, **Aring** (Å) with 129, and **trademark** (™) with 170.

Example 7.10 *Font encoding*

```
25 0 obj
<<
/Type /Encoding
/Differences [39 /quotesingle 96 /grave 128
/Adieresis /Aring /Ccedilla /Eacute /Ntilde
/Odieresis /Udieresis /acute /grave
/acircumflex /adieresis /atilde /aring /ccedilla
/eacute /egrave /ecircumflex /edieresis /iacute
/igrave /icircumflex /idieresis /ntilde /oacute
/ograve /ocircumflex /odieresis /otilde /uacute
/ugrave /ucircumflex /udieresis /dagger /degree
/cent /sterling /section /bullet /paragraph
/germandbls /registered /copyright /trademark
/acute /dieresis 174 /AE /Oslash 177 /plusminus 180
/yen /mu 187/ordfeminine /ordmasculine 190
```

```

/ae /oslash /questiondown /exclamdown /logicalnot
196 /florin 199 /guillemotleft /guillemotright
/ellipsis 203 /Agrave /Atilde /Otilde /OE /oe
/endash /emdash /quotedblleft /quotedblright
/quoteleft /quoteright /divide 216 /ydieresis
/Ydieresis /fraction /currency /guilsinglleft
/guilsinglright /fi /fl /daggerdbl
/periodcentered /quotesinglbase /quotedblbase
/perthousand /Acircumflex /Ecircumflex /Aacute
/Edieresis /Egrave /Iacute /Icircumflex
/Idieresis /Igrave /Oacute /Ocircumflex 241
/Ograve /Uacute /Ucircumflex /Ugrave /dotlessi
/circumflex /tilde /macron /breve /dotaccent
/ring /cedilla /hungarumlaut /ogonek /caron]
>>
endobj

```

7.10 CMaps

PDF 1.2

A CMap is used with a Type 0 composite font to define the mapping from character codes to a font number and a character selector. (A CMap serves the same function for a Type 0 font as an Encoding does for a Type 1 font.) The font number selects the font from the **DescendantFonts** array in the Type 0 font. The character selector selects the glyph from the descendant font. A character selector can be a CID or a character code. A CMap can handle encodings that use single-byte or multi-byte character codes.

Implementation note

In Acrobat 3.0 and 4.0, Type 0 fonts may contain only one descendant font, so the CMaps must always map character codes to font number 0.

The following table shows the predefined CMap names. These CMaps map character codes to CIDs in a single descendant CIDFont.

Table 7.20 *Predefined CJK CMap names*

<i>Name</i>	<i>Description</i>
	<i>Chinese (Simplified)</i>
GB-EUC-H	Microsoft Code Page 936 (IfCharSet 0x86), GB 2312-80 character set, EUC-CN encoding
GB-EUC-V	Vertical version of GB-EUC-H
GBpc-EUC-H	Macintosh, GB 2312-80 character set, EUC-CN encoding, Script Manager code 2
GBpc-EUC-V	Vertical version of GBpc-EUC-H
GBK-EUC-H	Microsoft Code Page 936 (IfCharSet 0x86), GBK character set, GBK encoding
GBK-EUC-V	Vertical version of GBK-EUC-H
UniGB-UCS2-H	Unicode (UCS-2) encoding for the Adobe-GB1 character collection
UniGB-UCS2-V	Vertical version of UniGB-UCS2-H .
	<i>Chinese (Traditional)</i>
B5pc-H	Macintosh, Big Five character set, Big Five encoding, Script Manager code 2

B5pc-V	Vertical version of B5pc-H
ETen-B5-H	Microsoft Code Page 950 (IfCharSet 0x88), Big Five character set with ETen extensions
ETen-B5-V	Vertical version of ETen-B5-H
ETenms-B5-H	Microsoft Code Page 950 (IfCharSet 0x88), Big Five character set with ETen extensions; this uses proportional forms for half-width Latin characters.
ETenms-B5-V	Vertical version of ETenms-B5-H
CNS-EUC-H	CNS 11643-1992 character set, EUC-TW encoding
CNS-EUC-V	Vertical version of CNS-EUC-H
UniCNS-UCS2-H	Unicode (UCS-2) encoding for the Adobe-CNS1 character collection
UniCNS-UCS2-V	Vertical version of UniCNS-UCS2-H .
<i>Japanese</i>	
83pv-RKSJ-H	Macintosh, JIS X 0208 character set with KanjiTalk6 extensions, Shift-JIS encoding, Script Manager code 1
90ms-RKSJ-H	Microsoft Code Page 932 (IfCharSet 0x80), JIS X 0208 character set with NEC and IBM extensions
90ms-RKSJ-V	Vertical version of 90ms-RKSJ-H
90msp-RKSJ-H	Same as 90ms-RKSJ-H , but replaces half-width Latin characters with proportional forms
90msp-RKSJ-V	Vertical version of 90msp-RKSJ-H
90pv-RKSJ-H	Macintosh, JIS X 0208 character set with KanjiTalk7 extensions, Shift-JIS encoding, Script Manager code 1
Add-RKSJ-H	JIS X 0208 character set with Fujitsu FMR extensions, Shift-JIS encoding
Add-RKSJ-V	Vertical version of Add-RKSJ-H
EUC-H	JIS X 0208 character set, EUC-JP encoding
EUC-V	Vertical version of EUC-H
Ext-RKSJ-H	JIS C 6226 (JIS78) character set with NEC extensions, Shift-JIS encoding
Ext-RKSJ-V	Vertical version of Ext-RKSJ-H
H	JIS X 0208 character set, ISO-2022-JP encoding
V	Vertical version of H
UniJIS-UCS2-H	Unicode (UCS-2) encoding for the Adobe-Japan1 character collection
UniJIS-UCS2-V	Vertical version of UniJIS-UCS2-H
UniJIS-UCS2-HW-H	Same as UniJIS-UCS2-H , but replaces proportional Latin characters with half-width forms
UniJIS-UCS2-HW-V	Vertical version of UniJIS-UCS2-HW-H
<i>Korean</i>	
KSC-EUC-H	KS X 1001:1992 character set, EUC-KR encoding
KSC-EUC-V	Vertical version of KSC-EUC-H
KSCms-UHC-H	Microsoft Code Page 949 (IfCharSet 0x81), KS X 1001:1992 character set plus 8,822 additional hangul, Unified Hangul Code (UHC) encoding
KSCms-UHC-V	Vertical version of KSCms-UHC-H
KSCms-UHC-HW-H	Same as KSCms-UHC-H , but replaces proportional Latin characters with half-width forms
KSCms-UHC-HW-V	Vertical version of KSCms-UHC-HW-H
KSCpc-EUC-H	Macintosh, KS X 1001:1992 character set with MacOS-KH extensions, Script Manager Code 3
UniKS-UCS2-H	Unicode (UCS-2) encoding for the Adobe-Korea1 character collection
UniKS-UCS2-V	Vertical version of UniKS-UCS2-H

Generic

Identity-H	The horizontal identity mapping for two-byte CIDs. This may be used with CIDFonts using any Registry, Ordering and supplement. It maps two-byte character codes from 1 to 65536 to the same two-byte CID value.
Identity-V	Vertical version of the Identity-H mapping. The mapping is the same as for Identity-H . However the writing mode (WMode) value is set to 1 to indicate vertical writing.

For character encodings that are not predefined, the PDF file must contain a stream that defines the CMap. The stream dictionary contains the entries defined in the table below in addition to the standard entries for streams. The data in the stream defines the mapping from character codes to a font number and a character selector. The data must follow the syntax defined in the *Adobe CMap and CIDFont File Specification* [4].

Table 7.21 CMap attributes

Key	Type	Semantics
Type	name	(Required) Must be CMap . (Note that although this object is the value of the key named Encoding in a Type 0 font, its type is CMap .)
CMapName	name	(Required) A name from Table 7.20 .
CIDSystemInfo	dictionary or array	(Required) This entry defines the Registry, Ordering, and Supplement attributes that define the character collection of the CIDFont; see Table 7.17 . If the CMap refers to a single CIDFont then the entry may be a dictionary. If this CMap refers to more than one descendant font, then an array must be supplied. The elements of the array correspond to the font numbers of the descendant fonts. If the descendant font is a CIDFont, a CIDSystemInfo dictionary must be supplied. If the descendant font is not a CIDFont, then null should be used for that element of the array.
WMode	integer	(Optional) An integer specifying the writing direction of the Type 0 font that uses this CMap. A value of 0 indicates a horizontal writing direction, and a value of 1 indicates a vertical writing direction. The default is 0.
UseCMap	name or stream	(Optional) The name of a predefined CMap, or a stream containing a CMap, that is to be used as the base for this CMap. This allows the CMap to be defined with only the character mappings that differ from the base CMap.

Below is a sample CMap for a Japanese Shift-JIS encoding. Character codes in this encoding can either be one or two bytes in length. In this CMap all character codes are mapped to CIDs in font number 0. This CMap could be used with a CIDFont that uses the same CID ordering as specified in the CIDSystemInfo entry. Note that several of the entries in the stream dictionary are also replicated in the stream data itself.

Example 7.11 *CMap Encoding*

```

22 0 obj <<
  /Type /CMap
  /CIDSystemInfo <<
    /Registry (Adobe)
    /Ordering (Japan1)
    /Supplement 2 >>
  /CMapName /90ms-RKSJ-H
  /WMode 0
  /Length 23 0 R
>>
stream
%!PS-Adobe-3.0 Resource-CMap
%%DocumentNeededResources: ProcSet (CIDInit)
%%IncludeResource: ProcSet (CIDInit)
%%BeginResource: CMap (90ms-RKSJ-H)
%%Title: (90ms-RKSJ-H Adobe Japan1 2)
%%Version: 1.402
%%Copyright: Copyright 1990-1994 Adobe Systems Inc.
%%Copyright: All Rights Reserved.
%%Copyright:
%%Copyright: Patents Pending
%%EndComments

/CIDInit /ProcSet findresource begin

12 dict begin

begincmap

/CIDSystemInfo 3 dict dup begin
  /Registry (Adobe) def
  /Ordering (Japan1) def
  /Supplement 2 def
end def

/CMapName /90ms-RKSJ-H def
/CMapVersion 1 def
/CMapType 1 def

/UIDOffset 950 def
/XUID [1 10 25343] def

/WMode 0 def

4 begincodespacerange
<00> <80>
<8140> <9FFC>

```

```

<A0> <DF>
<E040> <FCFC>
endcodespacerange

1 beginnotdefrange
<00> <1f> 231
endnotdefrange

100 begincidrange
<20> <7d> 231
<7e> <7e> 631
<8140> <817e> 633
<8180> <81ac> 696
<81b8> <81bf> 741
<81c8> <81ce> 749
... additional ranges ...

<fb40> <fb7e> 8518
<fb80> <fbfc> 8581
<fc40> <fc4b> 8706
endcidrange
endcmap
CMapName currentdict /CMap defineresource pop
end
end

%%EndResource
%%EOF

endstream
endobj

```

Note A CMap contained in a PDF file must not use the **usematrix**, **beginrearrangedfont**, or **endrearrangedfont** operators.

7.10.1 ToUnicode CMaps

In font dictionaries for Type 0 fonts, as well as for single-byte fonts (Type 1, Type 3, and TrueType), the value of the **ToUnicode** key is a CMap (see [Section 7.10 on page 215](#)) that defines the mapping from character codes to Unicode values. This entry is recommended for fonts that do not use one of the predefined CMaps or that use the **Identity-H** or **Identity-V** CMap with a CIDFont with an unknown character collection. If present, this entry allows text strings that use this font to be converted to Unicode values for export to other applications or plug-ins. The CMap defined in the **ToUnicode** entry must follow the syntax for CMaps defined in the *Adobe CMap and CIDFont File Specification* [4]. It must contain **begincodespacerange** and **endcodespace** range operators and use the **beginbfrange**, **endbfrange**, **beginbfchar**, and **endbfchar** operators to

define the mapping from character codes to 16-bit Unicode values. These operators have been extended to handle cases where a single character code maps to one or more Unicode values. See [Example 7.5 on page 201](#).

Implementation note

The known character collections in Acrobat 4.0 are **Adobe-Japan1-2**, **Adobe-Korea1-1**, **Adobe-CNS1-0**, and **Adobe-GB1-2**.

Following is an example of an embedded TrueType font that uses glyphIndices for character codes in the strings used by text operators in the page content stream.

Example 7.12 *Type 0 font that uses a **ToUnicode** entry*

```
14 0 obj <<
  /Type /Font
  /Subtype /Type0
  /BaseFont /Ryumin-Light
  /Encoding /Identity-H
  /DescendantFonts [15 0 R]
  /ToUnicode 16 0 R
>> endobj
```

The value of the **ToUnicode** entry is a stream object that contains the definition of the CMap. An example is shown below.

Example 7.13 *CMap used for a **ToUnicode** entry*

```
16 0 obj <<
  /Length 433
>>
stream
/CIDInit /ProcSet findresource begin
12 dict begin
begincmap
/CIDSystemInfo
<<
  /Registry (Adobe)
  /Ordering (UCS)
  /Supplement 0
>> def
/CMapName /Adobe-Identity-UCS def
/CMapType 2 def
1 begincodespacerange
<0000> <FFFF>
endcodespacerange

2 beginbfrange
<0000> <005e> <0020>
<005f> <0061> [<00660066> <00660069>
<00660066006c>]
endbfrange
endcmap
```

```

CMapName currentdict /CMap defineresource pop
end
end
endstream
endobj

```

The **begincodespacerange** and **endcodespacerange** operators in the above CMap define the source character code range to be the two-byte character codes from <00 00> to <FF FF>. The specific mapping for several of the character codes are shown. For example, <00 00> to <00 5E> are mapped to the Unicode values of U+0020 to U+007E. This is followed by the definition of a mapping used where the character code represents more than one Unicode value:

```

<005f> <0061> [<00660066> <00660069>
<00660066006c>]

```

In this case the original character codes are the glyphIndices for the glyphs for the ligatures *ff*, *fi*, and *ffl*. The entry defines the mapping from the character codes <00 5F>, <00 60>, and <00 61> to the string of Unicode values with a code for each character in the ligature: U+0066 U+0066 is the unicode for *ff*, U+0066 U+0069 is the Unicode for *fi*, and U+0066 U+0066 U+006c is the sequence for *ffl*.

The CMap example above illustrates several extensions to the way destination values can be defined. To support mappings from a source code to a string of destination codes, the following extension has been made to the ranges defined after a **beginbfchar** operator:

```

n beginbfchar
srcCode dstString
endbfchar

```

where *dstString* can be a string of up to 512 bytes. Likewise mappings after the **beginbfrange** operator maybe defined as:

```

n beginbfrange
srcCode1 srcCode2 dstString
endbfrange

```

In this case the last byte of the string will be incremented for each consecutive code in source code range. When defining ranges of this type, care must be taken to ensure that the value of the last byte in the string is less than or equal to $255 - (srcCode2 - srcCode1)$. This ensures that the last byte of the string will not be incremented past 255. If this occurs, the result of mapping is undefined and an error will occur.

To support more compact representations of mappings from a range of source character codes to a noncontiguous range of destination codes, the CMaps used for the **ToUnicode** entry may use the following syntax for the mappings following a **beginbfrange** definition:

```

n beginbfrange

```

```

srcCode1 srcCodeN
[ dstStr1 dstStr2 ... dstStrN ]
endbfrange

```

Consecutive codes starting with *srcCode1* and ending with *srcCodeN* are mapped to the destination strings in the array starting with *dstStr1* and ending with *dstStrN*.

7.11 Font descriptors

A font descriptor specifies a font's metrics and attributes. These metrics provide information for a viewing application to create a substitute multiple master font, where appropriate, or to select a similar font when the original font is unavailable. The font descriptor may also be used to embed the original font in the PDF file.

A font descriptor is a dictionary, as shown in [Table 7.22](#), whose keys specify various font attributes. All integer values are units in character space. The conversion from character space to text space depends on the type of font. See the discussion in [Section 7.7, "Fonts."](#)

Note For detailed information on the coordinate system in which characters are defined, see [Section 5.4 in the PostScript Language Reference Manual, Third Edition \[1\]](#) or [Section 3.1 in the Adobe Type 1 Font Format \[5\]](#).

Table 7.22 Attributes shared by all font descriptors

Key	Type	Semantics
Type	name	Object type. Always FontDescriptor .
Ascent	integer	<i>(Required)</i> The maximum height above the baseline reached by characters in this font, excluding the height of accented characters.
CapHeight	integer	<i>(Required)</i> The y-coordinate of the top of flat capital letters, measured from the baseline.
Descent	integer	<i>(Required)</i> The maximum depth below the baseline reached by characters in this font. Descent is a negative number.
Flags	integer	<i>(Required)</i> Collection of flags defining various characteristics of the font. See Table 7.24 .
FontBBox	Rectangle	<i>(Required)</i> The font's bounding box, which is the smallest rectangle enclosing the shape that would result if all characters in the font were placed with their origins coincident, and then painted.
FontName	name	<i>(Required)</i> The PostScript name of the font. (See page 204 for restrictions on the name.)

ItalicAngle	integer	<p><i>(Required)</i> Angle in degrees counterclockwise from the vertical of the dominant vertical strokes of the font. ItalicAngle is negative for fonts that slope to the right, as almost all italic fonts do.</p>	
StemV	integer	<p><i>(Required)</i> The width, measured in the x direction, of the dominant vertical stems of glyphs in the font.</p>	
AvgWidth	integer	<p><i>(Optional)</i> The average width of characters in this font. The default value is 0.</p>	
FontFile	stream	<p><i>(Optional)</i> A stream that defines a Type 1 font.</p>	
FontFile2	stream	<p><i>(Optional)</i> A stream that defines a TrueType font.</p>	
FontFile3	stream	<p><i>(Optional)</i> A stream that contains an embedded font. The format of the font is specified by the Subtype in the stream dictionary.</p>	
Leading	integer	<p><i>(Optional)</i> The desired spacing between lines of text. The default value is 0.</p>	
MaxWidth	integer	<p><i>(Optional)</i> The maximum width of characters in this font. The default value is 0.</p>	
MissingWidth	integer	<p><i>(Optional)</i> The width to use for unencoded character codes. The default value is 0.</p>	
StemH	integer	<p><i>(Optional)</i> The vertical width of the dominant horizontal stems of glyphs in the font. The default value is 0.</p>	
XHeight	integer	<p><i>(Optional)</i> The y-coordinate of the top of flat non-ascending lowercase letters, measured from the baseline. The default value is 0.</p>	
CharSet	string	<p><i>(Optional)</i> A string which lists the glyph names corresponding to the entries in the CharStrings dictionary if the font described is a subset font. Each name must be preceded by a slash. The names may appear in any order. The name .notdef should be omitted; it is assumed to exist in the font subset.</p>	

7.11.1 Font files

Implementation note

Font substitution is performed in Acrobat 1.0 through Acrobat 3.0 using multiple master Type 1 fonts. This substitution can be performed only for fonts that use the Adobe Roman Standard Character Set as defined in Appendix E.5 of the PostScript Language Reference Manual, Third Edition [1]. In Acrobat 3.0.1 and later, Type 0 fonts which use a CMap whose **CIDSystemInfo** defines the Adobe-Japan1, Adobe-Korea1, Adobe-GB1, or Adobe-CNS1 character collection can also be substituted. To make a document portable, it is necessary to embed fonts that cannot be substituted. The only exceptions are the fonts Symbol and ITC Zapf Dingbats, which are assumed to be present.

Fonts are embedded using the FontFile mechanism. If the embedded font is referred to by the **FontFile** key in Table 7.22, then the font must be in the standard Type 1 format. If the embedded font is referred to by the **FontFile2** key, then the

font must be in the TrueType format. If the embedded font is referred to by the **FontFile3** key, then the format of the font is specified by the **Subtype** key in the FontFile stream dictionary.

A standard Type 1 font definition, as described in the *Adobe Type 1 Font Format* [5], consists of three parts: a clear text portion, an encrypted portion, and a fixed-content portion. The fixed-content portion contains 512 ASCII zeros followed by a **cleartomark** operator, and perhaps followed by additional data. The stream dictionary for a font file contains the standard **Length** and **Filter** keys plus the additional keys shown in Table 7.23. While the encrypted portion of a standard Type 1 font may be in binary or ASCII hexadecimal format, PDF supports only the binary format. Example 7.14 shows the structure of an embedded standard Type 1 font.

Type 1 Compact fonts [7] (also known as CFF) may be embedded in a PDF 1.2 file using the **FontFile3** key. The **Subtype** value must be **Type1C**.

PDF 1.2

TrueType fonts are embedded using the FontFile2 mechanism. The font descriptor for an embedded TrueType font should contain a **FontFile2** key whose value is a stream that contains the TrueType font definition as described in *TrueType 1.0 Font Files*. The stream dictionary should include a **Length1** key as specified in Table 7.23; that key specifies the length in bytes of the font file after it has been decoded using the filters specified by the stream's **Filter** key. The **Length2** and **Length3** keys should not be used for TrueType fonts.

CIDFontType0 fonts are embedded using the **FontFile3** key with a **Subtype** value of **CIDFontType0C**. These fonts are embedded using the Compact Font Format described in the *Compact Font Format Specification* [7].

PDF 1.3

CIDFontType2 fonts are based on the TrueType font format and therefore are embedded using the **FontFile2** key. Glyphs in an embedded **CIDFontType2** font are accessed by glyphIndex; therefore, the embedded TrueType font does not require the **cmap** table to map from character code to glyphIndex. The following TrueType tables are required: **head**, **hhea**, **loca**, **maxp**, **cvt_**, **prep**, **glyf**, **hmtx**, and **fgpm**. For details on this format, see *The Type 2 Charstring Format* [8].

PDF 1.3

Due to the large number of characters in Chinese, Korean and Japanese fonts, the embedded font typically contains only the subset of glyphs that are used in the PDF document.

Table 7.23 Additional attributes for FontFile stream

Key	Type	Semantics
Length1	integer	(Required) Length in bytes of the ASCII portion of the Type 1 font file after it has been decoded using the filters specified by the stream's Filter key.
Length2	integer	(Required for Type 1 fonts) Length in bytes of the encrypted portion of the Type 1 font file after it has been decoded using the filters specified by the stream's Filter key.

Length3	integer	<i>(Required for Type 1 fonts)</i> Length in bytes of the portion of the Type 1 font file that contains the 512 zeros, plus the cleartomark operator, plus any following data. This is the length of the data after it has been decoded using the filters specified by the stream's Filter key. If Length3 is zero, it indicates that the 512 zeros and cleartomark have not been included in the FontFile and must be added.
Subtype	name	<i>(Required if in FontFile3)</i> The name specifies the format of the embedded font. The name must be Type1C for Type 1 Compact fonts. For CIDFontType0 Compact fonts, use CIDFontType0C . When additional font formats are added to PDF, more values will be defined for Subtype .

PDF 1.2

Note Font descriptors for Type 0 fonts contain additional attributes. See [Section 7.11.3](#) on page 227.

Example 7.14 *Embedded Type 1 font definition*

```

12 0 obj
<<
  /Filter /ASCII85Decode
  /Length 41116
  /Length1 2526R
  /Length2 32393
  /Length3 570
>>
stream
,p>`rDKJj'E+LaU0eP.@+AH9dBOu$hFD55nC
... omitted data ...
JJQ&Nt' ) <=^p&mGf (%:%h1%9c//K(/*o=.C>UXkbVGTrr~>
endstream
endobj

```

7.11.2 Font descriptor flags

The value of the **Flags** key in a font descriptor is a 32-bit integer that contains a collection of boolean attributes. These attributes are *true* if the corresponding bit is set to 1 in the integer. [Table 7.24](#) specifies the meanings of the bits, with bit 1 being the least significant. Reserved bits must be set to zero.

Table 7.24 *Font flags*

<i>Bit position</i>	<i>Semantics</i>
1	Fixed-width font
2	Serif font
3	Symbolic font
4	Script font
5	Reserved

6	Uses the Adobe Standard Roman Character Set
7	Italic
8–16	Reserved
17	All-cap font
18	Small-cap font
19	Force bold at small text sizes
20–32	Reserved

Note Bits 3 and 6 may not both be zero.

All characters in a *fixed-width font* have the same width, while characters in a *proportional font* have different widths. Characters in a *serif font* have short strokes drawn at an angle on the top and bottom of character stems, while *sans serif fonts* do not have such strokes. A *symbolic font* contains symbols rather than letters and numbers. Characters in a *script font* resemble cursive handwriting. An *all-cap font*, which is typically used for display purposes such as titles or headlines, contains no lowercase letters. It differs from a *small-cap font* in that characters in the latter, while also capital letters, have been sized and their proportions adjusted so that they have the same size and stroke weight as lowercase characters in the same typeface family. Figure 7.4 shows examples of these types of fonts.

Figure 7.4 Characteristics represented in the *Flags* field of a font descriptor

The quick brown fox jumped...	Fixed-width font
The quick brown fox jumped...	Sans serif font
The quick brown fox jumped...	Serif font
*** □♦*** ⚙️□□■ ❄️□ *◆○□**®	Symbolic font
<i>The quick brown fox jumped...</i>	Script font
<i>The quick brown fox jumped...</i>	Italic font
THE QUICK BROWN FOX JUMPED	All cap font
THE QUICK BROWN FOX JUMPED...	Small cap font

Bit 6 in the flags field indicates that the font’s character set is the Adobe Standard Roman Character Set, or a subset of that, and that it uses the standard names for those characters. The characters in the Adobe Standard Roman Character Set are shown in the first column of Table C.1 on page 448 (A, Æ, Á, etc.); the character names are shown in column 2 (A, AE, Aacute, etc.).

Finally, bit 19 is used to determine whether or not bold characters are drawn with extra pixels even at very small text sizes. Typically, when characters are drawn at small sizes on very low resolution devices such as display screens, features of bold

characters may appear only one pixel wide. Because this is the minimum feature width on a pixel-based device, ordinary non-bold characters also appear with one-pixel wide features, and cannot be distinguished from bold characters. If bit 19 is set, features of bold characters may be thickened at small text sizes.

Example 7.15 *Font descriptor*

```

7 0 obj
<<
  /Type /FontDescriptor
  /FontName /AGaramond-Semibold
  /Flags 262192                               Bits 5, 6, and 19
  /FontBBox [-177 -269 1123 866]
  /MissingWidth 255
  /StemV 105
  /StemH 45
  /CapHeight 660
  /XHeight 394
  /Ascent 720
  /Descent -270
  /Leading 83
  /MaxWidth 1212
  /AvgWidth 478
  /ItalicAngle 0
>>
endobj

```

7.11.3 Font descriptors for CID fonts

PDF 1.2

The following table lists the additional entries that may be found in the FontDescriptor dictionaries of Chinese, Japanese, or Korean CIDFonts.

Table 7.25 *Additional FontDescriptor attributes*

Key	Type	Semantics
Style	dictionary	<i>(Optional)</i> This entry is a dictionary that contains key-value pairs that describe the style of the glyphs in the font. See Table 7.26 for valid entries in this dictionary.
Lang	name	<i>(Optional)</i> This entry indicates the language of the font. It is used for encodings where the language is not implied by the encoding itself. The possible values are the two character names defined by ISO639; see Appendix I .
FD	dictionary	<i>(Optional)</i> The keys in this dictionary identify a subset of characters in a CIDFont. The values are dictionaries with entries that override the values in the FontDescriptor dictionary for the subset of characters.
CIDSet	stream	<i>(Optional)</i> This entry identifies which CIDs are present in the CIDFont file. It contains only a subset of the glyphs in the character collection defined by the CIDSystemInfo. If this entry is missing, then it is assumed that the CIDFont file contains all the glyphs for the character collection. The stream's length should be

rounded up to the nearest multiple of 8. The bits should be stored in bytes with the high-order bit first. Each bit corresponds to a CID. The first bit of the first byte corresponds to CID 0, the next bit corresponds to CID 1, and so on. If the subset contains a CID, the bit for that CID should be set. For compactness the stream may use one of the compression filters to encode the data.

Style

The Style dictionary contains a set of key-value pairs that define style attribute and value for the font. Currently only the **Panose** key is defined. The value of the **Panose** key is the 12-byte string containing the Family Class ID, Family SubClass ID, and 10 bytes for the Panose classification number for the font. For additional details on the Panose number see the *Japanese TrueType Font Property Selection Guidelines* by the TrueType Conference Technical Committee [36]. The following is an example of a **Style** entry in the FontDescriptor:

```
/Style
<< /Panose <01 05 02 02 03 00 00 00 00 00 00 00>
>>
```

FD

A CIDFont may be made up of different types of characters, each type requiring different sets of metrics. Numeric characters, for example, may require different metrics from ideographic characters. The font descriptor defines a set of default metrics that apply to all characters in the CIDFont; the **FD** entry in the font descriptor contains exceptions to these defaults. Each key in an **FD** dictionary is the name of a type of character; each type is defined as a particular subset of characters in the font.

It is strongly recommended that the **FD** dictionary contain at least the metrics for the proportional Roman characters. With the information for the proportional Roman characters, a more accurate substitution font can be created. For CIDFonts that use the Adobe-Japan1 character collection, the name for the proportional Roman characters is **Proportional**. For CIDFonts that use the Adobe-CNS1, Adobe-GB1, and Adobe-Korea1 character collections, the name for the proportional Roman characters is **Alphabetic**.

The following table gives the valid keys for the Adobe-Japan1, Adobe-Korea1, Adobe-Japan2, Adobe-GB1, and Adobe-CNS1 character collections.

Table 7.26 *Character Subsets in CJK fonts*

<i>Character collection</i>	<i>Key</i>	<i>Character subset</i>
Adobe-Japan1	AlphaNum	Numeric characters
	Alphabetic	Full-width Roman characters
	Dingbats	Special symbols
	HKana	Half-width Katakana characters.
	HRoman	Half-width Roman characters

Table 7.26 *Character Subsets in CJK fonts*

<i>Character collection</i>	<i>Key</i>	<i>Character subset</i>
Adobe-Korea1	Kana	Full-width Kana characters
	Kanji	Full-width ideographic characters
	Proportional	Proportional Roman characters
Adobe-Japan2	Alphabetic	Proportional Roman characters
	Dingbats	Special symbols
	Hangul	Full-width ideographic and Hangul characters
Adobe-GB1	Alphabetic	Roman characters
	Dingbats	Special symbols
	HojoKanji	Full-width ideographic characters
Adobe-CNS1	Alphabetic	Proportional Roman characters
	GBHanzi	Full-width ideographic characters
	Dingbats	Special symbols
	Kana	Japanese Kana characters
Adobe-CNS1	Alphabetic	Proportional Roman characters
	CNSHanzi	Full-width ideographic characters
	Dingbats	Special symbols
	Kana	Japanese Kana characters

Following is an example with two entries in the **FD** dictionary.

Example 7.16 *FD entry*

```

/ FD
<<
/Proportional 25 0 R
/HKana 26 0 R
>>

25 0 obj
<<
/Type /FontDescriptor
/FontName /HeiseiMin-W3-Proportional
/Flags 2
/AvgWidth 478
/MaxWidth 1212
/MissingWidth 250
/StemV 105
/StemH 45
/CapHeight 660
/XHeight 394
/Ascent 720
/Descent -270

```

```

/Leading 83
>>
endobj

26 0 obj
<<
/Type /FontDescriptor
/FontName /HeiseiMin-W3-HKana
/Flags 3
/Style Gothic
/AvgWidth 500
/MaxWidth 500
/MissingWidth 500
/StemV 50
/StemH 75
/Ascent 720
/Descent 0
/Leading 83
>>
endobj


```

7.12 Color spaces

A color space specifies how color values should be interpreted. While some PDF operators implicitly specify the color space they use, others require that a color space be specified explicitly. As shown in [Figure 7.5](#), PDF 1.3 supports eleven color space families:

PDF 1.3

Figure 7.5 *Color space families*



- three device-dependent color space families, **DeviceGray**, **DeviceRGB**, and **DeviceCMYK**;
- four device-independent color space families, **CalGray**, **CalRGB**, **Lab**, and **ICCBased**;
- and four special color space families, **Indexed**, **Pattern**, **Separation**, and **DeviceN**.

The color spaces follow the semantics described in Section 4.8 of the *PostScript Language Reference Manual, Third Edition* [1].

A *color space* is an instance of a *color space family*. For example, there are many possible color spaces in the **CalGray** family. A color space is specified by an array whose first element is the family name and whose remaining elements are parameters that define the particular instance. There is only one member of each of the device-dependent families, **DeviceGray**, **DeviceRGB**, and **DeviceCMYK**, so they are treated as color spaces and may be specified with just the name. For example, the **DeviceGray** color space may be written as either

[/DeviceGray]

or just

/DeviceGray

Pattern color spaces that represent colored tiling patterns or shading patterns may also be written with just the name, **/Pattern**. Uncolored tiling patterns, however, require an array; see [Section 7.12.11](#), “[Pattern color spaces](#).”

In a device-dependent color space, the color values are interpreted as specifying the percentage of device colorant to be used. This means that the exact color produced depends on the characteristics of the output device. For example, in the **DeviceRGB** color space, a value of 1 for the red component means “turn red all the way on.” If the output device is a monitor, the color displayed depends strongly on the settings of the monitor’s brightness, contrast, and color balance adjustments. In addition, the precise color displayed depends on the chemical composition of the compound used as the red phosphor in the particular monitor being used, the length of time the monitor has been turned on, and the age of the monitor.

In a device-independent color space, color values are defined by a mapping from the device-independent color space into a standard color space, the CIE (*Commission Internationale de l’Éclairage*) 1931 XYZ color space. Since the values in the XYZ space can be measured colorimetrically, this establishes a device-independent specification of the desired color. When a device-independent color value is rendered on a device, the rendered color is based on the device-independent color specification as well as the color characteristics of the device. This may or may not result in a true colorimetric rendering. Variations from a colorimetric rendering may occur as a consequence of gamut limitations and rendering intents. See the discussion in [Section 7.13.4](#), “[Color rendering intent](#).”

See the *PostScript Language Reference Manual, Third Edition* [1] for further explanation of device-independent color.

PDF 1.3

PDF 1.1

7.12.1 DeviceGray color spaces

Colors in the **DeviceGray** color space are specified by a single value: the intensity of achromatic light. In this color space, 0 is black, 1 is white, and intermediate values represent shades of gray.

7.12.2 DeviceRGB color spaces

Colors in the **DeviceRGB** color space are represented by three values in the range 0 to 1: the intensity of the red, green, and blue components in the output.

DeviceRGB is commonly used for video displays because they are generally based on red, green, and blue phosphors.

7.12.3 DeviceCMYK color spaces

Colors in the **DeviceCMYK** color space are represented by four values in the range 0 to 1. These values are the amounts of the cyan, magenta, yellow, and black components in the output. This color space is commonly used for color printers, where they are the colors of the inks traditionally used for four-color printing.

Only cyan, magenta, and yellow are strictly necessary, but black is generally also used in printing because black ink produces a better black than a mixture of cyan, magenta, and yellow inks, and because black ink is less expensive than the other inks.

Note In PDF 1.1, a color space named **CalCMYK** was partially defined with the expectation that its definition would be completed in a future version of PDF. That is no longer being considered. PDF 1.3 supports calibrated 4-component color spaces by means of ICC profiles; see [Section 7.12.7, “ICCBased color spaces.”](#) PDF consumers should ignore **CalCMYK** color space attributes and render colors specified in this color space as if they had been specified using **DeviceCMYK**.

PDF 1.3

7.12.4 CalGray color spaces

PDF 1.1

Colors in a **CalGray** color space are represented by a single value. Input values are in the range 0 to 1, where 0 is black, 1 is white and intermediate values are gray.

A **CalGray** color space is specified by an array of the form

[/CalGray dict]

where the contents of *dict* are described in [Table 7.27](#).

Table 7.27 *CalGray attributes*

Key	Type	Semantics
WhitePoint	array	<i>(Required)</i> Three numbers $[X_w Y_w Z_w]$ that specify the CIE 1931 (XYZ)-space tristimulus value of the diffuse white point. The numbers X_w and Z_w must be positive, and Y_w must be equal to 1. See discussion in 4.8.3 in the <i>PostScript Language Reference Manual, Third Edition</i> [1] for further details.
BlackPoint	array	<i>(Optional)</i> Three numbers $[X_b Y_b Z_b]$ that specify the CIE 1931 (XYZ)-space tristimulus value of the diffuse black point. The numbers must be non-negative. The default value is $[0 0 0]$. See discussion in 4.8.3 in the <i>PostScript Language Reference Manual, Third Edition</i> [1] for further details.
Gamma	number	<i>(Optional)</i> Defines the exponential relationship between the gray component and Y . The governing equation is $Y = \text{gray}^{\text{Gamma}}$. <i>Gamma</i> must be positive and is generally greater than or equal to 1. The default value is 1.

7.12.5 CalRGB color spaces

PDF 1.1

Colors in a **CalRGB** color space are represented by three values: the red, green, and blue components of the color. Each value is in the range 0 to 1.

A **CalRGB** color space is specified by an array of the form:

$[/CalRGB dict]$

where the contents of *dict* are described in [Table 7.28](#).

Table 7.28 *CalRGB attributes*

Key	Type	Semantics
WhitePoint	array	<i>(Required)</i> Same as for CalGray.
BlackPoint	array	<i>(Optional)</i> Same as for CalGray.
Gamma	array	<i>(Optional)</i> Three numbers $[G_r G_g G_b]$ that specify the gamma for the red, green, and blue components respectively. The governing equations are $R' = R^{G_r}$, $G' = G^{G_g}$, and $B' = B^{G_b}$, where R , G , and B are the input calibrated RGB values, and R' , G' , and B' are the gamma-modified values. The default value is $[1 1 1]$.
Matrix	array	<i>(Optional)</i> Nine numbers $[X_r Y_r Z_r X_g Y_g Z_g X_b Y_b Z_b]$ that specify the linear interpretation of the gamma-modified red, green, and blue components, R' , G' , and B' . The default value is the identity matrix, $[1 0 0 0 1 0 0 0 1]$. The transformation from $R'G'B'$ to XYZ is given by: $X = R' \times X_r + G' \times X_g + B' \times X_b$ $Y = R' \times Y_r + G' \times Y_g + B' \times Y_b$ $Z = R' \times Z_r + G' \times Z_g + B' \times Z_b$

An example of a **CalRGB** color space is shown here for D65 white point, 1.8 gammas, and Trinitron phosphor chromaticities.

```

12 0 obj
[/CalRGB
<<
/WhitePoint [0.9505 1 1.0890]
/Gamma [1.8 1.8 1.8]
/Matrix [0.4497 0.2446 0.0252
         0.3163 0.6720 0.1412
         0.1845 0.0833 0.9227]
>> ]
endobj

```

7.12.6 Lab color spaces

PDF 1.1

Colors in a **Lab** color space are represented by three values: the L^* , a^* and b^* components of the color. The ranges of each of the three values are specified under the **Range** key in [Table 7.29](#). For details about Lab color spaces, see Section 4.8.3 of *The PostScript Language Reference Manual, Third Edition* [1].

A **Lab** color space is specified by an array of the form:

```
[/Lab dict ]
```

where the contents of *dict* are described in [Table 7.29](#).

Table 7.29 *Lab attributes*

Key	Type	Semantics
WhitePoint	array	(Required) Same as for CalGray .
BlackPoint	array	(Optional) Same as for CalGray .
Range	array	(Optional) Four numbers [a_{\min} a_{\max} b_{\min} b_{\max}] specifying the range of the a^* and b^* components, which must be between -128 and 127. That is, a^* and b^* are limited by $-128 \leq a_{\min} \leq a^* \leq a_{\max} \leq 127$, and $-128 \leq b_{\min} \leq b^* \leq b_{\max} \leq 127$. The default value is [-100 100 -100 100]. The range of L^* is always 0 to 100.

PDF 1.3

7.12.7 ICCBased color spaces

PDF 1.3

An **ICCBased** color space contains a *color profile* as defined by the International Color Consortium (ICC). See the ICC specification [29] for details of the profile format.

An **ICCBased** color space is specified as an array:

```
[/ICCBased stream ]
```

The stream contains the ICC profile. In addition to keys that are defined for any stream, the stream has the following keys:

Table 7.30 *ICCBased attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
N	integer	<i>(Required)</i> The number of color components in the color space described by ICC profile data. This number must match the number of components actually in the ICC profile. In PDF 1.3, N must be 1, 3, or 4, and there are additional restrictions; see below.
Alternate	color space	<i>(Optional)</i> Provides an alternate color space to be used in the case where the color space specified in the stream data isn't supported. It can be any valid color space, except a Pattern color space, that has the same number of components as the color space specified by the stream data (N). For example, one might want to provide an alternate space if using an ICCBased color space not allowed in PDF 1.3, but permitted in a later version of PDF. If this key is omitted, and the ICC profile data is not understood by the PDF processor, and the value of the N key is 1, 3 or 4, then a PDF consumer treats the color space as being DeviceGray , DeviceRGB , or DeviceCMYK , respectively. Note that there is no conversion of color values, such as a tint transformation, when using the alternate color space. Color values that are within the range of the ICCBased color space may not be within the range of the alternate color space; in this case, the nearest values within the alternate's range will be used.
Range	array	<i>(Optional)</i> An array of $2 \times \mathbf{N}$ numbers, [min_0 max_0 min_1 max_1 ...], specifying the minimum and maximum values that the corresponding color components may have. This must match the information in the ICC profile. Default: [0 1 0 1 ...] .

The ICC specification is an evolving standard. The **ICCBased** color spaces supported in PDF 1.3 are based on version 3.3 of the ICC specification. Early versions of the ICC specification are also supported. (The version number is available in the ICC profile's header.)

PDF 1.3 supports only the profile types shown in the following table; other types may be supported in the future. (Profiles must satisfy *both* the criteria shown in the table. The terminology is taken from the ICC specifications.)

Table 7.31 *ICC profile types*

<i>Header field</i>	<i>Required value</i>
deviceClass	icSigInputClass ('scnr') icSigDisplayClass ('mnr')
	icSigOutputClass ('prtr')
or	icSigColorSpaceClass ('spac')

```
colorSpace      icSigGrayData ('GRAY')
                icSigRgbData ('RGB ')
                icSigCmykData ('CMYK')
or              icSigLabData ('Lab ')
```

An example of an **ICCBased** color space is shown here. The profile's data has been encoded in hexadecimal representation for readability.

Example 7.17 *ICCBased color space*

```
12 0 obj
[/ICCBased 15 0 R]
endobj

15 0 obj
<<
/N 3
/Alternate /DeviceRGB
/Length 524
/Filter /ASCIIHexDecode
>>
stream
00 00 02 0C 61 70 70 6C 02 00 00 00 6D 6E 74 72
52 47 42 20 58 59 5A 20 07 CB 00 02 00 16 00 0E
00 22 00 2C 61 63 73 70 41 50 50 4C 00 00 00 00
61 70 70 6C 00 00 04 01 00 00 00 00 00 00 02
00 00 00 00 00 00 F6 D4 00 01 00 00 00 00 D3 2B
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 09 64 65 73 63 00 00 00 F0 00 00 00 71
72 58 59 5A 00 00 01 64 00 00 00 14 67 58 59 5A
00 00 01 78 00 00 00 14 62 58 59 5A 00 00 01 8C
00 00 00 14 72 54 52 43 00 00 01 A0 00 00 00 0E
67 54 52 43 00 00 01 B0 00 00 00 0E 62 54 52 43
00 00 01 C0 00 00 00 0E 77 74 70 74 00 00 01 D0
00 00 00 14 63 70 72 74 00 00 01 E4 00 00 00 27
64 65 73 63 00 00 00 00 00 00 00 00 17 41 70 70 6C
65 20 31 33 22 20 52 47 42 20 53 74 61 6E 64 61
72 64 00 00 00 00 00 00 00 00 00 00 17 41 70
70 6C 65 20 31 33 22 20 52 47 42 20 53 74 61 6E
64 61 72 64 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 58 59 5A 58 59 5A 20 00 00 00 00 00 00 63 0A
00 00 35 0F 00 00 03 30 58 59 5A 20 00 00 00 00
00 00 53 3D 00 00 AE 37 00 00 15 76 58 59 5A 20
00 00 00 00 00 00 40 89 00 00 1C AF 00 00 BA 82
```

```

63 75 72 76 00 00 00 00 00 00 01 01 CC 63 75
63 75 72 76 00 00 00 00 00 00 01 01 CC 63 75
63 75 72 76 00 00 00 00 00 00 01 01 CC 58 59
58 59 5A 20 00 00 00 00 00 00 F3 1B 00 01 00 00
00 01 67 E7 74 65 78 74 00 00 00 00 20 43 6F 70
79 72 69 67 68 74 20 41 70 70 6C 65 20 43 6F 6D
70 75 74 65 72 73 20 31 39 39 34 00
endstream
endobj

```

The terminology used in PDF color spaces and ICC color profiles is similar, but sometimes the same terms are used with different meanings. For example, the default value for each component in an **ICCBased** color space is zero. The range of each color component is a function of the color space specified by the profile, and is indicated in the ICC specification. The ranges for several ICC color spaces are shown in the following table; not all of these are supported in PDF.

Table 7.32 *Ranges for typical ICC color spaces*

<i>ICC Color space</i>	<i>Component ranges</i>
Gray	[0 1]
RGB	[0 1]
CMYK	[0 1]
L*a*b*	L*: [0 100]; a* and b*: [-128 127]
L*a*b* (16-bit)	L*: [0 100]; a* and b*: [-128 127 + $\left(\frac{255}{256}\right)$]
XYZ	[0 1 + $\left(\frac{32767}{32768}\right)$]

Note that the ranges on color spaces described by Colorspace Conversion Profiles, e.g., L*a*b* and XYZ, are a bit unusual and depend on the bit depth of the profiles themselves.

Since the **ICCBased** color space is being used for the specification of a source color space, only the “to CIE” (*AToB* in ICC parlance) profile information is used; the “from CIE” (*BToA*) information is ignored when present.

The size of the ICC profile in the PDF file can be reduced by using a lossless compression filter, e.g., Flate (page 47) or LZW (page 45).

PDF 1.1 introduced support for device-independent color spaces: **CalGray**, **CalRGB**, and **Lab**. In PDF 1.3, the **ICCBased** color space also provides for device-independent color specifications. The **CalGray**, **CalRGB**, and **Lab** color spaces of PDF 1.1 provide a more compact representation of the color spaces they specify. The **ICCBased** color space provides support for additional color spaces. In those cases where the desired color space can be expressed by either a PDF 1.1 device-independent color space, or by the **ICCBased** color space, the resulting colors are expected to be rendered similarly, independent of the method selected for representation.

PDF 1.1 also introduced support for rendering intents. The rendering intent for the graphics state is set using the **ri** operator (see [page 333](#)). Images also have an optional **Intent** field. The default rendering intent in both cases is **RelativeColorimetric**. ICC profiles also contain a rendering intent in their header. The rendering intent specified in the ICC profile's header is ignored in PDF, in order that the same color space information can be used for different objects in a PDF with different rendering intents. The rendering intents specified by the **ri** operator for the graphics state, by the **Intent** field for images, or by default in either case, specify the rendering intent.

7.12.8 Separation color spaces

PDF 1.2

Colors in a **Separation** color space are represented by a single value, called a *tint*, in the range of 0 to 1. The value 0 represents the minimum amount of colorant; 1 represents the maximum amount.

A **Separation** color space is specified as a 4-element array:

[/Separation name alternateSpace tintTransform]

name is the name of the separation or colorant. At the moment when the color space is set to a separation color space on a printer or viewer, if the device can produce the named separation, then it does so, and the remaining elements of the array are ignored.



If the device cannot produce the named separation, then subsequent painting operations are performed in the color space specified by *alternateSpace*. The alternate color space can be a device-dependent color space or a device-independent color space, but not a special color space. The color used in the alternate space is determined by applying *tintTransform* to the tint value. *tintTransform* is a Function (see [page 262](#)) whose input is the tint and whose output is a set of color values, one for each of the color components in the alternate color space.

7.12.9 DeviceN color spaces

PDF 1.3

A **DeviceN** color space can be thought of as a generalization of a **Separation** color space from one colorant to *n* colorants. It provides for the composite specification of an arbitrary number of device-dependent color components. This color space provides greater control with respect to overprint between the components. Previously, for example, to paint only the cyan, magenta, and yellow colorants on a CMYK device and leave the black colorant unaltered required using the **Separation** color space once each for cyan, magenta, and yellow (and turning overprint on); a composite description of this color is not permitted using the **DeviceCMYK** color space without altering the black component. In contrast, the **DeviceN** color space allows you to create a composite description for the cyan, magenta, and yellow color components, while leaving the black component unaltered.

It is also often desirable to use data for a single component to paint multiple color components. For example, a single gray image can be used to produce a duotone. The **DeviceN** color space, used in conjunction with an **Indexed** color space, provides this functionality.

A **DeviceN** color space is specified as follows:

[/DeviceN names alternateSpace tintTransform]

or

[/DeviceN names alternateSpace tintTransform attributes]

Color values in the **DeviceN** color space are tint components in the range 0 to 1. The value 0 represents the minimum amount of colorant; the value 1 represents the maximum. The **scn** and **SCN** operators set the current color in the graphics state to a set of tint values; the initial value is 1 for each tint. A sampled image can also be treated as a source of tint values.

The number of color components is derived from the length of the *names*, which is an array containing name objects, specifying the individual colorants. It is an error for the *names* array to contain any name more than once. When **DeviceN** is the current color space, the **scn** and **SCN** operators take a number of arguments equal to the length of the *names* array.

At the moment the color space is set to **DeviceN** on a rendering device (printer or viewer), if all named colorants are available on that device, then subsequent painting operations use the named device colorants, and the parameters *alternateSpace* and *tintTransform* are ignored. If not all the named colorants are available, then subsequent painting operations are performed using the *alternateSpace* color space. The *alternateSpace* color space enables the colors to be approximated by colors in some other color space.

The *alternateSpace* parameter may be any device-dependent or device-independent color space, but not a special color space.

The *tintTransform* procedure is a Function (see [page 262](#)) that provides an *n*-to-*m* transformation, where *m* is the number of color components in the *alternateSpace*. During subsequent painting operations, if the named colorants are not available on the device, this procedure is used to transform *n* tint values into *m* color component values in the alternate color space.

In a **DeviceN** color space, one or more of the colorant names (in the *names* array) can be **None**. This indicates that the corresponding color component is never marked on the page, consistent with the behavior of a **Separation** color space for the **None** colorant. When a **DeviceN** color space is marking the named device colorants directly, color components corresponding to **None** colorants are discarded. However, when the **DeviceN** color space reverts to its alternate color space, those components are passed to the *tintTransform* function, which may use them in any desired manner.



The optional *attributes* parameter is a dictionary containing additional information about the color space. The only attribute (key) currently defined for this dictionary is **Colorants**. The value of this key is a dictionary whose keys are names of colorants and whose values are Separation color spaces corresponding to those colorants. The Colorants dictionary may contain entries for any colorants, not just the ones used in the DeviceN color space; it may be shared among several DeviceN color spaces using different sets of colorants, so long as all of them use consistent definitions for any colorants.

Example 7.18 A DeviceN color space with a Colorants dictionary

```

30 0 obj
  [/DeviceN [/Orange /Green /None]
   /DeviceCMYK
   tintTransform1
  <</Colorants 45 0 R>>
]
endobj

45 0 obj
<</Orange [/Separation /Orange
           /DeviceCMYK tintTransform2 ]
   /Green  [/Separation /Green
           /DeviceCMYK tintTransform3 ]
   /PANTONE#20131 [/Separation /PANTONE#20131
                 /DeviceCMYK tintTransform4 ]
>>
endobj

```

In this example, the DeviceN color space, object 30, has an attributes dictionary that contains an indirect reference to a Colorants dictionary, object 45, which might also be referenced by other attributes dictionaries. *tintTransform1* maps 3 color components (orange, green, and “none”) to 4 color components (CMYK). *tintTransform2* maps a single color component (an orange tint) to 4 CMYK components. Likewise, *tintTransform3* maps a green tint to CMYK, and *tintTransform4* maps a tint of Pantone 131 to CMYK.

7.12.10 Indexed color spaces

PDF 1.2

Indexed color spaces allow colors to be specified by small integers that are used as indexes into a table of color values. The values in this table are colors specified in some other “base” color space. For example, an indexed color space can have white as color number 1, dark blue as color number 2, turquoise as color number 3, and black as color number 4.

An **Indexed** color space is specified as follows:

```
[/Indexed base hival lookup]
```


The base color space is specified by *base*. In PDF 1.0 and PDF 1.1 it must be either **DeviceRGB** or **DeviceCMYK**. In PDF 1.2 it may be any color space except a **Pattern** color space, another **Indexed** color space, or a **Separation** color space. In PDF 1.3, it may be any color space except a **Pattern** color space or another **Indexed** color space; in particular, **Separation**, **DeviceN**, and **ICCBased** color spaces are allowed. The maximum valid index value, specified by *hival*, is determined by the number of colors desired in the **Indexed** color space; it may not be greater than 255.

Colors are specified by integers in the range 0 to *hival*. The color table values are contained in *lookup*, which must be a stream in PDF 1.0 and 1.1 but may also be a string in later versions of PDF. The stream or string contains $m \times (hival + 1)$ bytes where m is the number of color components in the base color space. Each byte is an unsigned integer in the range 0 to 255 that is scaled to the range of the corresponding color component in the base; that is, 0 corresponds to the minimum value in the range for that component, and 255 corresponds to the maximum value in the range.²

The color components for each entry in the table are adjacent in the stream. For example, if the base color space is **DeviceRGB** and the indexed color space contains two colors, the order of bytes in the stream is: $R_0 G_0 B_0 R_1 G_1 B_1$, where letters are the color component and numbers are the table entry.

[Example 7.19](#) shows an indexed color space. Colors in the table are specified in the **DeviceRGB** color space, and the table contains 256 entries. The stream containing the table has been LZW and ASCII base-85 encoded.

Example 7.19 *Indexed color space*

```
12 0 obj
[/Indexed /DeviceRGB 255 13 0 R]
endobj

13 0 obj <<
/Filter [/ASCII85Decode /LZWDecode]
/Length 554
>>
stream
J3Vsg-=dE=!]* )rE$,8^$P%cp+RI0B1)A)g_;FLE.V9
... omitted data ...
bS/5%"OmlTJ=PC!c2]]^rh(A~>
endstream
endobj
```

2. This is different than lookup-strings in Indexed color spaces in PostScript, where the values are always divided by 255, yielding a color component value in the range 0 to 1. Lookup-strings in PostScript can therefore be used only for color spaces whose components are all in the range 0 to 1. For wider ranges, such as $L^*a^*b^*$, PostScript uses procedures in place of strings.

7.12.10.1 Indexed DeviceN (duotones)

DeviceN color spaces can be used for objects that have an arbitrary number of color components. A duotone image is a 1-component image that is printed with two inks to add tonal depth. Similarly, monotones, tritones, and quadtones use 1, 3, and 4 inks. A duotone image can be easily represented, therefore, with an Indexed DeviceN color space. An Indexed color space has 1 component; its lookup table maps the single index into the N color components of the DeviceN color space, whose colorants are the desired inks.

For example, an image that uses two inks, Black and Gold (a “spot” color), might use the following color space:

```
[/Indexed
  [/DeviceN [/Black /Gold]
   [/CalRGB <</WhitePoint [1 1 1]
    /Gamma [2.2 2.2 2.2] >> ]
   tintTransform ]
255
lookupTable ]
```

In this example, the lookup table would have 512 bytes (256 entries in the table \times 2 bytes per entry), mapping each index to a 2-component (black + gold) color. The alternate for this DeviceN color space is a calibrated RGB; this would be used when printing to a device that does not support black and gold colorants (and when previewing the image in Acrobat). *tintTransform* is a function that converts a 2-component color in the DeviceN color space to a 3-component color in the CalRGB color space.

Specifying a 2-in, 3-out tint transform function may require a Type 0 function with a large number of sample points to represent the function accurately. If each input variable were sampled in 256 places between 0 and 1, that would require 256^2 samples. But in the case of an indexed image, there are only 256 possible combinations of input values, so a more compact way to represent this information is to put the alternate color values directly into the lookup table along with the DeviceN color values. Doing this with the previous example would result in the following color space:

```
[/Indexed
  [/DeviceN [/Black /Gold /None /None /None]
   [/CalRGB <</WhitePoint [1 1 1]>> ]
   tintTransform ]
255
lookupTable ]
```

In this example, each entry in the lookup table has five components, two for the black and gold colorants, and three for the CalRGB colors (the “None” colorants). If the DeviceN color space is used when printing, the None components are ignored. If the alternate color space is used, the tint transform will be used to convert a 5-component color into a 3-component CalRGB color, but by

construction, the third, fourth, and fifth components *are* the CalRGB components, so the tint transform function can merely discard the first two components. The way to express that is with a Type 4 (“PostScript calculator”) function:

```
20 0 obj <<
  /FunctionType 4
  /Length 21
>> stream
{5 3 roll pop pop}
endstream
endobj
```

Example 7.20 shows a quadtone (4-component) image that uses Indexed DeviceN.³

Example 7.20 *Quadtone image using Indexed DeviceN*



1-component (“grayscale”) image



Quadtone image

In this example, we start with the grayscale image shown on the left. We want to print this with four inks: black and three Pantone spot color. The alternate color space is a simple calibrated RGB. The DeviceN color space, then, will have seven components: the four inks plus the three components of the alternate. Here is the PDF object representing the quadtone image:

```
4 0 obj <<
  /Type /XObject
  /Subtype /Image
  /Width 288
  /Height 288
  /BitsPerComponent 8
  /ColorSpace 8 0 R
>>
```

3. This method of representing duotones is used by Adobe Photoshop[®] 5.0.2 and subsequent versions, when exporting EPS files. Acrobat 4.0 exports Level 3 EPS files using this method, and it can also export Level 1 EPS files that use the “Level 1 separation” conventions of Adobe Tech Note 5044 to emit multi-tone images as calls to “customcolorimage” with overprinting, which can then be placed in page layout applications such as Adobe PageMaker[®], Adobe InDesign[™], and QuarkXPress.

```

/Length 105278
/Filter /ASCII85Decode
>> stream
... The data for the grayscale image ...
endstream
endobj

```

This is the color space used by the quadtone image:

```

8 0 obj
[/Indexed
  7 0 R           The base color space, shown below.
  255            The table has 256 entries.
  12 0 R ]       The lookup table, shown below.
endobj

7 0 obj
[/DeviceN
 [ /Black        Four inks (black + 3 spot colors).
  /PANTONE#20216#20CVC
  /PANTONE#20409#20CVC
  /PANTONE#202985#20CVC
  /None          Three components for the alternate.
  /None
  /None ]
  6 0 R           The alternate color space.
10 0 R ]         The tint transform function.
endobj

6 0 obj          The alternate color space.
[/CalRGB << /WhitePoint [1 1 1] >> ]
endobj

10 0 obj <<    The tint transform function.
/FunctionType 4
/Domain        7 input values (the DeviceN components)
  [0 1 0 1 0 1 0 1 0 1 0 1 0 1]
/Range         3 output values (the CalRGB components)
  [0 1 0 1 0 1]
/Length 28
>> stream
{7 3 roll pop pop pop pop} Discard the first 4 values.
endstream
endobj

12 0 obj <<   The lookup table for the Indexed color space
/Length 1975
/Filter [/ASCII85Decode /FlateDecode]

```

```
>> stream
8;T1BB2"M7*!"psYBt1k\gY1T<D&tO]r*F7Hga*
... Each entry in the table has 7 components.
endstream
endobj
```

7.12.11 Pattern color spaces

PDF 1.2

Colors in a **Pattern** color space are represented by Patterns (see [page 287](#)). A pattern is either a *tiling* pattern (PatternType 1) or a *shading* pattern (PatternType 2). A tiling pattern is either *colored*, in which case the colors it uses are contained in the pattern's contents stream; or it is *uncolored*, in which case its color must be specified whenever the pattern is used, and an underlying color space must be specified with the pattern color space.

A **Pattern** color space for an uncolored tiling pattern is specified as a 2-element array:

```
[/Pattern base]
```

base is the underlying color space, which may be any color space except a **Pattern** color space. A **Pattern** color space for a colored tiling pattern or a shading pattern is specified as the name **Pattern** or an array containing just the name **Pattern**. It may also be specified by the 2-element array shown above; in this case, *base* is ignored.

7.12.12 Default color spaces

PDF 1.1

When a rendering device is performing a marking operation, there is always a current color space, which is established using the operators of [Section 8.5.2](#), “Color operators,” the **ColorSpace** key of an Image XObject, or the **CS** key of an in-line image. When the current color space is **DeviceGray**, the ColorSpace dictionary of the current Resources dictionary is checked for the presence of the **DefaultGray** key. If this key is present, then the color space that is the value of that key is used as the color space for the operation currently being performed. The value of the **DefaultGray** key may be the **DeviceGray** color space, a **CalGray** color space, a **Separation** color space, or (in PDF 1.3) a 1-component **ICCBased** or **DeviceN** color space.

PDF 1.3

Similarly, when the current color space is **DeviceRGB**, the ColorSpace dictionary of the current Resources dictionary is checked for the presence of the **DefaultRGB** key. If this key is present, then the color space that is the value of that key is used as the color space for the operation currently being performed. The value of the **DefaultRGB** key may be the **DeviceRGB** color space, a **CalRGB** color space specification, or (in PDF 1.3) a 3-component **ICCBased** or **DeviceN** color space.

PDF 1.3

In PDF 1.3, when the current color space is **DeviceCMYK**, the ColorSpace dictionary of the Resources dictionary is checked for the presence of the **DefaultCMYK** key. If this key is present, then the color space that is the value of

PDF 1.3

that key is used as the color space for the operation currently being performed. The value of the **DefaultCMYK** key may be the **DeviceCMYK** color space or a 4-component **ICCBased** or **DeviceN** color space.

7.13 XObjects

XObjects are named resources. PDF supports three types of XObjects: images, forms, and pass-through PostScript language fragments. In the future it may support other object types.



XObjects are passed by name to the **Do** operator, described on [page 348](#). The action taken by the **Do** operator depends on the subtype of XObject passed to it. In the case of images and forms, the **Do** operator draws the XObject.

7.13.1 Images

An Image is an XObject whose **Subtype** is **Image**. Images allow a contents stream to specify a sampled image or image mask. PDF supports image masks, 1-bit, 2-bit, 4-bit, and 8-bit grayscale images, and color images with 1, 2, 4, or 8 bits per component. Color images may have one color component (indexed-color values or separation tints), three color components (RGB, CalRGB, or Lab), four color components (CMYK), or an arbitrary number (DeviceN).

All PDF images have a size of 1×1 unit in user space, and the data is specified left to right, top to bottom. PDF images are sized and positioned by adjusting the current transformation matrix in the contents stream.⁴

An Image XObject is specified by a stream object. The stream dictionary must include the standard keys required of all streams as well as additional ones described in [Table 7.33](#).⁵

An image mask has 1-bit samples. Each bit indicates whether paint should be applied at that position; by default, a sample value of 1 indicates that the sample is *masked*, that is, that no paint should be applied there, so the color already at that position “shows through the mask.” A sample value of 0 indicates that paint is applied there. (The **Decode** array can be used to reverse the meaning of the bit.) Used by itself, an image mask paints using the current color. In PDF 1.3, an image mask can also be used to indicate where another *image* should be painted; the mask is specified as the value of the **Mask** key of that other image. Alternatively, the **Mask** key can specify which *colors* in an image should be painted and which ones should not.

PDF 1.3







4. The format and interpretation of the sample data conform to the conventions required by the PostScript **image** and **imagemask** operators. However, PostScript images have a separate matrix, while PDF images are always defined in a 1×1 square. Also, PostScript image data is specified bottom to top, while PDF image data is specified top to bottom.

5. Several of the keys are the same as those required by the PostScript language **image** and **imagemask** operators. Matching keys have the same semantics.

Table 7.33 *Image XObject attributes*

Key	Type	Semantics
Type	name	Object type. Always XObject .
Subtype	name	(<i>Required</i>) XObject subtype. Always Image .
Name	name	(<i>Required only for compatibility with PDF 1.0</i>) Resource name, used as an operand of the Do operator. Name must match the name used in the XObject dictionary within the page's Resources dictionary.
Width	integer	(<i>Required</i>) Width of the source image in samples.
Height	integer	(<i>Required</i>) Height of the source image in samples.
BitsPerComponent	integer	(<i>Required</i>) The number of bits used to represent each color component. The value must be 1, 2, 4, or 8.
ColorSpace	name or array	(<i>Required for images, not allowed for image masks</i>) Color space used for the image samples. This may be any type of color space except Pattern . Separation color spaces are not permitted before PDF 1.2. DeviceN color spaces are not permitted before PDF 1.3.
Decode	array	(<i>Optional</i>) An array of numbers specifying the mapping from sample values in the image to values appropriate for the current color space. The number of elements in the array must be twice the number of color components in the color space specified in the ColorSpace key. The default value results in the image sample values being used directly. Decode arrays are described further on page 249 . For an image mask, the Decode array is either [0 1] or [1 0] . If it is [0 1] (the default), then 0-bits in the image data indicate where paint is applied; 1-bits are masked (not painted). If the array is [1 0] , the opposite happens: 0-bits are masked, and 1-bits are painted.
Interpolate	boolean	(<i>Optional</i>) If <i>true</i> , requests that image interpolation be performed. Interpolation attempts to smooth transitions between sample values. Interpolation may be performed differently by different devices, and not at all by some. The default value is <i>false</i> .
ImageMask	boolean	(<i>Optional</i>) Specifies whether the image should be treated as a mask. If <i>true</i> , the image is treated as a mask: BitsPerComponent must be 1, ColorSpace should not be provided, and the mask is drawn using the current fill color. If <i>false</i> , the image is not treated as a mask. The default value is <i>false</i> .
Intent	name	(<i>Optional</i>) A name which is a color rendering intent indicating the style of color rendering that should occur. For example, one might want to render images in a perceptual or pleasing manner while rendering line art colors with exact color matches. Intents are meaningful only for the device-independent color spaces. For further details, see page 252 .


PDF 1.2
PDF 1.1

OPI	dictionary	<i>(Optional)</i> An OPI dictionary. See page 258 . This key is ignored if ImageMask is <i>true</i> .	
Mask	stream or array	<i>(Optional)</i> If Mask is a stream, then the image is masked <i>by position</i> , as described in Section 7.13.5, “Masking images by position.” If it is an array, then the image is masked <i>by color</i> , as described in Section 7.13.6, “Masking images by color.” If the value of the ImageMask key is <i>true</i> , the Mask key is not permitted.	 
Alternates	array	<i>(Optional)</i> Alternate versions of this image. See page 250 . This key can only appear in a base image. It must not appear in alternate images. The order of elements in this array has no significance. Note that the attributes of the base image (e.g., OPI) do not apply to the alternates.	
ID	string	<i>(Optional; indirect reference preferred)</i> An ID that can be used to locate a SpiderContentSet array in the IDs name tree. See page 160 .	
PZ (Preferred Zoom)	number	<i>(Optional)</i> A scaling factor by which the image can be scaled to achieve the “natural” zoom level. See page 160 .	

[Example 7.21](#) shows an image object. It is a monochrome image (1 bit per component, **DeviceGray**) that is 24 samples wide and 23 samples high. Interpolation is not requested and the default decode array is used.

Example 7.21 *Image with length specified as an indirect object*

```

5 0 obj
<<
  /Type /XObject
  /Subtype /Image
  /Width 24
  /Height 23
  /BitsPerComponent 1
  /ColorSpace /DeviceGray
  /Filter /ASCIIHexDecode
  /Length 6 0 R
>>
stream
003B00 002700 002480 0E4940 114920 14B220 3CB650
75FE88 17FF8C 175F14 1C07E2 3803C4 703182 F8EDFC
B2BBC2 BB6F84 31BFC2 18EA3C 0E3E00 07FC00 03F800
1E1800 1FF800>
endstream
endobj

6 0 obj
174
endobj

```


7.13.2 Decode arrays

A **Decode** array can be used to invert the colors in an image or to compress or expand the range of values specified in the image data. Each pair of numbers in a **Decode** array specifies the upper and lower values to which the range of sample values in the image is mapped. A **Decode** array contains one pair of numbers for each component in the color space specified in the image. The mapping for each color component is a linear transformation. That is, it uses the following formula for linear interpolation:

$$y = \text{Interpolate}(x, x_{\min}, x_{\max}, y_{\min}, y_{\max})$$

$$= y_{\min} + (x - x_{\min}) \times \left(\frac{y_{\max} - y_{\min}}{x_{\max} - x_{\min}} \right)$$

Note Generally this is used to convert a value x between x_{\min} and x_{\max} to a corresponding value y between y_{\min} and y_{\max} , projecting along the line defined by the points (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) . While this formula applies to values outside the range x_{\min} to x_{\max} , and does not require that $x_{\min} < x_{\max}$, interpolation used for color-conversion, such as the **Decode** array, does require that $x_{\min} < x_{\max}$, and it “clips” x values to this range, so that for all $x \leq x_{\min}$, we have $y = y_{\min}$, and for all $x \geq x_{\max}$, we have $y = y_{\max}$.

For a **Decode** array of the form $[D_{\min} D_{\max}]$, this can be written as:

$$y = \text{Interpolate}(x, 0, 2^n - 1, D_{\min}, D_{\max})$$

$$= D_{\min} + x \times \frac{D_{\max} - D_{\min}}{2^n - 1}$$

where:

n is the value of **BitsPerComponent**

x is the input value, in the range 0 to $2^n - 1$

D_{\min} and D_{\max} are the values specified in the **Decode** array

y is the output value, to be interpreted in the color space of the image.

Samples with a value of zero are mapped to D_{\min} , samples with a value of $2^n - 1$ are mapped to D_{\max} , and samples with intermediate values are mapped linearly between D_{\min} and D_{\max} . The default **Decode** array for each color component is **[0 1]**, causing sample values in the range 0 to $2^n - 1$ to be mapped to color values in the range 0 to 1. [Table 7.34](#) shows the default Decode arrays for various color spaces.

Table 7.34 Default Decode arrays for various color spaces

Color space	Default Decode array
DeviceGray	[0 1]

DeviceRGB	[0 1 0 1 0 1]
DeviceCMYK	[0 1 0 1 0 1 0 1]
CalGray	[0 1]
CalRGB	[0 1 0 1 0 1]
Lab	[0 100 a_{\min} a_{\max} b_{\min} b_{\max}] where a_{\min} , a_{\max} , b_{\min} , and b_{\max} correspond to the entries in the Range array of the image's color space. 0 and 100 are the first two entries since the range of L^* is always 0 to 100.
Separation	[0 1]
DeviceN	[0 1 0 1 ...] where there are n pairs of 0 and 1.
Indexed	[0 N] where $N = 2^n - 1$
Pattern	(cannot be used with images)

As an example of a **Decode** array, consider a **DeviceGray** image with 8 bits per component. The color of each sample in a **DeviceGray** image is represented by a single number. The default **Decode** array, [0 1], maps a sample value of 0 to a color value of 0 and a sample value of 255 to a color value of 1. A negative image is produced by specifying a **Decode** array of [1 0], which maps a sample value of 0 to a color value of 1 and a sample value of 255 to a color value of 0. If the image contains only values from 0 to 63 and is to be displayed using the full gray range of 0 to 1, a **Decode** array of [0 4] should be used. With this **Decode** array, a sample value of 0 maps to a color value of 0, a sample value of 255 maps to a color value of 4, and a sample value of 63 (the maximum value in the example) maps to a color value of 0.99.

7.13.3 Alternate images

PDF 1.3

Alternate images provide a straightforward and backwards-compatible way to have at least one viewing-resolution and at least one printing-resolution version of each Image XObject in a PDF file. The primary goal is to substantially reduce the need for authors to maintain two versions of each PDF document (one version with low-resolution images for viewing on the Web, and a separate version with high-resolution images for printing).

In PDF 1.3, a *base image* (i.e., the Image XObject referred to in a Resources dictionary) can contain an **Alternates** key. The value of this key is an array that specifies alternate versions of the image. These alternates may differ, for example, in resolution or in color space. This permits a single PDF file to contain low-resolution versions of each image (for on-line use), and high-resolution versions (for printing).

Each alternate version is a dictionary that specifies the Image XObject for that alternate, as well as properties of the alternate. [Table 7.35](#) shows the contents of the dictionary representing an alternate image.

Table 7.35 *Alternate image dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Image	stream	<i>(Required)</i> The Image XObject for the alternate.
DefaultForPrinting	boolean	<i>(Optional; may be true in at most one alternate for each base image)</i> If <i>true</i> , this alternate is the default for printing. If no alternate has DefaultForPrinting set to <i>true</i> , then the base image is the default image for printing.

[Example 7.22](#) shows an image with a single alternate. In this example, the base image is a grayscale image. The alternate is a high-resolution RGB image stored on a Web server.

Example 7.22 *Image XObject with a single alternate*

```

10 0 obj <<
  /Type /XObject
  /Subtype /Image
  /Width 100
  /Height 200
  /BitsPerComponent 8
  /ColorSpace /DeviceGray
  /Filter /DCTDecode
  /Length 2167
  /Alternates 18 0 R
>>
stream
... image data ...
endstream
endobj

18 0 obj
[ << /Image 98 0 R /DefaultForPrinting true >> ]
endobj

98 0 obj <<
  /Type /XObject
  /Subtype /Image
  /Width 1000
  /Height 2000
  /BitsPerComponent 8
  /ColorSpace /DeviceRGB
  /FFilter /DCTDecode
  /Length 0
  /F <<
  /FS /URL
  /F (http://www.myserver.mycorp.com/im\

```

This is an external stream.

```

ages/extttest.jpg) >>
>>
stream
endstream
endobj

```

7.13.4 Color rendering intent

PDF 1.1



The supported color rendering intents and their meanings are given below in [Table 7.36](#). Other intents are permitted, but a viewer based on the PDF 1.1 specification will most likely ignore other values. The default intent is **RelativeColorimetric**.

Table 7.36 *Color rendering intents*

<i>Name</i>	<i>Semantics</i>
AbsoluteColorimetric	Requests an exact color (hue, saturation, and brightness) match. This is appropriate for uses such as some line art or spot colors. If the exact color cannot be displayed, the closest available one is substituted.
RelativeColorimetric	Requests an exact hue/saturation match, but scales the brightness range so that all brightnesses fit into the display device's brightness range. This is often appropriate for line art and spot color. As a result of the brightness scaling, the exact colors produced will differ on devices having different brightness range capabilities. If the exact hue/saturation cannot be displayed, the closest available one is substituted.
Perceptual	Scales the hue, saturations, and brightness ranges so that all values can be displayed on the output device. This generally provides a pleasing rendering of scanned images. As a result of the scaling, all colors are modified somewhat.
Saturation	Emphasizes saturation. This is appropriate for business graphics.

7.13.5 Masking images by position

PDF 1.3


Within the graphics arts industry and page-layout applications, it is common to take an image in a photograph, for example, “mask” out the background, and then place the cropped image on a different background. This is commonly done by drawing a clipping path between the pixels of the source image. Such paths need many line segments. Depending on the capability of the device on which the image is to be printed, this can lead to errors.

Masked images enable these constructs to be represented and processed more efficiently. Conventional images contain n color components per pixel; masked images contain n color components plus one mask component per pixel. Image XObjects that are masked in this way (i.e., by position) have a **Mask** key whose value is an imagemask, i.e., another Image XObject whose **ImageMask** key is *true*. This imagemask contains only the one mask component; the n color components are stored in the first Image XObject in the normal way.


Although the **Height** and **Width** of the imagemask are independent of the **Height** and **Width** of the image, both the image and the imagemask are defined on the unit square, as all PDF images are, so the rasterizer determines exactly which pixels are masked, when the image and imagemask have been mapped to device space.

The following example shows the use of an image masked by position.

Example 7.23 *An image masked by position*



The base image



The imagemask



The masked image

In this example, a bitmapped (1-component, 1 bit per component) image of a flower (center) is used to mask a color image of clouds (left), producing a “cloud-colored flower” (right).

```
5 0 obj <<
  /Type /XObject
  /Subtype /Image
  /Mask 4 0 R
  /Width 144 /Height 144
  /BitsPerComponent 8
  /ColorSpace /DeviceRGB
  /Length 78959
  /Filter /ASCII85Decode
>> stream
... RGB data for the clouds ...
endstream
endobj
```

The masked image

*The mask, shown below.
144 samples in x and y*

```
4 0 obj <<
  /Type /XObject
  /Subtype /Image
  /Width 900 /Height 900
  /BitsPerComponent 1
  /ImageMask true
>>
```

The image that serves as the mask

900 samples in x and y

```

/Length 129083
/Filter /ASCII85Decode
>>
stream
... bitmap data for the flower...
endstream
endobj

```

Note that in this example, the mask has a much higher resolution than the base image, 900 samples per unit as opposed to 144. Since all images in PDF are defined on the unit square, a single scaling factor will apply uniformly to the image and its mask.



7.13.6 Masking images by color

PDF 1.3

In addition to masking certain *positions* in an image, it is also useful to mask certain *colors* in an image. For example, an image containing a particular shade of blue might be painted over a background. If that shade of blue were the masked color, then the background would show through wherever the image contained blue pixels.

If the value of the **Mask** key is an array, then it is an array of $2 \times n$ integers, $[min_1 max_1 \dots min_n max_n]$, where n is the number of color components in the image's color space. Each integer must be in the range 0 to $2^{\text{BitsPerComponent}} - 1$, representing color values *before* decoding. Any image sample whose color value, before decoding, is $c_1 \dots c_n$, where $min_i \leq c_i \leq max_i$ for $1 \leq i \leq n$, is masked (not painted).

The following example shows an image that is masked by color.

Example 7.24 *An image masked by color*



In this example, an image of a world map (left) is printed, and then a separate image of a weather man (center) is printed at the same location. If there were no masking, then the center picture would completely obscure the map. But if the center image specifies a blue color mask, then when it is printed on top of the map, the blue samples are not painted, so the map shows through (right). In this example, the center image uses a DeviceRGB color space, and the Image XObject contains this key-value pair:

```
/Mask [0 2 40 42 136 139]
```

This specifies a narrow range of RGB colors centered around dark blue.

Why doesn't the weather man ever wear a blue tie?

This effect of masking by color is also known as *chroma-key* or *blue-screen* masking, and it is commonly used in video, where a TV weather man, for example, stands in front of a blue screen. The blue is being masked and is therefore transparent, so he appears to be standing in front of a large map or other image coming from a separate source. When an image is masked by color, *all* the image samples within that color range are transparent, regardless of where they appear in the image. [Example 7.25](#) shows what would happen if the weather man wore a tie that was within the range of masked colors.

Example 7.25 *Unintended transparency*



7.13.7 Form XObjects

A *form* is a self-contained description of any text, graphics, or sampled images that may be drawn on several pages or more than once on a single page.

Note In PDF 1.2, the term “form” also refers to a completely different object, a database described on [page 129](#) and stored in the file’s Catalog under the name **AcroForm**. There is only one of those per document. The form described here is a subtype of XObject, corresponding to forms in the PostScript language. There can be any number of these forms in a document. These forms are referred to as “Form XObjects.”

A Form XObject is specified by a PDF stream. It is a contents stream (see [page 195](#)). As usual, the stream must also include a **Length** key and may include **Filter** and **DecodeParms** keys if the stream is encoded. [Table 7.37](#) describes the attributes of a Form XObject.⁶

6. The keys in the stream dictionary correspond to the keys in a PostScript language Form dictionary. Unlike a PostScript language Form dictionary, however, the PDF Form dictionary does not contain a **PaintProc** key. Instead, the stream contents specify the painting procedure.

To draw a Form XObject, the **Do** operator is used, with the name of the Form XObject to be drawn given as an operand. As discussed in the introduction on [page 195](#), this name is mapped to an object using the current Resources dictionary.

Table 7.37 *Form XObject attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always XObject .
Subtype	name	<i>(Required)</i> XObject subtype. Always Form .
BBox	Rectangle	<i>(Required)</i> The form's bounding box, in the form coordinate system. This bounding box is used to clip the output of the form and to determine its size for caching.
FormType	integer	<i>(Required)</i> Must be 1.
Matrix	matrix	<i>(Required)</i> A transformation matrix that maps the form's coordinate space into user space.
Name	name	<i>(Required only for compatibility with PDF 1.0)</i> Resource name, used as an operand of the Do operator. Name must match the name used in the XObject dictionary within the Resources dictionary of the contents stream where the form appears.
Resources	dictionary	<i>(Optional but strongly recommended)</i> A list of the resources such as fonts and images required by this form.

PDF 1.2

In PDF 1.0 and 1.1, all named resources used in the form must be included in the Resources dictionary of each Page object on which the form appears, regardless of whether or not they also appear in the Resources dictionary of the form. It can be useful to specify them in the form's own Resources dictionary as well, in order to determine which resources are used inside the form. If a resource is included in both dictionaries, it should have the same name in both locations.

In PDF 1.2 and later versions, forms can be independent of the contents streams in which they appear, and this is strongly recommended although not required. In an independent form, the Resources dictionary of the form is required and contains all the named resources used by the form. Those resources are not "promoted" to the outer contents stream's Resources dictionary, although that stream's Resources dictionary will refer to the form itself.

XUID	array	<i>(Optional)</i> An ID that uniquely identifies the form. This allows the form to be cached after the first time it has been drawn in order to improve the speed of subsequent redraws.
-------------	-------	--

XUID arrays may contain any number of elements, all of which are integers. The first element in an XUID array is the organization ID. Forms that are used only in closed environments may use 1000000 as the organization ID. Any value can be used for subsequent elements, but the same values must not be used for different forms. Organizations that plan to distribute forms widely and wish to use XUIDs

must obtain an organization ID from Adobe Systems Incorporated, as described in [Appendix E](#). Section 5.6 of the *PostScript Language Reference Manual, Third Edition* [1] provides a further explanation of XUIDs.

OPI	dictionary	(Optional) An OPI dictionary. See Section 7.13.9 on page 258 .
PieceInfo	dictionary	(Optional) A PieceInfo dictionary. PieceInfo is described in Section 7.13.10 , “Page-Piece Dictionary.”
LastModified	Date	(Optional unless PieceInfo key is present) Time of the most recent alteration of the content of the Form XObject.

PDF 1.3

PDF 1.3

Example 7.26 Form XObject

```

6 0 obj
<<
  /Type /XObject
  /Subtype /Form
  /FormType 1
  /BBox [0 0 1000 1000]
  /Matrix [1 0 0 1 0 0]
  /Length 38
>>
stream
0 0 m 0 1000 1 1000 1000 1 1000 0 1 f
endstream
endobj

```

7.13.8 PostScript XObjects

PDF 1.1

PDF enables a document to include PostScript language fragments in a page description. These fragments are printer-dependent and take effect only when printing on a PostScript printer. They have no effect either when viewing the file or when printing to a non-PostScript printer. In addition, applications that understand PDF are unlikely to be able to interpret the PostScript language fragments. Hence, this capability should be used only if there is no other way to achieve the same result.

A PostScript XObject is an XObject whose **Subtype** key has the value **PS**. When a document is printed to a PostScript printer, the contents of the XObject’s stream replace the **Do** command that references the XObject. This stream is copied without interpretation and may include PostScript comments. In any other case, the XObject is ignored. When printing to a PostScript Level 1 printer, if the XObject contains a **Level1** key, the value of that key, which must be a stream, is used instead of the contents of the PostScript XObject’s stream.

The PostScript fragment may use Type 1 and TrueType fonts listed in the current Resources dictionary. It may not use Type 3 fonts.

Note PostScript XObjects should be used with extreme caution and only to obtain results not otherwise possible in PDF. Inappropriate use of PostScript XObjects can cause PDF files to print incorrectly.

Table 7.38 PostScript XObject attributes

Key	Type	Semantics
Type	name	Object type. Always XObject .
Subtype	name	(Required) XObject subtype. Always PS .
Level1	stream	(Optional) A contents stream to be used in place of the XObject's stream when printing to a Level 1 PostScript printer.

The PostScript XObject is not compatible with 1.0 viewers.

7.13.9 OPI dictionary

PDF 1.2

In PDF 1.2, Image and Form XObjects can contain an **OPI** key. OPI stands for *Open Prepress Interface*. It is a mechanism for representing a placeholder for an image, typically a high-resolution image. The placeholder often includes a low-resolution image or *proxy*, which might be a downsampled version of the high-resolution image. Before a document containing OPI references is printed, it typically passes through a filter known as an OPI server, which replaces the proxies with the high-resolution images.

The workflow in a prepress environment often involves several applications in areas such as graphic design, photo manipulation, word processing, page layout, and document construction. As pieces of the final document are moved from one application to another, it is useful to maintain an “indirect reference” or external pointer to the data of high-resolution images; they can be quite large, many times the size of the rest of the document.

The Open Prepress Interface, originally developed by Aldus Corporation, is represented in PostScript language programs by a set of comments surrounding the PostScript code for the proxy. In PDF, the proxy may be represented as an Image XObject if the proxy is, in fact, a single image, with no changes in the graphics state; otherwise the proxy, which can be any set of graphic objects such as a gray rectangle with text or nothing at all, is represented as a Form XObject. (For OPI 2.0, a Form is always used, because there will be changes in the graphics state.) In either case, the information that corresponds to the OPI comments in a PostScript file is stored in the XObject's dictionary under the name **OPI**.



Table 7.39 *OPI dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
<i>version</i>	dictionary	There are two standard versions of OPI, the older version 1.3 and the newer version 2.0, which is replacing 1.3. While they serve the same purpose, they represent information in different ways; an OPI dictionary contains one of the keys 1.3 or 2.0 . The value of each key is a version-specific dictionary containing the OPI information, described below.

Note The keys in an OPI dictionary are name objects, as all keys are in PDF dictionaries. An OPI dictionary would be written thus:

```

<<                                     OPI 1.3 dictionary
/1.3 20 0 R
>>

<<                                     OPI 2.0 dictionary
/2.0 27 0 R
>>

```

The following two tables describe the version-specific OPI dictionaries. For details on the meaning of these keys and the effect they have on OPI servers, refer to the *Open Prepress Interface Specifications* for both versions 1.3 and 2.0. These are available on Adobe's Web site. The tables show the relevant OPI comment for each attribute.

Table 7.40 *OPI 1.3 dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics and OPI comment</i>
Type	name	Object type. Always OPI .
Version	number	<i>(Required)</i> The version of OPI used in this dictionary. Must be the number 1.3 (not the name 1.3 , as in the OPI dictionary).
F (File) File specification		<i>(Required)</i> An external file containing image data. In PDF 1.3, the file specification must be a dictionary (strings are not permitted), and the dictionary must have a Type key with the value Filespec . (%ALDImageFilename)
CropRect	array	<i>(Required)</i> An array of integers [left top right bottom] specifying the portion of the external image that is to be included. (%ALDImageCropRect)
Color	array	<i>(Optional)</i> An array of 4 numbers and a string, [C M Y K colorname] . Default is [0 0 0 1 (Black)] . (%ALDImageColor)
ColorType	name	<i>(Optional)</i> Must be one of Spot , Process , or Separation . Default is Spot . (%ALDImageColorType)
Comments	Text	<i>(Optional)</i> Documentation. (%ALDObjectComments)



PDF 1.3

CropFixed	array	(Optional) An array of real numbers [left top right bottom] with the same semantics as CropRect . (%ALDImageCropFixed)
GrayMap	array	(Optional) An array of 2^n 16-bit numbers, where n is the number of bits per sample (1, 4, or 8), recording changes made to the brightness or contrast of the image. (%ALDImageGrayMap)
ID	string	(Optional) An identifier for TIFF images in the external file. (%ALDImageID)
ImageType	array	(Optional) An array of 2 integers, [samplesPerPixel bitsPerSample] . (%ALDImageType)
Overprint	boolean	(Optional) Default is <i>false</i> . (%ALDImageOverprint)
Position	array	(Required) An array [ll_x ll_y ul_x ul_y ur_x ur_y lr_x lr_y] of user-space coordinates, for the placement of the image on the page. (%ALDImagePosition)
Resolution	array	(Optional) An array [horizRes vertRes] specifying the resolution of the image, in samples per inch. (%ALDImageResolution)
Size	array	(Required) An array of 2 integers, [pixelsWide pixelsHigh] . (%ALDImageDimensions)
Tags	array	(Optional) An array of pairs [tagNumber tagName ...] describing the TIFF ASCII tag values in the external file. Each <i>tagNumber</i> is an integer; each <i>tagName</i> is a string. (%ALDImageASCIITag<NNN>)
Tint	number	(Optional) Default is 1. (%ALDImageTint)
Transparency	boolean	(Optional) Default is <i>true</i> . (%ALDImageTransparency)

Table 7.41 OPI 2.0 dictionary

Key	Type	Semantics
Type	name	Object type. Always OPI .
Version	number	(Required) The version of OPI used in this dictionary. Must be the number 2 or 2.0 (not the name 2.0 , as in the OPI dictionary).
F (File) File specification		(Required) An external file containing image data. (%ImageFileName)
Size	array	(Optional; see note) An array of two positive numbers, [width height] , specifying the width and height of the proxy image. (%ImageDimensions)
CropRect	array	(Optional; see note) An array of numbers [left top right bottom] specifying the portion of the external image that is to be included. It is required that $0 \leq \textit{left} < \textit{right}$ and that $0 \leq \textit{top} < \textit{bottom}$. (Note that in this coordinate space, the positive y axis extends vertically <i>downward</i> . Hence the requirement that $\textit{top} < \textit{bottom}$.) (%ImageCropRect)

Note **Size** and **CropRect** should both be specified, or both omitted. If they are specified, it is required that $right \leq width$ and that $bottom \leq height$.

IncludedImageDimensions

array (Optional) An array of two nonnegative integers, [*width height*].
(%%IncludedImageDimensions)

IncludedImageQuality

number (Optional) Legal values are 1, 2, and 3. (%%IncludedImageQuality)

Inks

name or array (Optional) Legal values include:

/full_color
/registration
[/monochrome name tint name tint ...]

where each *name* is a string and each *tint* is a real number. (%%ImageInks)

MainImage

string (Optional) (%%MainImage)

Overprint

boolean (Optional) Default is *false*. (%%ImageOverprint)

Tags

array (Optional) An array of pairs [*tagNumber tagName ...*], describing the TIFF ASCII tag values in the external file. Each *tagNumber* is an integer; each *tagName* is be either a string or an array of strings. (%%TIFFASCIITag)

7.13.10 Page-Piece Dictionary

PDF 1.3

In PDF 1.3, Form XObjects and Page objects can contain a **PieceInfo** key. The PieceInfo is a mechanism for associating an application and its private data with the Form or Page. This enables the Form or Page to retain all of the information that the application associates with the content. An example of such application data is the Layer information of Adobe Photoshop®.

Note A Form XObject drawn on a page with a **Do** operator inherits the graphics state that exists at the point of the **Do**. A Form XObject that is a Page-Piece should, in its contents stream, explicitly set all of the graphics state that it requires unless its semantics specifically include the notion of inheriting graphic state from the environment.

Table 7.42 PieceInfo dictionary

Key	Type	Semantics
<application name>	dictionary	An Application Data dictionary. See Table 7.43 .

There may be several key-value pairs in the **PieceInfo** dictionary. Each is keyed by a distinct *application name* that identifies the type of the data in the Application Data Dictionary.

Although we call the identifier an application name, it may be the name of an application or the name of a popular data-type understood by a family of applications.

Table 7.43 *Application Data dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
LastModified	Date	(Required) Time of the most recent alteration of the Form XObject or Page by this application.
Private	<i>any, typically dictionary</i>	(Optional) Private data values.

The **LastModified** value indicates when this application last altered the Form or Page. If there are several Application Data dictionaries in the PieceInfo, their **LastModified** dates can be compared with the **LastModified** date of the Form or Page to ascertain which Application Data corresponds to the current Form or Page.

If two or more Application Data dictionaries have a **LastModified** date that agrees with the Form or Page, then all of them correspond to the current content. This can be used to identify extensions of data formats. For example, if the data format **PictureEdit2** is extended and becomes **PictureEdit3**, and it is a pure extension, then by storing two Application Data dictionaries, one keyed as **PictureEdit2** and one as **PictureEdit3additions**, the Form or Page can be correctly processed either by an application looking for **PictureEdit2** data or by one looking for **PictureEdit3** data as the pairing {**PictureEdit2**, **PictureEdit3additions**}.

Note Some platforms on which editing occurs may only have an approximate value for the time. Others may not understand about time zones. This is why we arrange things so that **LastModified** dates need to be compared only for equality.

The **Private** value may be any data type. Typically it will be a dictionary. It contains all of the application's private data: the data, other than the PDF page or Form content, that represents the page-piece to the application.

7.14 Functions

PDF 1.2

PDF is not a programming language; a PDF file is not a program. However, PDF does provide several types of *function objects* that represent parameterized classes of functions, including mathematical formulas and sampled representations with

arbitrary resolution. Functions are used in two areas: device-dependent rasterization information for high-quality printing (halftoning and transfer functions), and color specification for smooth shading.

Each class of function has a *function type* that specifies the representation of the function, the set of attributes that parameterize that representation, and the additional data needed by that representation.

Functions in PDF represent static, self-contained numerical transformations. A function to add two numbers has two input values and one output value:

$$f(x_0, x_1) = x_0 + x_1$$

Similarly, a function that computes the average of two numbers and the square root of their product could be viewed as a function of two input values and two output values:

$$f(x_0, x_1) = \frac{x_0 + x_1}{2}, \sqrt{x_0 \times x_1}$$

In general, a function can take any number (m) of input values and produce any number (n) of output values:

$$f(x_0, \dots, x_{m-1}) = y_0, \dots, y_{n-1}$$

In PDF functions, all the input values and all the output values are numbers, and functions have no side effects.

Each function definition includes a *domain*, the set of legal values for the input. Some types of function also define a *range*, the set of legal values for the output. Values passed to the function are clipped to the domain, and values produced by the function are clipped to the range. For example, suppose the function $f(x) = x + 2$ is defined with a domain of $[-1 \ 1]$. If the function is called with the value 6, that value is replaced with the nearest value in the defined domain, 1, before the function is evaluated, and the result is therefore 3. Similarly, if the function $f(x_0, x_1) = 3 \times x_0 + x_1$ is defined with a range of $[0 \ 100]$, and if the values -6 and 4 are passed to the function (and are within its domain), then the value produced by the function, -14 , is replaced with 0 , the nearest value in the defined range.

Some function objects are represented as dictionaries; others are represented as streams. Since every stream includes a dictionary, every function has a dictionary.

Four types of functions are available, as indicated by the dictionary's

FunctionType entry:

- A *sampled function* (type 0) uses a table of *sample values* to represent the function. Various techniques are used to interpolate values between the sample values.
- An *exponential interpolation function* (type 2) defines a set of coefficients for an exponential function.

- A *stitching function* (type 3) is a combination of other functions, partitioned across a domain.
- A “*PostScript calculator*” function (type 4) uses a few operators from the PostScript language to describe an arithmetic expression.

All function dictionaries share the following attributes:

Table 7.44 *Attributes shared by all functions*

Key	Type	Semantics								
FunctionType	integer	(Required) The function type: <table style="margin-left: 2em;"> <tr><td>0</td><td>Sampled function</td></tr> <tr><td>2</td><td>Exponential interpolation function</td></tr> <tr><td>3</td><td>Stitching function</td></tr> <tr><td>4</td><td>PostScript calculator function</td></tr> </table>	0	Sampled function	2	Exponential interpolation function	3	Stitching function	4	PostScript calculator function
0	Sampled function									
2	Exponential interpolation function									
3	Stitching function									
4	PostScript calculator function									
Domain	array	(Required) An array of $2 \times m$ numbers, where m is the number of input values. For each i from 0 to $m - 1$, Domain _{$2i$} must be less than or equal to Domain _{$2i+1$} , and the i th input value, x_i , must lie in the interval Domain _{$2i$} $\leq x_i \leq$ Domain _{$2i+1$} . Input values outside the declared domain are clipped to the nearest boundary value.								
Range	array	(Required for type 0 functions, optional otherwise; see below) An array of $2 \times n$ numbers, where n is the number of output values. For each j from 0 to $n - 1$, Range _{$2j$} must be less than or equal to Range _{$2j+1$} , and the j th output value, y_j , must lie in the interval Range _{$2j$} $\leq y_j \leq$ Range _{$2j+1$} . Output values outside the declared range are clipped to the nearest boundary value. If the Range entry is absent, no clipping is done.								

In addition, each type of function dictionary must include attributes appropriate to the particular function type. The number of output values can usually be inferred from other attributes of the function; if not (as is always the case for type 0 functions), the **Range** attribute is required. The dimensionality of the function implied by the **Domain** and **Range** attributes must be consistent with the dimensionality implied by other attributes of the function.

7.14.1 Sampled functions (Function Type 0)

Type 0 function dictionaries use a sequence of sample values to provide an approximation for functions whose domains and ranges are bounded. The samples are organized as an m -dimensional table in which each entry has n components.

Sampled functions are highly general and offer reasonably accurate representations of arbitrary analytic functions at low expense. For example, a 1-input sinusoidal function can be represented over the range [0 180] with an average

error of only 1%, using just ten samples and linear interpolation. Two-input functions require significantly more samples, but usually not a prohibitive number, so long as the function does not have high frequency variations.

The dimensionality of a sampled function is restricted only by implementation limits. However, the number of samples required to represent high-dimensionality functions multiplies rapidly unless the sampling resolution is very low. Also, the process of multilinear interpolation becomes computationally intensive if m is greater than 2. The multidimensional spline interpolation is even more computationally intensive.

A function of Type 0 includes the following attributes:

Table 7.45 *Attributes of sampled functions (FunctionType 0)*

Key	Type	Semantics
FunctionType	integer	<i>(Required)</i> Specifies a sampled function. Always 0.
Domain	array	<i>(Required)</i> As in Table 7.44 .
Range	array	<i>(Required)</i> As in Table 7.44 .
Size	array	<i>(Required)</i> An array of m positive integers specifying the number of samples in each input dimension of the sample table.
BitsPerSample	integer	<i>(Required)</i> Specifies the number of bits used to represent each sample value. Limited to 1, 2, 4, 8, 12, 16, 24, or 32.
Order	integer	<i>(Optional; default is 1)</i> The order of interpolation between samples. Allowed values are 1 or 3, specifying linear or cubic-spline interpolation, respectively.
Encode	array	<i>(Optional)</i> An array of $2 \times m$ numbers specifying the linear mapping of input values into the domain of the function's sample table. Default value: $[0 (\mathbf{Size}_0 - 1) 0 (\mathbf{Size}_1 - 1) \dots]$.
Decode	array	<i>(Optional)</i> An array of $2 \times n$ numbers specifying the linear mapping of sample values into the range of values appropriate for the function's output values. Default value: Same as the value of Range .
other stream attributes		<i>(Optional)</i> The sample values that specify the function are provided in a stream. The stream's attributes are included in the function dictionary, as appropriate. See below for details.



The **Domain**, **Encode**, and **Size** attributes determine how the function's input variable values are mapped into the sample table. For example, if **Size** is $[21\ 31]$, the default **Encode** array is $[0\ 20\ 0\ 30]$, which maps the entire domain into the full set of sample table entries. Other values of **Encode** may be used.

To explain the relationship between **Domain**, **Encode**, **Size**, **Decode**, and **Range**, we use the following notation:

$$y = \text{Interpolate}(x, x_{\min}, x_{\max}, y_{\min}, y_{\max}) = (x - x_{\min}) \times \frac{(y_{\max} - y_{\min})}{(x_{\max} - x_{\min})} + y_{\min}$$

For a given value of x , `Interpolate` calculates the y value on the line defined by the two points (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) .

When a sampled function is called, each input value x_i , for $0 \leq i < m$, is clipped to the domain:

$$x'_i = \min(\max(x_i, \text{Domain}_{2i}), \text{Domain}_{2i+1})$$

That value is encoded:

$$e_i = \text{Interpolate}(x'_i, \text{Domain}_{2i}, \text{Domain}_{2i+1}, \text{Encode}_{2i}, \text{Encode}_{2i+1})$$

That value is clipped to the size of the sample table in that dimension:

$$e'_i = \min(\max(e_i, 0), \text{Size}_i - 1)$$

The encoded input values are real numbers, not restricted to integers. Interpolation is then used to determine output values from the nearest surrounding values in the sample table. Each output value r_j , for $0 \leq j < n$, is then decoded:

$$r'_j = \text{Interpolate}(r_j, 0, 2^{\text{BitsPerSample}} - 1, \text{Decode}_{2j}, \text{Decode}_{2j+1})$$

Finally, each decoded value is clipped to the range:

$$y_j = \min(\max(r'_j, \text{Range}_{2j}), \text{Range}_{2j+1})$$

Sample data is represented as a stream of unsigned 8-bit bytes (integers in the range 0 to 255). The bytes constitute a continuous bit stream, with the high-order bit of each byte first. Each sample value is represented as a sequence of **BitsPerSample** bits. Successive values are adjacent in the bit stream; there is no padding at byte boundaries.

For a function with multidimensional input (more than one input variable), the sample values in the first dimension vary fastest, and the values in the last dimension vary slowest. For example, for a function $f(a, b, c)$, where a , b , and c vary from 0 to 9 in steps of 1, the sample values would appear in this order: $f(0, 0, 0), f(1, 0, 0), \dots, f(9, 0, 0), f(0, 1, 0), f(1, 1, 0), \dots, f(9, 1, 0), f(0, 2, 0), f(1, 2, 0), \dots, f(9, 2, 0), f(0, 0, 1), f(1, 0, 1)$, and so on.

For a function with multidimensional output (more than one output value), the values are stored in the same order as **Range**.

[Example 7.27](#) illustrates a sampled function with 4-bit samples in an array containing 21 columns and 31 rows. The function takes two arguments, x and y , in the domain $[1, 1]$, and returns one value, z , in that same range.

Example 7.27 *Sampled function*

```

14 0 obj
<<
/FunctionType 0
/Domain [-1 1 -1 1]
/Size [21 31]
/Encode [0 20 0 30]
/BitsPerSample 4
/Range [-1 1]
/Decode [-1 1]
/Length ...
/Filter ...
>>
stream
... 651 sample values ...
endstream
endobj

```

The x argument is linearly transformed by the encoding to the domain [0 20] and the y argument to the domain [0 30]. Using bilinear interpolation between sample points, the function computes a value for z , which (because **BitsPerSample** is 4) will be in the range [0 15], and the decoding transforms z to a number in the range [-1 1] for the result. The sample array is stored in a string of 326 bytes, calculated as follows (rounded up):

$$326 \text{ bytes} = 31 \text{ rows} \times 21 \text{ samples/row} \times 4 \text{ bits/sample} \div 8 \text{ bits/byte}$$

The first byte contains the sample for the point (1, 1) in the high-order 4 bits and the sample for the point (0.9, 1) in the low-order 4 bits.


The **Decode** array can be used creatively to increase the accuracy of encoded samples corresponding to certain values in the range. For example, if the desired range of the function is [-1 1] and **BitsPerSample** is 4, the usual value of **Decode** would be [-1 1] and the sample values would be integers in the interval [0 15] (as shown in [Figure 7.6](#)). But if these values were used, the midpoint of the range, 0, would not be represented exactly by any sample value, since it would fall halfway between 7 and 8. On the other hand, if the **Decode** array were **[-1 1.1429]** ($1.1429 \cong 16/14$) and the sample values supplied were in the interval [0 14], then the desired effective range of [-1 1] would be achieved, and the range value 0 would be represented by the sample value 7.

7.14.2 Exponential Interpolation Function (FunctionType 2)**PDF 1.3**

Type 2 functions include a set of parameters that define an exponential interpolation of one input value and n output values:

$$f(x) = y_0, \dots, y_{n-1}$$

In addition to the entries listed in [Table 7.44](#), a type 2 function dictionary includes the entries listed in [Table 7.46](#).

Figure 7.6 Mapping with the Decode array**Table 7.46** Entries specific to a type 2 function dictionary

Key	Type	Semantics
FunctionType	integer	(Required) Must be 2, for an exponential interpolation function.
C0	array	(Optional) An array of n numbers defining the function result when $x = 0$ (hence the “0” in the name). Default value: [0] .
C1	array	(Optional) An array of n numbers defining the function result when $x = 1$ (hence the “1” in the name). Default value: [1] .
N	number	(Required) The interpolation exponent. Each input value x will return n values, given by $y_j = \mathbf{C0}_j + x^{\mathbf{N}} \times (\mathbf{C1}_j - \mathbf{C0}_j)$, for $0 \leq j < n$.

Values of **Domain** must constrain x in such a way that if **N** is not an integer, all values of x must be greater than or equal to 0, and if **N** is negative, no value of x may be 0.

For typical use as an interpolation function, **Domain** will be declared as **[0 1]**, and **N** will be a number greater than 0. The **Range** parameter is optional and can be used to clip the output to a desired range.

Note that when **N** is 1, the function performs a *linear* interpolation between **C0** and **C1**. This can also be expressed as a sampled function (**FunctionType** 0).

7.14.3 1-Input Stitching Function (FunctionType 3)

Type 3 functions define a “stitching” of the subdomains of several 1-input functions to produce a single new 1-input function. Since the resulting *stitching function* is a 1-input function, the domain is given by a two-element array, **[Domain₀ Domain₁]**. This domain is partitioned into *k* subdomains, as indicated by the dictionary’s **Bounds** array, which is an array of *k* – 1 numbers that obey the following relationships (with exceptions as noted below):

$$\text{Domain}_0 \leq \text{Bounds}_0 \leq \text{Bounds}_1 \leq \dots \leq \text{Bounds}_{k-2} \leq \text{Domain}_1$$

The **Bounds** array describes a series of half-open intervals, closed on the left, open on the right (except the last, which is closed on the right as well). The value of the **Functions** entry is an array of *k* functions. The first function applies to *x* values in the first subdomain, **Domain₀ ≤ x < Bounds₀**; the second function applies to *x* values in the second subdomain, **Bounds₀ ≤ x < Bounds₁**; and so on. The last function applies to *x* values in the last subdomain, which includes the upper bound: **Bounds_{k-2} ≤ x ≤ Domain₁**.

Note The **Bounds** array may contain equal elements, but a pair of equal bounds describes an interval containing no points and is therefore ignored.

The **Encode** array contains 2 × *k* numbers. A value *x* from the *i*th subdomain is encoded as follows:

$$x' = \text{Interpolate}(x, \text{Bounds}_{i-1}, \text{Bounds}_i, \text{Encode}_{2i}, \text{Encode}_{2i+1})$$

for 0 ≤ *i* < *k*. In this equation, **Bounds₋₁** means **Domain₀**, and **Bounds_{k-1}** means **Domain₁**. If the last bound, **Bounds_{k-2}**, is equal to **Domain₁**, then *x*’ is defined to be **Encode_{2i}**.

The value of *k* may be 1, in which case the **Bounds** array is empty and the single item in the **Functions** array applies to all *x* values, **Domain₀ ≤ x ≤ Domain₁**.

In addition to the entries listed in [Table 7.22](#), a type 3 function dictionary includes the entries listed in [Table 7.47](#).

Table 7.47 Entries specific to a type 3 function dictionary

Key	Type	Semantics
FunctionType	integer	(Required) Must be 3, for a stitching function.
Functions	array	(Required) An array of <i>k</i> 1-input functions making up the stitching function. The output dimensionality of all functions must be the same, and compatible with the value of Range if Range is present.
Bounds	array	(Required) An array of <i>k</i> – 1 numbers that, in combination with Domain , define the intervals to which each function from the Functions array applies. Bounds elements must be in order of increasing magnitude, and each value must be within the value of Domain .

Encode array *(Required)* An array of $2 \times k$ numbers that, taken in pairs, map each subset of the domain defined by **Domain** and the **Bounds** array to the domain of the corresponding function.

Domain must be of size 2 (that is, $m = 1$). Note that **Domain**₀ must be strictly less than **Domain**₁ unless $k = 1$.

The stitching function is designed to make it easy to combine several functions to be used within one shading pattern, over different parts of the shading's domain. The same effect could be achieved by creating separate shading dictionaries for each of the functions, with adjacent domains. However, since each shading would have similar parameters, and because the overall effect is one shading, it is more convenient to have a single shading with multiple function definitions.

Also, **FunctionType 3** provides a general mechanism for inverting the domains of 1-input functions. For example, consider a function f with a **Domain** of **[0 1]**, and a stitching function g with a **Domain** of **[0 1]**, a **Functions** array containing f , and an **Encode** array of **[1 0]**. In effect, $g(x) = f(1 - x)$.

7.14.4 PostScript Calculator Function (FunctionType 4)

PDF 1.3

A function of Type 4 is represented as a stream. The contents of the stream is a procedure written with a very small subset of the PostScript language, defined below. While the semantics are those of the corresponding PostScript operators, a PostScript interpreter is not required.

The language contains expressions involving only integers, real numbers, and booleans. There are also procedures, but these are treated as part of the syntax, not as objects. There are no composite data structures such as strings or arrays. In particular, procedures are not arrays, as they are in PostScript, and the operand stack is merely the source of input values (the operands) and destination for output values (the results). There are no variables and no names.

Table 7.48 *Operators*

arithmetic operators	abs add and atan bitshift ceiling cos cvi cvr div exp floor idiv ln log mod mul neg not or round sin sqrt sub truncate xor
boolean constants	true false
boolean operators	and eq ge gt if ifelse le lt ne not or
stack operators	copy dup exch index pop roll

The syntax includes the use of curly braces to define procedures. The only operators that apply to procedures are **if** and **ifelse**. Also, the entire definition is written as a procedure.

The syntax of the **if** and **ifelse** operators is the same as in PostScript:

boolean { *expression* } **if**

boolean { *expression* } { *expression* } **ifelse**

Example 7.28 *A halftone spot function*

```
10 0 obj
<< /FunctionType 4
/Domain [-1 1 -1 1]
/Range [-1 1]
/Length 50
>>
stream
{ 360 mul sin 2 div exch 360 mul sin 2 div add }
endstream
endobj
```

Note In Type 4 functions, the **copy** operator is defined only to operate on a non-negative integer, duplicating that number of items on the stack. The other use of **copy** in the PostScript language, copying elements between composite objects, is not defined for Type 4 functions.

The **Domain** and **Range** keys are required. The input variables constitute the initial operand stack; the items remaining on the operand stack after execution of the function are the output variables; it is an error for the number of remaining operands to differ from the number of output variables specified by the **Range**, or for any of them to be objects other than numbers.

Implementations of Type 4 functions must provide a stack of at least 100 in length. No implementation is required to provide a larger stack, and it is an error to overflow the stack.

Errors in Type 4 functions

The code that reads a Type 4 function (the “parser”) must detect and report syntactic errors. It may also be able to detect some errors that will occur when the function is used, although this is not always possible. Any errors detected by the parser are considered to be errors in the PDF file itself and are handled like other errors in the file.

The code that executes a Type 4 function (the “interpreter”) must detect and report errors. PDF does not define a representation for the errors; those details are provided by the application that processes the PDF file. The following types of errors can occur:

- stack overflow
- stack underflow
- type error (e.g., applying **not** to a real number)
- range error (e.g., applying **sqrt** to a negative number)

- undefined result (e.g., dividing by zero)

Other types of errors may also be detected.

Usage

Like all functions in PDF, Type 4 functions operate on numerical data. While any function can be sampled (Type 0), and others may be described with exponential functions (Type 2), Type 4 offers greater flexibility, and potentially greater accuracy, through a small programming language. For example, a tint-transformation function for a hexachrome (6-component) **DeviceN** color space, with an alternate color space of **DeviceCMYK**, requires a 6-in, 4-out function. If such a function were sampled with m values for each input variable, the number of samples, $4 m^6$, could be prohibitively large. In practice, such functions are often written as simple, short PostScript functions.

Type 4 functions also make it possible to include a wide variety of halftone spot functions without the loss of accuracy that comes from sampling, and without adding to the list of predefined spot functions (see [Section 7.16.3, “Predefined spot functions”](#)). All of the predefined spot functions can be written as Type 4 functions. (Some spot functions are written using the **repeat** operator, which is not supported in Type 4, but they can easily be rewritten without it. [Example 7.28](#) shows how the **DoubleDot** spot function could be written as a Type 4 function.)

7.15 Extended graphics states

PDF 1.2

A Resources dictionary may include an **ExtGState** key (“extended graphics state”), whose value is a dictionary that specifies a set of parameters in the graphics state, as shown in [Table 7.49](#). Other parameters may be added in the future. Note that many of these parameters have device-dependent effects. They should not be used in a page description that is intended to be device-independent.

To set these parameters in the graphics state, use the **gs** operator; see [Section 8.4.7, “Generic Graphics State operator.”](#) These parameters are also implicitly set in Type 2 patterns (see [page 293](#)).

Table 7.49 *ExtGState* attributes

Key	Type	Semantics
Type	name	Object type. Always ExtGState .
SA	boolean	<i>(Optional)</i> Stroke adjustment.
OP	boolean	<i>(Optional)</i> Overprint for stroke. For devices that support Separation and DeviceN color spaces, this controls whether marks made on the “color planes” specified by the colorants remove (“knock out”) marks in the corresponding positions in the unselected color planes. If overprint is <i>false</i> , they do; if it is <i>true</i> , they do not. The default is <i>false</i> .

op	boolean	<i>(Optional)</i> Overprint for fill. If this key is omitted, and the OP key is supplied, then the OP key specifies overprinting for fill as well as stroke.	PDF 1.3
OPM	integer	<i>(Optional)</i> The overprint mode. If OPM is 1, and overprinting is on (see op and OP , above), and the current color space is DeviceCMYK , then a color component value of 0 indicates that that component overprints the other colorants; otherwise it knocks out the other planes. For example, .2 0 0 .9 k has the same effect as .2 .9 scn in a DeviceN color space whose colorants are [/Cyan /Black] , since the magenta and yellow components are zero. The default value of OPM is 0.	PDF 1.3
BG	function	<i>(Optional)</i> Black-generation function, which maps the interval [0 1] to the interval [0 1].	
BG2	function or name	<i>(Optional)</i> This is the same as BG , but the value maybe also be the name Default , which specifies the black-generation function that was in effect at the start of the page. If both BG amd BG2 are specified, BG2 takes precedence.	PDF 1.3
UCR	function	<i>(Optional)</i> Undercolor removal function, which maps the interval [0 1] to [-1 1].	
UCR2	function or name	<i>(Optional)</i> This is the same as UCR , but the value maybe also be the name Default , which specifies the undercolor removal function that was in effect at the start of the page. If both UCR and UCR2 are specified, UCR2 takes precedence.	
TR	function, array, or name	<i>(Optional)</i> Transfer function, which maps the interval [0 1] to [0 1]. The value is either a single function ⁷ or an array of four functions ⁸ . The name Identity may be used to represent the identity function. (See Example 7.30 on page 282.)	
TR2	function, array, or name	<i>(Optional)</i> This is the same as TR , but the value maybe also be the name Default , which specifies the transfer function that was in effect at the start of the page. If both TR and TR2 are specified, TR2 takes precedence.	PDF 1.3
HT	halftone or name	<i>(Optional)</i> The halftone dictionary or stream; see Section 7.16, “Halftones.” Specifying the name Default has the effect of resetting the halftone parameter to its default value.	
HTP	array	<i>(Optional)</i> Halftone phase, specified as an array of two integers. (<i>Used by viewers that are based on Display PostScript</i>)	
SM	number	<i>(Optional)</i> Smoothness, specified as a number in the range 0 to 1. This parameter controls the quality of smoothly shaded output, and thus indirectly controls the rendering performance. See Section 8.4.15 on page 329.	PDF 1.3
Font	array	<i>(Optional)</i> An array [<i>font size</i>]. <i>font</i> is an indirect reference to a font See Section 7.7 on page 198. <i>size</i> is a number expressed in text-space units.	PDF 1.3
LW (Line Width)	number	<i>(Optional)</i> The line width. See the description of the w operator on page 326.	PDF 1.3
LC (Line Cap)	integer	<i>(Optional)</i> The line cap. See the description of the J operator on page 324.	PDF 1.3

7. This corresponds to the PostScript **settransfer** operator.

8. This correspond to the PostScript **setcolortransfer** operator.

LJ (Line Join)	integer	(<i>Optional</i>) The line join. See the description of the j operator on page 326 .	PDF 1.3
ML (Miter Limit)	number	(<i>Optional</i>) The miter limit. See the description of the M operator on page 327 .	PDF 1.3
D (Dash Pattern)	array	(<i>Optional</i>) An array [<i>dashesAndGaps phase</i>]. <i>dashesAndGaps</i> is itself an array. See the description of the d operator on page 325 .	PDF 1.3
FL (Flatness)	number	(<i>Optional</i>) The flatness. See the description of the i operator on page 323 .	PDF 1.3
RI (Rendering Intent)name		(<i>Optional</i>) The name of the rendering intent. See the description of the ri operator on page 333 .	PDF 1.3

[Example 7.29](#) shows two ExtGStates. In the first one, stroke adjustment is on, and it includes a transfer function that inverts its value, $f(x) = 1 - x$. In the second one, overprint is turned off, and it includes a parabolic transfer function,

$f(x) = (2x - 1)^2$, with a sample of 21 values. The domain of the transfer function, [0 1], is mapped to [0 20], and the range of the sample values, [0 255], is mapped to the range of the transfer function, [0 1].

Example 7.29 *ExtGStates*

```

2 0 obj
<<
  /Type /Page
  /Parent 6 0 R
  /Resources 33 0 R
  /Contents 3 0 R
>>
endobj

33 0 obj
<<
  /ProcSet [/PDF /Text]
  /Font << /F1 5 0 R >>
  /ExtGState << /GS1 10 0 R /GS2 12 0 R >>
>>
endobj

10 0 obj
<< /Type /ExtGState /SA true /TR 11 0 R >>
endobj

11 0 obj
<< /Size 2 /Length 7 /Filter /ASCIIHexDecode >>
stream
01 00>
endstream
endobj

```

```

12 0 obj
<< /Type /ExtGState /OP false /TR 14 0 R >>
endobj

14 0 obj
<< /Size 21 /Length 63 /Filter /ASCIIHexDecode >>
stream
FF CE A3 7C 5B 3F 28 16 0A 02 00 02 0A 16 28 3F 5B
7C A3 CE FF>
endstream
endobj

```

7.16 Halftones

PDF 1.2

PDF supports halftones of type 1, 5, 6, 10, and 16, corresponding to those same types in the PostScript language. A halftone is represented by a dictionary or a stream, depending on the type. The halftone dictionary (or the stream's dictionary) contains the same key-value pairs as a PostScript language halftone dictionary (see Section 7.4 of the *PostScript Language Reference Manual, Third Edition* [1]), with the following exceptions:

- The dictionary contains a **Type** key.
- Spot functions and transfer functions are represented by function objects.
- Threshold arrays are specified as streams.
- In Type 5 halftone dictionaries, the keys for colorants must be names; they may not be strings.

7.16.1 HalftoneName

PDF 1.3

A halftone dictionary has an optional key, **HalftoneName**, that is the name of the halftone. In PDF 1.3, if the **HalftoneName** key is specified, then all the other keys, including **HalftoneType**, are optional. At rendering time, if the device has a halftone with that name, that halftone is used. Otherwise, if the other keys are specified, such as **HalftoneType**, that halftone is used. Otherwise, the default halftone is used.

Some printer manufacturers supply halftones with proprietary threshold arrays, so it is unlikely that such halftones would be defined, with their threshold arrays, in a PDF file. However, by using the *name* of the halftone, a PDF file can refer to that halftone.

7.16.2 Type 1 halftones

Table 7.50 *Entries in a Type 1 halftone dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Always Halftone .
HalftoneType	number	<i>(Required)</i> Must be 1.
HalftoneName	string	<i>(Optional)</i> If present, this supplies the name of the halftone dictionary.
Frequency	number	<i>(Required)</i>
Angle	number	<i>(Required)</i>
AccurateScreens	boolean	<i>(Optional; the default is false)</i>
SpotFunction	function or name	<i>(Required)</i> The spot function, which maps points in “spot space” (the square whose corners are at ± 1 in x and y) to values in the interval $[-1\ 1]$. SpotFunction is represented as a function with a 2-dimensional array of sample values; or it may be the name of a predefined spot function (see 7.16.3).
TransferFunction	function	<i>(Optional)</i> The transfer function, which maps the interval $[0\ 1]$ to the interval $[0\ 1]$. If present, this overrides any transfer function specified by the TR key in the ExtGState dictionary. This is required in a Type 1 halftone dictionary that is used as an element of a Type 5 halftone dictionary for a non-primary color component. The name Identity may be used to indicate the identity function.

7.16.3 Predefined spot functions



Table 7.51 shows the predefined spot-function names. The description shows the function corresponding to the name, as well as the PostScript code. The image on the left shows the relative values of the function in spot space, indicating approximately the order in which pixels are whitened in the halftone cell, darker points corresponding to pixels that are whitened later than lighter points.

Table 7.51 *Predefined spot functions*

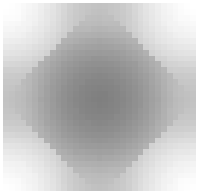
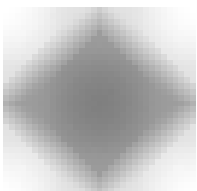
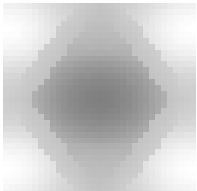
Name	Definition
Round 	<pre data-bbox="521 430 1299 556"> {abs exch abs 2 copy add 1 le { dup mul exch dup mul add 1 exch sub } { 1 sub dup mul exch 1 sub dup mul add 1 sub } ifelse } </pre> $\text{if } x + y \leq 1 \text{ then } 1 - (x^2 + y^2) \text{ else } (x - 1)^2 + (y - 1)^2 - 1$
Diamond 	<pre data-bbox="521 709 1299 934"> { abs exch abs 2 copy add .75 le { dup mul exch dup mul add 1 exch sub } { 2 copy add 1.23 le { .85 mul add 1 exch sub } { 1 sub dup mul exch 1 sub dup mul add 1 sub } ifelse } ifelse } </pre> $\text{if } x + y \leq 0.75 \text{ then } 1 - (x^2 + y^2)$ $\text{else if } x + y \leq 1.23 \text{ then } 1 - (0.85 x + y)$ $\text{else } (x - 1)^2 + (y - 1)^2 - 1$
Ellipse 	<pre data-bbox="521 1102 1299 1417"> { abs exch abs 2 copy 3 mul exch 4 mul add 3 sub dup 0 lt { pop dup mul exch .75 div dup mul add 4 div 1 exch sub } { dup 1 gt { pop 1 exch sub dup mul exch 1 exch sub .75 div dup mul add 4 div 1 sub } { .5 exch sub exch pop exch pop } ifelse } ifelse } </pre> $\text{let } w = 4 x + 3 y - 3$ $\text{if } w < 0 \text{ then } 1 - \frac{x^2 + \left(\frac{ y }{0.75}\right)^2}{4}$ $\text{else if } w > 1 \text{ then } \frac{(1 - x)^2 + \left(1 - \frac{ y }{0.75}\right)^2}{4} - 1$ $\text{else } 0.5 - w$

Table 7.51 *Predefined spot functions*

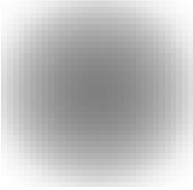
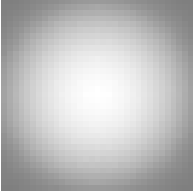
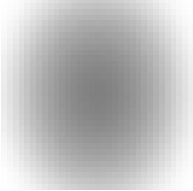
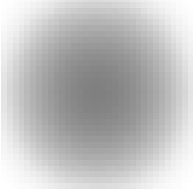
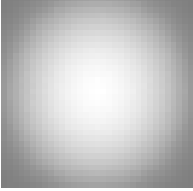
Name	Definition
EllipseA 	<pre data-bbox="521 373 1252 401">{ dup mul .9 mul exch dup mul add 1 exch sub }</pre> $1 - (x^2 + 0.9y^2)$
InvertedEllipseA 	<pre data-bbox="521 651 1170 678">{ dup mul .9 mul exch dup mul add 1 sub }</pre> $x^2 + 0.9y^2 - 1$
EllipseB 	<pre data-bbox="521 928 1284 989">{ dup 5 mul 8 div mul exch dup mul exch add sqrt 1 exch sub }</pre> $1 - \sqrt{x^2 + \frac{5}{8}y^2}$
EllipseC 	<pre data-bbox="521 1205 1252 1232">{ dup mul exch dup mul .9 mul add 1 exch sub }</pre> $1 - (0.9x^2 + y^2)$
InvertedEllipseC 	<pre data-bbox="521 1482 1170 1509">{ dup mul exch dup mul .9 mul add 1 sub }</pre> $0.9x^2 + y^2 - 1$

Table 7.51 *Predefined spot functions*

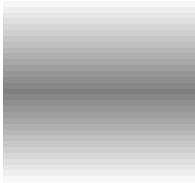
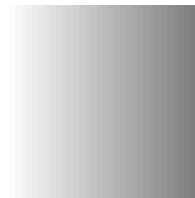

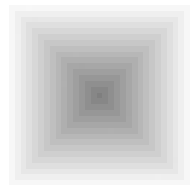
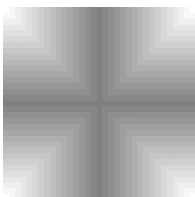
<i>Name</i>	<i>Definition</i>
Line 	<pre data-bbox="521 373 837 401">{ exch pop abs neg }</pre> $- y $
LineX 	<pre data-bbox="521 653 634 680">{ pop }</pre> x
LineY 	<pre data-bbox="521 926 708 953">{ exch pop }</pre> y
Square 	<pre data-bbox="521 1199 1252 1226">{ abs exch abs 2 copy lt { exch } if pop neg }</pre> $-\max(x , y)$
Cross 	<pre data-bbox="521 1472 1252 1499">{ abs exch abs 2 copy gt { exch } if pop neg }</pre> $-\min(x , y)$

Table 7.51 *Predefined spot functions*

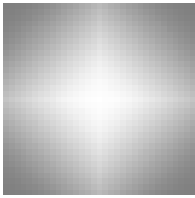
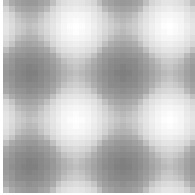
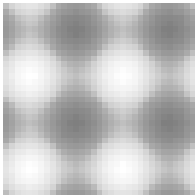
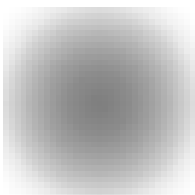
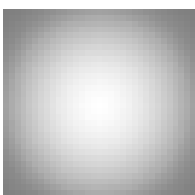
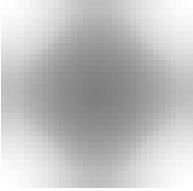
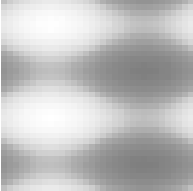
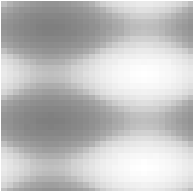
Name	Definition
Rhomboid 	<pre data-bbox="521 373 1062 401">{ abs exch abs 0.9 mul add 2 div }</pre> $\frac{0.9 x + y }{2}$
DoubleDot 	<pre data-bbox="521 653 1192 680">{ 2 {360 mul sin 2 div exch } repeat add }</pre> $\frac{\sin(360x)}{2} + \frac{\sin(360y)}{2}$
InvertedDoubleDot 	<pre data-bbox="521 932 1256 959">{ 2 {360 mul sin 2 div exch } repeat add neg }</pre> $-\left(\frac{\sin(360x)}{2} + \frac{\sin(360y)}{2}\right)$
SimpleDot 	<pre data-bbox="521 1211 1143 1239">{ dup mul exch dup mul add 1 exch sub }</pre> $1 - (x^2 + y^2)$
InvertedSimpleDot 	<pre data-bbox="521 1491 1062 1518">{ dup mul exch dup mul add 1 sub }</pre> $x^2 + y^2 - 1$

Table 7.51 *Predefined spot functions*

Name	Definition
CosineDot 	<pre data-bbox="521 373 1192 401">{ 180 mul cos exch 180 mul cos add 2 div }</pre> $\frac{\cos(180x) + \cos(180y)}{2}$
Double 	<pre data-bbox="521 651 1256 709">{ exch 2 div exch 2 { 360 mul sin 2 div exch } repeat add }</pre> $\frac{\sin\left(360\frac{x}{2}\right)}{2} + \frac{\sin(360y)}{2}$
InvertedDouble 	<pre data-bbox="521 928 1256 987">{ exch 2 div exch 2 { 360 mul sin 2 div exch } repeat add neg }</pre> $-\left(\frac{\sin\left(360\frac{x}{2}\right)}{2} + \frac{\sin(360y)}{2}\right)$

[Example 7.30](#) shows an ExtGState that includes a Type 1 halftone.

Example 7.30 *Halftone with spot-function dictionary*

```

28 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 120
/Angle 30
/SpotFunction /CosineDot
/TransferFunction /Identity
>>
endobj

```

7.16.4 Type 5 halftones

Type 5 halftones allows specification of individual halftones for an arbitrary number of color components.

Table 7.52 *Entries in a Type 5 halftone dictionary*

Key	Type	Semantics
Type	name	Object type. Must be Halftone .
HalftoneType	number	<i>(Required)</i> Must be 5.
HalftoneName	string	<i>(Optional)</i> If present, this supplies the name of the halftone dictionary.
<i>separation</i>	halftone	<i>(Required, one per separation)</i> The separation key is the name of a separation or colorant, and the value is a halftone for that separation. The halftone must be of type 1, 6, 10, or 16. Note that the key must be a name object; strings are not permitted here, although they are in PostScript type 5 halftone dictionaries. (In Example 7.31 , the separation names are /Cyan , /Magenta , /Yellow , /Black , and /Default .)
Default	halftone	<i>(Required)</i> The halftone that is to be used for any separation that does not have its own entry. This may not be a Type 5 halftone. If there are any non-primary separations, then the Default halftone must have a TransferFunction .

[Example 7.31](#) shows an ExtGState dictionary that includes a Type 5 halftone dictionary with the primary colorants for a CMYK device. In this example, the halftone dictionaries for the colorants and for the default all use the same spot function.

Example 7.31 *Halftone dictionary for type 5*

```

27 0 obj
<<
/Type /Halftone

```

```
/HalftoneType 5
/Cyan 31 0 R
/Magenta 32 0 R
/Yellow 33 0 R
/Black 34 0 R
/Default 35 0 R
>>
endobj

31 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 89.827
/Angle 15
/SpotFunction /Round
/AccurateScreens true
>>
endobj

32 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 89.827
/Angle 75
/SpotFunction /Round
/AccurateScreens true
>>
endobj

33 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 90.714
/Angle 0
/SpotFunction /Round
/AccurateScreens true
>>
endobj

34 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 89.803
/Angle 45
```

```

/SpotFunction /Round
/AccurateScreens true
>>
endobj

35 0 obj
<<
/Type /Halftone
/HalftoneType 1
/Frequency 90
/Angle 45
/SpotFunction /Round
/AccurateScreens true
>>
endobj

```

7.16.5 Type 6 halftones

A Type 6 halftone defines a halftone screen directly by specifying a threshold array at device resolution. The halftone is represented as a stream; the threshold values are stored in the stream, which must contain $width \times height$ bytes. For more details, see the *PostScript Language Reference Manual, Third Edition* [1]. The attributes are stored in the stream's dictionary.

Table 7.53 *Type 6 halftone attributes*

Key	Type	Semantics
Type	name	Object type. Must be Halftone .
HalftoneType	number	<i>(Required)</i> Must be 6.
HalftoneName	string	<i>(Optional)</i> If present, this specifies the name of the halftone dictionary.
Width	integer	<i>(Required)</i> Width of the threshold array, in pixels.
Height	integer	<i>(Required)</i> Height of the threshold array, in pixels.
TransferFunction	function	<i>(Optional)</i> Transfer function, which maps the interval [0 1] to the interval [0 1]. If present, this overrides the transfer function specified by the TR key in the ExtGState dictionary specified by the gs operator. This is required in a Type 6 halftone dictionary that is used as an element of a Type 5 halftone dictionary for a non-primary color component. The name Identity may be used to indicate the identity function.

7.16.6 Type 10 halftones

A Type 10 halftone can be used to specify a threshold array that represents a halftone cell with a non-zero screen angle. A Type 6 halftone can be used to specify a threshold array representing a zero-angle halftone cell, but there is no provision for other angles. Zero-angle halftone cells are easy to specify because they line up nicely with scan lines and because it is not difficult to determine where a sampled point goes. The Type 10 halftone applies a simple transformation to the halftone cell that converts it into two squares, thus making it easier to specify non-zero angle cells.

The halftone is represented as a stream; the threshold values are stored in the stream. There must be $Xsquare^2 + Ysquare^2$ bytes in the threshold array. For more details, see the *PostScript Language Reference Manual, Third Edition* [1]. The attributes are stored in the stream's dictionary.

Table 7.54 *Type 10 halftone attributes*


Key	Type	Semantics
Type	name	Object type. Must be Halftone .
HalftoneType	number	<i>(Required)</i> Must be 10.
HalftoneName	string	<i>(Optional)</i> If present, this specifies the name of the halftone dictionary.
Xsquare	integer	<i>(Required)</i> Length of one side of the upper square, in pixels.
Ysquare	integer	<i>(Required)</i> Length of one side of the lower square, in pixels.
TransferFunction	function	<i>(Optional)</i> Transfer function, which maps the interval [0 1] to the interval [0 1]. If present, this overrides the transfer function specified by the TR key in the ExtGState dictionary specified by the gs operator. This is required in a Type 10 halftone dictionary that is used as an element of a Type 5 halftone dictionary for a non-primary color component. The name Identity may be used to indicate the identity function.

7.16.7 Type 16 halftones

PDF 1.3

A Type 16 halftone dictionary defines a halftone screen directly, by specifying a threshold array at device resolution. This is similar to a Type 6 or Type 10 halftone dictionary, but each element of the threshold array is 16 bits wide instead of 8. This allows the threshold array to specify 65,536 levels of color instead of 256.

The threshold array is composed of one or two rectangles. The first rectangle is defined by **Width** and **Height**. The second rectangle, which is optional, is defined by **Width2** and **Height2**. If two rectangles are specified, they tile device space as illustrated in [Figure 7.7](#). For more details, see the *PostScript Language Reference Manual, Third Edition* [1].

Figure 7.7 Type 16 halftone tiling

As in Type 6 and Type 10 halftones, a Type 16 halftone is represented as a stream; the threshold values are stored in the stream. Each value is represented as two bytes, high-order byte first. If there is only one rectangle, then the array must contain $2 \times \text{Width} \times \text{Height}$ bytes. If there are two rectangles, then the array must contain $(2 \times \text{Width} \times \text{Height}) + (2 \times \text{Width2} \times \text{Height2})$ bytes; the first $2 \times \text{Width} \times \text{Height}$ bytes define the first rectangle.

Table 7.55 Type 16 halftone attributes

Key	Type	Semantics
Type	name	Object type. Must be Halftone .
HalftoneType	number	<i>(Required)</i> Must be 16.
Width	integer	<i>(Required)</i> Width of the threshold array for the first rectangle, in pixels.
Height	integer	<i>(Required)</i> Height of the threshold array for the first rectangle, in pixels.
Width2	integer	<i>(Optional)</i> Width of the threshold array for the second rectangle, in pixels.
Height2	integer	<i>(Optional)</i> Height of the threshold array for the second rectangle, in pixels.
HalftoneName	string	<i>(Optional)</i> If present, this specifies the name of the halftone dictionary.
TransferFunction	function	<i>(Optional)</i> Transfer function, which maps the interval [0 1] to the interval [0 1]. If present, this overrides the transfer function specified by the TR key in the ExtGState dictionary. The name Identity may be used to indicate the identity function.

Width2 and **Height2** must both be defined, or both omitted. If they are omitted, then the threshold array has only one rectangle.

7.17 Patterns

There are two major types of patterns in PDF. The first type is called a *tiling* pattern; the second is called a *smooth shading* pattern. Tiling patterns are procedural; they include a sequence of explicit marking instructions, which are repeated once per tile. Smooth shading patterns are declarative; they include a description of the desired effect in terms of transitions between colors across a certain area.

All patterns are treated as *colors*; a Pattern color space is installed with the **cs** and **CS** operators just as for other color spaces, and a particular pattern is installed as the current color with the **scn** and **SCN** operators. See [Section 8.5.2, “Color operators.”](#)



7.17.1 The pattern matrix

Every pattern has a **Matrix** key whose value is a transformation matrix that maps the pattern’s coordinate space to the default user space of the contents stream in which the pattern appears as a resource. For example, if a pattern is used on a page, the pattern will be listed in the Resources dictionary of that page, and the matrix maps pattern space to the default (initial) coordinate space of the page. Subsequent changes to the page’s transformation matrix, such as rotation and scaling, have no effect on the pattern; it maintains its original relationship to the page. Similarly, if a pattern is used within a Form XObject, the matrix maps pattern space to the form’s default user space, that is, the form coordinate space at the time the form is drawn with the **Do** operator. Finally, a pattern may be used within another pattern. (This is possible only if that “outer” pattern is a Type 1 pattern. A Type 2 pattern is not a contents stream and therefore has no Resources dictionary.) The inner pattern’s matrix indicates its relationship to the initial user space of the outer pattern.

Note *PostScript allows a pattern to be used in one context but defined in another. For example, a pattern might be defined on a page (that is, its matrix maps the pattern space to the user space of the page) but used in a form on that page, so that its relationship to the page is independent of each placement of the form. PDF does not support this feature. In PDF, all patterns are local to the context in which they are defined.*

7.17.2 Tiling patterns (PatternType 1)

PDF 1.2

Like a form, a tiling pattern is a contents stream: a self-contained description of text, graphics, or sampled images. Unlike forms, which are drawn one at a time, patterns are replicated (“tiled”) at fixed intervals in *x* and *y* to cover the areas to be painted.

A Type 1 Pattern is specified by a PDF stream. The additional attributes of the stream’s dictionary are shown in the following table.


Table 7.56 *Type 1 Pattern attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Type	name	Object type. Must be Pattern .
Resources	dictionary	<i>(Required)</i> A list of all the named resources required by the pattern. See page 195 .
PatternType	integer	<i>(Required)</i> Must be 1.
PaintType	integer	<i>(Required)</i> Determines how the color of the pattern cell is to be specified. The choices are: <ol style="list-style-type: none"> 1 <i>Colored pattern.</i> The stream itself specifies the colors used to paint the pattern cell. 2 <i>Uncolored pattern.</i> The stream does not specify any color information. Instead, the entire pattern cell is painted with a separately specified color each time the pattern is used.
TilingType	integer	<i>(Required)</i> Controls adjustments to the tiling to quantize it to the device pixel grid. The choices are: <ol style="list-style-type: none"> 1 <i>Constant spacing.</i> Pattern cells are spaced consistently—that is, by a multiple of a device pixel. 2 <i>No distortion.</i> The pattern cell is not distorted, but the spacing between pattern cells may vary by as much as one device pixel in both <i>x</i> and <i>y</i> dimensions when the pattern is painted. This achieves the spacing requested by XStep and YStep on average, but not for individual pattern cells. 3 <i>Constant spacing and faster tiling.</i> Like TilingType 1, but with additional distortion of the pattern cell permitted to enable a more efficient implementation.
BBox	Rectangle	<i>(Required)</i> The pattern cell bounding box, which is used to clip the pattern cell and to determine its size for caching.
XStep	number	<i>(Required)</i> The desired horizontal spacing between pattern cells, measured in the pattern cell coordinate system.
YStep	number	<i>(Required)</i> The desired vertical spacing between pattern cells, measured in the pattern cell coordinate system.
Matrix	array	<i>(Optional; the default is the identity matrix)</i> A transformation matrix that maps the pattern's coordinate space into default user space of the contents stream in which the pattern is defined. (See Section 7.17.1 , “ The pattern matrix ,” for details.)
XUID	array	<i>(Optional)</i> An extended unique ID that uniquely identifies the pattern; see section 5.6.2, “Extended Unique ID Numbers,” of the <i>PostScript Language Reference Manual, Third Edition</i> [1] for more details.

Note **XStep** and **YStep** may differ from the dimensions of the pattern cell implied by the **BBox** entry. This enables tiling with irregularly shaped figures. **XStep** and **YStep** may be either positive or negative, but not zero.

Figure 7.8 and Example 7.32 show a “bitmap pattern,” represented by a 4-by-4 cell that contains an imagemask. The pattern is used twice, once to fill a circle and once to fill a square. This example uses an uncolored tiling pattern with a DeviceGray base color space, so the shade of gray is set outside the definition of the pattern, each time the pattern is used.

Figure 7.8 *Bitmapped pattern*



The Pattern color space, **/CS9**, and the pattern, **/P1**, are defined in the Page’s Resources dictionary (object 3). The pattern’s contents stream paints an image; its Resources dictionary (object 10) therefore includes the **ImageB** ProcSet; the Page’s Resources dictionary does not.

Example 7.32 *Bitmap pattern*

```
1 0 obj <<
  /Type /Page
  /Parent 7 0 R
  /Resources 3 0 R
  /Contents 2 0 R
>> endobj
```

```
2 0 obj <<
  /Length 189
>> stream
  /Cs9 cs
  0 /P1 scn
  44 72 m
```

Set the color space.

Set the gray color (0) and the pattern (P1)

Draw a circular path.

```

144 111.744 111.744 144 72 144 c
32.256 144 0 111.744 0 72 c
0 32.256 32.256 0 72 0 c
111.744 0 144 32.256 144 72 c
f                               Fill the path with the pattern.
0.8 /P1 scn                     Set a lighter shade (.8) and the same pattern
60 60 100 100 re                 Fill the square with the pattern.
f
endstream
endobj

3 0 obj <<                       The Resources dictionary for the page.
/ProcSet [/PDF]
/Pattern << /P1 4 0 R >>
/ColorSpace << /CS1 [/Pattern /DeviceGray] >>
>>
endobj

4 0 obj <<                       The definition of the pattern.
/Type /Pattern
/PatternType 1                   A tiling pattern.
/PaintType 2                     An uncolored pattern.
/Matrix [18 0 0 18 0 0]         The pattern is scaled.
/Resources 10 0 R
/TilingType 1
/BBox [0 0 1 1]                 The pattern space is the unit square.
/XStep 1
/YStep 1
/Length 28
>>
stream
q
1 0 0 -1 0 1 cm                 The image data is stored top to bottom.
/Im1 Do                           Paint the image.
Q
endstream
endobj

10 0 obj <<                       The Resources dictionary for the pattern.
/ProcSet [/PDF /ImageB]
/XObject
  << /Im1 11 0 R >> >>         IM1 is the bitmap image.
endobj

11 0 obj <<                       The bitmap image
/Type /XObject
/Subtype /Image
/Width 4

```


```

/Height 4
/BitsPerComponent 1
/ImageMask true
/Length 10
/Filter /ASCIISHexDecode
>>
stream
6fcf6fcf
endstream
endobj

```

Figure 7.9 and Example 7.33 show an example of a colored tiling pattern⁹. The pattern is used to fill a rectangle and also to show text.

Figure 7.9 *Star pattern*



Example 7.33 *Star pattern*

```

1 0 obj <<
/Type /Page
/Parent 8 0 R
/Resources 3 0 R
/Contents 2 0 R
>> endobj

```

9. This is the same as Example 4.21 on page 255 of the *PostScript Language Reference Manual, Third Edition* [1]. The painting procedure is more detailed in PDF because it does not have the programming constructs such as **repeat** that PostScript does. Consequently, the positions, which involve constants like $\sin \pi/5$, are expressed as literal numbers, as opposed to expressions that evaluate to those numbers.

```

2 0 obj
<< /Length 18 0 R>>
stream
/Pattern cs /P1 scn
120 120 184 120 re
f
0 G
120 120 184 120 re
S
BT
/F1 1 Tf
270 0 0 270 160 100 Tm
0.9 g
(A)Tj
/CS1 cs /P1 scn
0 0 TD
(A)Tj
ET
endstream
endobj

```

*The Page contents
Set the pattern.*

Fill the rectangle with it.

Set the pattern

Show text with it

```

3 0 obj
<<
/ProcSet [/PDF /Text]
/Font << /F1 4 0 R >>
/Pattern << /P1 5 0 R >>
/ColorSpace << /CS1 [/Pattern /DeviceGray] >>
>>
endobj

```

The Page's resources

```

4 0 obj
<<
/Type /Font
/Subtype /Type1
/BaseFont /Times-Roman
>>
endobj

```

```

5 0 obj
<<
/Type /Pattern
/PatternType 1
/Resources 9 0 R
/PaintType 1 /TilingType 1
/BBox [0 0 60 60]
/XStep 60 /YStep 60
/Length 33 0 R
>>

```

The pattern

```

stream
0.3 g
15 27 m
7.947 5.292 1
26.413 18.708 1
3.587 18.708 1
22.053 5.292 1
f
45 57 m
37.947 35.292 1
56.413 48.708 1
33.587 48.708 1
52.053 35.292 1
f
0.7 g
15 57 m
7.947 35.292 1
26.413 48.708 1
3.587 48.708 1
22.053 35.292 1
f
45 27 m
37.947 5.292 1
56.413 18.708 1
33.587 18.708 1
52.053 5.292 1
f
endstream
endobj

9 0 obj
<< /ProcSet [/PDF] >>
endobj
    
```

The pattern's Resources dictionary.

7.17.3 Smooth shading patterns (PatternType 2)

PDF 1.3

Table 7.57 gives the entries in the pattern dictionary used to define the complex paints used for painting gradients. See Section 7.18, “Smooth Shading.”

Table 7.57 *Entries in a type 2 pattern dictionary*

Key	Type	Semantics
Type	name	Object type. Must be Pattern .
PatternType	integer	(<i>Required</i>) Must be 2.

Shading	dictionary	<i>(Required)</i> Defines the color transition. A Shading dictionary consists of the entries in Table 7.58 plus those in one of Table 7.59 through Table 7.64 .
XUID	array	<i>(Optional)</i> An extended unique ID that identifies the pattern; see Section 5.6.2 in the <i>PostScript Language Reference Manual, Third Edition</i> [1].
ExtGState	dictionary	<i>(Optional)</i> An Extended Graphics State dictionary representing the graphics state that is put in effect when the shading pattern is drawn. If this key is omitted, or if any key within this dictionary is omitted, the unspecified graphics state parameters will have their default values, as specified in each of the operator descriptions in Chapter 8 .
Matrix	array	<i>(Optional; the default is the identity matrix)</i> A transformation matrix that maps the pattern's coordinate space into default user space of the contents stream in which the pattern is defined. (See Section 7.17.1 , “ The pattern matrix .”)

Note Type 2 patterns do not tile. To create a tiling pattern containing a gradient, use the **sh** operator in the stream of a Type 1 pattern.

7.18 Smooth Shading

PDF 1.3

Just as the stroke adjustment parameter (see [Section 7.15](#), “[Extended graphics states](#)”) provides a resolution-independent method for requesting uniform line thickness, shading provides a way to request smooth transitions between paths and between colors, without specifying the number of steps in the transition.

Smooth shadings of objects and regions are defined in terms of *gradients*, where a gradient is a complex paint that may be used with the **fill**, **stroke**, **show**, or **imagemask** operators to paint a path or mask, using a smooth transition between colors across the areas painted. A Type 2 **Pattern** dictionary (see [Section 7.17.3](#), “[Smooth shading patterns \(PatternType 2\)](#)”) is used to define such paints. Such **Pattern** dictionaries include a **Shading** subdictionary to define the color transition used.

In many cases, the geometry of an object to be painted with a gradient will be the same as the geometry of the gradient itself. In this case, the **sh** operator is used. (See [page 338](#).) In some implementations, using **sh** with a shading dictionary may be more efficient than using **fill** with a Type 2 pattern.

All **Shading** dictionaries contain the entries listed in [Table 7.58](#), plus the entries associated with the specified **ShadingType** value ([Table 7.59](#) through [Table 7.64](#)).

Shading dictionaries also include a **Function** attribute. It is required in some cases, optional in others. In cases where it is present, the **Shading** dictionary usually defines the geometry of the shading, while the function defines the color transitions across that geometry. Functions with high spatial frequency (or discontinuous) color transitions may display aliasing effects when imaged at low effective resolutions.

Table 7.58 *Entries common in all shading dictionaries*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
ShadingType	integer	<p>(<i>Required</i>) Must be a defined ShadingType value. The attributes associated with each value of ShadingType are defined in Table 7.59 to Table 7.64. The shading types are defined as follows:</p> <ol style="list-style-type: none"> 1 Function-based shadings define the color of every point in the domain using a mathematical function. See Table 7.59. 2 Axial shadings define a color blend along a line between two points, optionally extended beyond the points by continuing the boundary colors. See Table 7.60. 3 Radial shadings define a blend between two circles, optionally extended beyond the points by continuing the boundary colors. Radial shadings are commonly used to represent three-dimensional spheres and cones. See Table 7.61. 4 Free-form Gouraud-shaded triangle meshes define a common construct used by many three-dimensional applications to represent complex colored and shaded shapes. The vertices are specified in free-form geometry. See Table 7.62. 5 Lattice-form Gouraud-shaded triangle meshes are based on the same geometrical construct as Type 4, but with vertices specified in terms of a pseudo-rectangular lattice. See Table 7.63. 6 Coons patch meshes construct shadings from one or more color patches, each bounded by four Bézier curves. See Table 7.64. 7 Tensor product patch meshes are similar to Type 6 meshes, but with 16 control points in each patch instead of 12. See Section 7.18.2.7 on page 311.
ColorSpace	name or array	<p>(<i>Required</i>) May be a color space of any type except Pattern. (Indexed color spaces require some special considerations, as defined below in Section 7.18.1, “ColorSpace: Special Considerations.”)</p>
Background	array	<p>(<i>Optional</i>) An array of color components appropriate to the ColorSpace key, specifying a single color value. If this key is present, this color is used before any painting operation involving the shading, to fill the entire area to be painted. The effect is as if the painting operation were performed twice: first with the background color and then again with the shading. (This key is used by the f operator but ignored by the sh operator.)</p>
BBox	Rectangle	<p>(<i>Optional</i>) A rectangle in the coordinate space that is current at the time the shading is imaged. If present, this bounding box is applied as a temporary clipping boundary when the shading is painted, in addition to the current clipping path and any other clipping boundaries in effect at that time.</p>
AntiAlias	boolean	<p>(<i>Optional</i>) A flag indicating whether to filter the shading function to prevent aliasing artifacts. The shading operators sample shading functions at a rate determined by the resolution of the output device. Aliasing can occur if the function is not smooth—that is, if it has a high spatial frequency relative to the sampling rate. Anti-aliasing can be computationally expensive and is usually unnecessary, since most shading functions</p>

are smooth enough, or are sampled at a high enough frequency, to avoid aliasing effects. This feature may not be implemented on some devices, in which case this flag is ignored. Default value: *false*.

7.18.1 ColorSpace: Special Considerations

Conceptually, a shading determines a color value for each individual point within the area to be painted. In practice, however, the shading may actually be used to compute color values for only a subset of the points in the target area, with the colors of the intervening points determined by interpolation between the ones computed. Implementations are free to use this strategy as long as the interpolated color values approximate those defined by the shading to within the tolerance specified by the smoothness parameter in the graphics state (see the description of the **SM** key in Table 7.49). The **ColorSpace** key common in all shading dictionaries defines not only the color space in which color values are specified in the shading, but also the color space in which color interpolation is performed.

Some shading types (4, 5, 6, and 7) perform interpolation on a parametric value supplied as input to the shading's color mapping function, as described in the relevant sections below. This form of interpolation is conceptually distinct from the interpolation described here, which operates on the output color values produced by the color mapping function and takes place within the shading's target color space.

Gradient fills between colors defined by most shadings are implemented using a variety of interpolation algorithms, and these algorithms are sensitive to the characteristics of the color space. Linear interpolation, for example, may have observably different results when applied in a **DeviceCMYK** color space than in a **Lab** color space, even if the starting and ending colors are perceptually identical. The difference arises because the two color spaces are not linear relative to each other. Shadings are rendered according to the following rules:

- If the value of the **ColorSpace** key is device-dependent and different from the native color space of the device, then color values will be converted to device colors using standard conversion formulas. To optimize performance, these conversions may take place at any time (either before or after any interpolation on the color values in the shading). Thus, any shadings defined with device-dependent color spaces may have color gradients that are somewhat device-dependent. (This does not apply to any of the vignette shadings (**ShadingType** 2 and 3), since these perform gradient calculations on a single variable and then convert to parametric colors.)
- If the value of the **ColorSpace** key is device-independent, then all gradient fill calculations will occur in the device-independent color space. Conversion to device colors will occur only after all interpolation calculations are performed. Thus, the color gradients will be device-independent for the colors generated at each point.
- If **ColorSpace** is a **Separation** or **DeviceN** color space and the specified colorants are supported, then no color conversion calculations are needed. Otherwise, the gradient calculations will be performed in the designated

Separation or **DeviceN** color space prior to conversion to the alternate color space. Thus, nonlinear *tintTransform* functions will be accommodated for the best possible representation of the shading.

- If the **ColorSpace** is an **Indexed** color space, then all color values specified in the shading will be converted to the base color space. Depending on whether the base color space is device-dependent or device-independent, gradient calculations will be performed as stated above. Interpolation never occurs in the **Indexed** color space, which is quantized and inappropriate for calculations that assume a continuous range of colors. For similar reasons, an **Indexed** color space is not allowed in any shading whose color values are generated by a function; this applies to any shading dictionary that contains a **Function** entry.

7.18.2 Shading Dictionary Entries Associated with Each Defined ShadingType

In addition to the entries listed in [Table 7.58](#), a **Shading** dictionary must have entries specific to the value of the **ShadingType** key. This section defines the available shading types.

7.18.2.1 ShadingType 1: Function-Based Shading

A Type 1 shading contains a function that maps every point in the domain to a color. This is the most general of the available shading types, and is useful for shadings that cannot be adequately described with any of the other types.

(A Type 1 shading is similar in some ways to an image, which maps the points on a fixed-resolution grid to colors.)

In addition to the common entries listed in [Table 7.58](#), a **Shading** dictionary for **ShadingType** 1 includes the following entries:

Table 7.59 *Entries in a Type 1 Shading dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
ShadingType	integer	<i>(Required)</i> Always 1. Specifies function-defined shading.
Domain	array	<i>(Optional)</i> An array of four numbers, [<i>xmin xmax ymin ymax</i>], specifying the rectangular domain of arguments with which the Function is called. Default value: [0 1 0 1] .
Matrix	array	<i>(Optional)</i> A transformation matrix. If the shading is being used as a pattern, then the matrix maps the shading’s Domain value into pattern space, and it is concatenated with the pattern’s Matrix key (see Table 7.57) to map the Domain into default user space of the current contents stream (see Section 7.17.1 , “ The pattern matrix ”). If the shading is being used with the sh operator, then the matrix maps the shading’s Domain value to the current user space. For example, if the shading’s Domain value is [0 1 0 1] and it must be mapped to a one-inch square with lower-left corner at (100, 100) in user space, the Matrix value would be [72 0 0 72 100 100] . The default is the identity matrix.

Function	dictionary or array	<i>(Required)</i> The function maps an (x,y) location to a color. It may be either a single 2-in, <i>n</i> -out function, or an array of <i>n</i> 2-in, 1-out functions, where <i>n</i> is the number of color components in the shading's ColorSpace . The function's Domain must be a superset of the shading's Domain . If the values returned by the functions are out of range for a given color component, the values will be adjusted to the nearest value in range.
-----------------	------------------------	--

Any points within the shading's **BBox** value but outside the shading's **Domain** value, as mapped into user space by **Matrix**, will be painted with the **Background** color, if that is specified, or left unpainted otherwise.

7.18.2.2 ShadingType 2: Axial Shading

ShadingType 2 shadings define a color blend that varies along a linear axis between two endpoints and extends infinitely perpendicular to that axis. The shading may optionally extend beyond the endpoints by continuing the boundary colors. In addition to the common attributes listed in [Table 7.58](#), a **Shading** dictionary for **ShadingType 2** includes the following entries:

Table 7.60 *Entries in a Type 2 Shading dictionary*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
ShadingType	integer	<i>(Required)</i> Always 2.
Coords	array	<i>(Required)</i> Array of four numbers [<i>x₀ y₀ x₁ y₁</i>] specifying the starting and ending coordinate pairs. These coordinates are expressed in pattern space (if this is being used as a pattern) or current user space (if this is being used with the sh operator).
Domain	array	<i>(Optional)</i> Array of two numbers [<i>t₀ t₁</i>]. A parametric variable <i>t</i> is considered to vary linearly between these two values, as the blend varies between the start and end points, respectively. <i>t</i> becomes the argument with which the Function is called. Default value: [0 1].
Function	dictionary or array	<i>(Required)</i> The function maps the parametric variable <i>t</i> to a color. It may be a single 1-in, <i>n</i> -out function, or an array of <i>n</i> 1-in, 1-out functions, where <i>n</i> is the number of components in the shading's ColorSpace . The function is called with values of <i>t</i> in the shading's value of Domain . The function's Domain must be a superset of the shading's Domain . If the values returned by the function is out of range for a given color component, the values will be adjusted to the nearest value in range.
Extend	array	<i>(Optional)</i> An array of two boolean values specifying whether to extend the start and end colors past the start and end points, respectively. Default value: [<i>false false</i>].

The color blend is accomplished by linearly mapping each point (x, y) along the axis between the endpoints (x_0, y_0) and (x_1, y_1) to a corresponding point in the domain specified by the shading dictionary's **Domain**. The point $(0, 0)$ in the domain corresponds to (x_0, y_0) on the axis, and $(1, 0)$ corresponds to (x_1, y_1) . Since all points along a line in domain space perpendicular to the line from $(0, 0)$ to $(1, 0)$ will have the same color, only the new value of x needs to be computed:


$$x' = \frac{(x_1 - x_0)(x - x_0) + (y_1 - y_0)(y - y_0)}{(x_1 - x_0)^2 + (y_1 - y_0)^2}$$

The value of the parametric variable t is then determined from x' as follows:

- For $0 \leq x' \leq 1$, $t = t_0 + (t_1 - t_0)x'$.
- For $x' < 0$, if the first value in the **Extend** array is *true*, then $t = t_0$; otherwise, t is undefined and the point is left unpainted.
- For $x' > 1$, if the second value in the **Extend** array is *true*, then $t = t_1$; otherwise, t is undefined and the point is left unpainted.

The value of t is then passed as the input argument to the **Function**, yielding the component values of the color with which to paint the point (x, y) .

The **ColorSpace** of this type of shading may not be an **Indexed** color space.



Note This shading is called “axial” rather than “linear” because, while the mapping from (x, y) to t is linear, the color transition defined by the **Function** key need not be.

7.18.2.3 ShadingType 3: Radial Shading

ShadingType 3 shadings define a color blend that varies between two circles. They are commonly used to depict three-dimensional spheres and cones. In addition to the attributes in [Table 7.58](#), a **Shading** dictionary for **ShadingType 3** includes the following entries:

Table 7.61 *Entries in a Type 3 Shading dictionary*

Key	Type	Semantics
ShadingType	integer	<i>(Required)</i> Always 3. Specifies radial shading.
Coords	array	<i>(Required)</i> An array of six numbers [x_0 y_0 r_0 x_1 y_1 r_1] specifying the centers and radii of the starting and ending circles. These coordinates are expressed in pattern space (if this is being used as a pattern) or current user space (if this is being used with the sh operator). r_0 and r_1 must be greater than or equal to 0. If one radius is 0, the corresponding circle is treated as a point; if both are 0, nothing is painted.
Domain	array	<i>(Optional)</i> An array of two numbers [t_0 t_1] specifying the limiting values of a parametric variable t . The variable is considered to vary linearly between these two values as the color gradient varies between the starting and ending circles. The variable t becomes the argument with which Function is called. Default value: [0 1] .
Function	dictionary or array	<i>(Required)</i> The function maps a parametric variable t to a color. It may be a single 1-in, n -out function, or an array of n 1-in, 1-out functions, where n is the number of components in the shading's ColorSpace . The function is called with values of t in the shading's Domain . The function's Domain must be a superset of the shading's Domain . If the values returned by the function are out of range for a given color component, the values will be adjusted to the nearest values in range.
Extend	array	<i>(Optional)</i> An array of two boolean values specifying whether to extend the shading beyond the starting and ending circles, respectively. Default value: [false false] .

This type of shading may not be used with the **Indexed** color space.

The color blend is based on a family of *blend circles* interpolated between the starting and ending circles that are defined by the shading dictionary's **Coords** entry. The blend circles are defined in terms of a subsidiary parametric variable

$$s = \frac{t - t_0}{t_1 - t_0}$$

which varies linearly between 0.0 and 1.0 as t varies across the domain from t_0 to t_1 , as specified by the dictionary's **Domain** entry. The center and radius of each blend circle are given by the parametric equations

$$\begin{aligned}
 x_c(s) &= x_0 + s \times (x_1 - x_0) \\
 y_c(s) &= y_0 + s \times (y_1 - y_0) \\
 r(s) &= r_0 + s \times (r_1 - r_0)
 \end{aligned}$$

Each value of s between 0.0 and 1.0 determines a corresponding value of t , which is then passed as the input argument to the shading's **Function**. This yields the component values of the color with which to fill the corresponding blend circle. For values of s not lying between 0.0 and 1.0, the boolean values in the shading dictionary's **Extend** entry determine whether and how the shading will be extended. If the first of the two boolean values is *true*, the shading is extended beyond the defined starting circle to values of s less than 0.0; if the second boolean value is *true*, the shading is extended beyond the defined ending circle to s values greater than 1.0.

Note that either of the starting or ending circles may be larger than the other. If the shading is extended at the smaller end, the family of blend circles continues as far as that value of s for which the radius of the blend circle $r(s) = 0$; if the shading is extended at the larger end, the blend circles continue as far as that s value for which $r(s)$ is large enough to encompass the shading's entire bounding box (**BBox**). Extending the shading can thus cause painting to extend beyond the areas defined by the two circles themselves.

Conceptually, all of the blend circles are painted in order of increasing values of s , from smallest to largest. Blend circles extending beyond the starting circle are painted in the same color defined by the shading's **Function** for the starting circle ($s = 0.0, t = t_0$); those extending beyond the ending circle are painted in the color defined for the ending circle ($s = 1.0, t = t_1$). The painting is opaque, with the color of each circle completely overlaying those preceding it; thus if a point lies within more than one blend circle, its final color will be that of the last of the enclosing circles to be painted, corresponding to the greatest value of s . Note the following points:

- If one of the starting and ending circles entirely contains the other, the shading will depict a sphere.
- If neither circle contains the other, the shading will depict a cone. If the starting circle is larger, the cone will appear to point out of the page; if the ending circle is larger, the cone will appear to point into the page.

Example 7.34 *Radial shading*

In this example, every leaf is filled with the same radial shading:

```
7 0 obj <<
/ShadingType 3
/ColorSpace /DeviceCMYK
/Function 9 0 R
/Coords [0 0 0.096 0 0 1]
/Extend [true true] >>
endobj
```

Concentric circles

The function is a stitching function (see [page 269](#)):

```
9 0 obj <<
/FunctionType 3
/Domain [0 1]
/Functions [10 0 R 11 0 R ]
/Encode [1 0 0 1]
/Bounds [0.708] >>
endobj
```


Its two sub-functions are both exponential interpolation functions (see [page 267](#)):

```
10 0 obj <<
/FunctionType 2
/Domain [0 1]
/C0 [0.929 0.357 1 0.298]
/C1 [0.631 0.278 1 0.027]
/N 1.048 >>
endobj
11 0 obj <<
/FunctionType 2
/Domain [0 1]
/C0 [0.929 0.357 1 0.298]
/C1 [0.941 0.4 1 0.102]
/N 1.374 >>
endobj
```

Each leaf is drawn as a path and then filled with the shade:

```
316.789 140.311 m
303.222 146.388 282.966 136.518 279.122 121.983 c
277.322 120.182 l
285.125 122.688 291.441 121.716 298.156 119.386 c
336.448 119.386 l
331.072 128.643 323.346 137.376 316.789 140.311 c
W n
q 27.7843 0 0 -27.7843 310.2461 121.1521 cm
/Sh1 sh Q
```

Sh1 is object 7, the Shading dictionary.



7.18.2.4 ShadingType 4: Free-Form Gouraud-Shaded Triangle Meshes

ShadingType 4 shadings are commonly used to represent complex colored and shaded three-dimensional shapes. The area to be shaded is defined by a path composed entirely of triangles. The color at each vertex of the triangles is specified, and a technique known as Gouraud interpolation is used to color the interiors. The interpolation functions defining the shading may be linear or nonlinear.

Unlike shading types 1, 2, and 3, shading types 4 through 7 are represented as streams. Each stream contains a sequence of vertex coordinates and color data that specify the triangle mesh. Each vertex is represented by the following sequence:

$$f\ x\ y\ c_1\ \dots\ c_n$$

where f is the *edge flag* (described below), x and y are the vertex coordinates, and c is a color value with n components. If the dictionary specifies a **Function** key, then only one color component is permitted in each sequence (see below). The vertex coordinates are expressed in pattern space (if this is being used as a pattern) or current user space (if this is being used with the **sh** operator).

The stream's dictionary contains the standard keys for a stream (e.g., **Length**), the standard keys for a shading dictionary (see [Table 7.58](#)), and the following keys:

Table 7.62 *Entries in a Type 4 Shading dictionary*

Key	Type	Semantics
ShadingType	integer	<i>(Required)</i> Always 4. Specifies free-form Gouraud-shaded triangle mesh.
BitsPerCoordinate	integer	<i>(Required)</i> Specifies the number of bits used to represent each vertex coordinate. The data is decoded based on the first four elements of the Decode key, in much the same way as color component data is decoded in image dictionaries. Allowed values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	<i>(Required)</i> Specifies the number of bits used to represent each color component. The data is decoded based on all but the first four elements of the Decode key, as in image dictionaries. Allowed values are 1, 2, 4, 8, 12, and 16.
BitsPerFlag	integer	<i>(Required)</i> The number of bits used to represent the edge flag for each vertex. The allowed values of BitsPerFlag are 2, 4, and 8. Allowed values for the edge flag itself are 0, 1, and 2.
Decode	array	<i>(Required)</i> Specifies how to decode coordinate and color component values into the ranges of numbers appropriate for these values. The ranges are specified as: $[x_{\min}\ x_{\max}\ y_{\min}\ y_{\max}\ c_{1,\min}\ c_{1,\max}\ \dots\ c_{n,\min}\ c_{n,\max}]$ where n is either the number of color components, or 1 if Function is specified.


Function

dictionary or array (Optional) A single 1-in, n -out function, or an array of n 1-in, 1-out functions, where n is the number of components in the shading's **ColorSpace**. If this key is present, the color data for each vertex must be specified by a single parametric variable rather than by n separate color components; the designated functions will be called with each interpolated value of the parametric variable to determine the actual color at each point. The function's **Domain** value must be a superset of the range specified in the shading's **Decode** array for the color value. The **Function** key may not be used with an **Indexed** color space. If the values returned by the functions are out of range for a given color component, the values will be adjusted to the nearest values in range.

Triangles are built up as described below.

The first vertex must have an edge flag value 0, which means “start a new triangle.” The edge flags of the next two vertices are ignored. Subsequent triangles are defined by a single new vertex and an edge of the preceding triangle (see [Figure 7.10](#)). Given triangle (v_a, v_b, v_c) , where the vertices v_a , v_b , and v_c appear in that order in the stream, the next vertex, v_d , forms either the new triangle (v_b, v_c, v_d) if $f = 1$ or the triangle (v_a, v_c, v_d) if $f = 2$.

Figure 7.10 How the value of edge flag, f , affects which edge is used



The stream must provide a whole number of triangles with appropriate vertex edge flags. The encoding is specified by the **BitsPerFlag**, **BitsPerCoordinate**, **BitsPerComponent**, and **Decode** attributes. Each vertex edge flag is expressed in **BitsPerFlag** bits, each vertex coordinate pair is expressed in $2 \times$ **BitsPerCoordinate** bits, and each vertex color tuple is expressed in $n \times$ **BitsPerComponent** bits. Each set of vertex data takes an integer number of bytes; if the total number of bits in the vertex data is not divisible by eight, the vertex data is padded with ignored bits inserted between the coordinate and color data. The coordinate and color data is decoded based on the **Decode** array.

If a value of the **Function** key is present, the mesh's vertex color data must be specified by single values (t), rather than color tuples. All linear interpolation within the triangle mesh will be done using the values of t , and after interpolation, the function will be called with each value to determine the color of each point.

Note Using free-form Gouraud-shaded triangle meshes differs from using an **Indexed** color space for the shading. If an **Indexed** color space is used, the vertex color value is converted to the base color space first, and linear interpolation occurs in that color space. Thus, there is no opportunity to effect a nonlinear interpolation using an **Indexed** color space.

7.18.2.5 ShadingType 5: Lattice-Form Gouraud-Shaded Triangle Meshes

ShadingType 5 shadings are similar to **ShadingType** 4 shadings, but instead of using free-form geometry, their vertices are arranged in a pseudorectangular lattice, which is topologically equivalent to a rectangular grid. The vertices are organized into “rows.” (The rows do not need to be geometrically linear.) Given m rows of n vertices, triangles are constructed using the following triplets of vertices:

$$(V_{i,j}, V_{i,j+1}, V_{i+1,j}) \quad \text{for } 0 \leq i \leq m-2, 0 \leq j \leq n-2$$

$$(V_{i,j+1}, V_{i+1,j}, V_{i+1,j+1}) \quad \text{for } 0 \leq i \leq m-2, 0 \leq j \leq n-2$$

In addition to the common attributes listed in [Table 7.58](#), a **Shading** dictionary for **ShadingType** 5 includes the following entries:

Table 7.63 Entries in a Type 5 Shading dictionary

Key	Type	Semantics
ShadingType	integer	(Required) Always 5. Specifies lattice-form Gouraud-shaded triangle mesh.
BitsPerCoordinate	integer	(Required) Specifies the number of bits used to represent each vertex coordinate. The data is decoded according to the value of Decode , in much the same way as color component data is decoded in image dictionaries. Allowed values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	(Required) Specifies the number of bits used to represent each color component. The data is decoded according to the value of Decode . Allowed values are 1, 2, 4, 8, 12, and 16.
Decode	array	(Required) Specifies how to decode coordinate and color component values into the ranges of numbers appropriate for these values. The ranges are specified as: $[x_{\min} x_{\max} y_{\min} y_{\max} c_{1,\min} c_{1,\max} \dots c_{n,\min} c_{n,\max}]$ where n is the number of color components.
VerticesPerRow	integer	(Required) The number of vertices in each “row” of the lattice; it must be greater than or equal to 2. The number of rows does not need to be specified.
Function	dictionary or array	(Optional) A single 1-in, n -out function, or an array of n 1-in, 1-out functions, where n is the number of components in the ColorSpace entry. If present, then the mesh’s vertex color data must be specified by single values rather than color tuples, and the function will be called with each interpolated color value to

determine the actual color of each point. The function's **Domain** must be a superset of the range specified in the shading's **Decode** array for the color value. **Function** may not be used with unencoded vertex data, nor with **ColorSpace** set to **Indexed** color space. If the values returned by the functions are out of range for a given color component, the values will be adjusted to the nearest allowed value.

7.18.2.6 ShadingType 6: Coons patch meshes


ShadingType 6 shadings are constructed from one or more color patches, each bounded by four Bézier curves. Degenerate Bézier curves are allowed and are useful for certain graphical effects. At least one complete patch must be specified. A primary use of this feature is to allow the specification of “conical” vignettes and other complex blends as patch meshes with nonlinear interpolation functions. The geometry of these patches is defined below, followed by the specification of the entries in the dictionary associated with **ShadingType 6**.

A Coons patch generally has two independent aspects, a color specification and a coordinate mapping. These are defined as follows:

- Colors are specified for each of the corners of the unit square, and bilinear interpolation is used to fill in colors over the entire unit square.
- Coordinates are mapped from the unit square into a four-sided patch whose sides are not necessarily linear. The mapping is continuous; the corners of the unit square map to corners of the patch, and the sides of the unit square map to sides of the patch. (See [Figure 7.11](#).)

Thus, **ShadingType 6** shading results from coloring the unit square and then mapping it.

Figure 7.11 *Coons patch meshes: coordinate mapping from a unit square to a four-sided patch*



The sides of the patch are given by four cubic Bézier curves, C_1 , C_2 , D_1 , and D_2 , defined over a pair of parametric variables u and v that vary horizontally and vertically across the unit square. The four corners of the Coons patch satisfy the equations

$$C_1(0) = D_1(0)$$

$$C_1(1) = D_2(0)$$

$$C_2(0) = D_1(1)$$

$$C_2(1) = D_2(1)$$

Two surfaces can be described that are linear interpolations between the boundary curves. Along the u axis, the surface S_C is defined by

$$S_C(u, v) = (1 - v) \times C_1(u) + v \times C_2(u)$$

Along the v axis, the surface S_D is given by

$$S_D(u, v) = (1 - u) \times D_1(v) + u \times D_2(v)$$

A third surface is the bilinear interpolation of the four corners:

$$S_B(u, v) = (1 - v) \times [(1 - u) \times C_1(0) + u \times C_1(1)] \\ + v \times [(1 - u) \times C_2(0) + u \times C_2(1)]$$


The coordinate mapping for the shading is given by the surface S , defined as

$$S = S_C + S_D - S_B$$

This defines the geometry of each patch. A patch mesh is constructed from a sequence of one or more such colored patches.

Patches can sometimes appear to fold over on themselves—for example, if a boundary curve intersects itself. As the value of parameter u or v increases in parameter space, the location of the corresponding pixels in device space may change direction, so that new pixels are mapped onto previous pixels already mapped. If more than one point (u, v) in parameter space is mapped to the same point in device space, the point selected will be the one with the largest value of v ; if multiple points have the same v , the one with the largest value of u will be selected. If one patch overlaps another, the patch that appears later in the data source paints over the earlier one.

Note also that the patch is a control surface, rather than a painting geometry. The outline of a projected square (that is, the painted area) may not be the same as the patch boundary if, for example, the patch folds over on itself, as shown in the following figure.



In addition to the common attributes listed in [Table 7.58](#), a **Shading** dictionary associated with **ShadingType** 6 includes the following entries:

Table 7.64 Entries in a Type 6 Shading dictionary

Key	Type	Semantics
ShadingType	integer	<i>(Required)</i> Always 6. Specifies a Coons patch mesh.
BitsPerCoordinate	integer	<i>(Required)</i> Specifies the number of bits used to represent each geometric coordinate. The data is decoded based on the Decode entry, in the same way color component data is decoded in image dictionaries. Allowed values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	<i>(Required)</i> Specifies the number of bits used to represent each color component. The data is decoded based on the value of the Decode entry, as in image dictionaries. Allowed values are 1, 2, 4, 8, 12, and 16.
BitsPerFlag	integer	<i>(Required)</i> The number of bits used to represent the edge flag for each patch. The allowed values of BitsPerFlag are 2, 4, and 8. Allowed values for the edge flag itself are 0, 1, 2, and 3.
Decode	array	<i>(Required)</i> Specifies how to decode coordinate and color component values into the ranges of numbers appropriate for these values. The ranges are specified as: $[x_{\min} x_{\max} y_{\min} y_{\max} c_{1,\min} c_{1,\max} \cdots c_{n,\min} c_{n,\max}]$ where n is the number of color components.
Function	dictionary or array	<i>(Optional)</i> A single 1-in, n -out function, or an array of n 1-in, 1-out functions, where n is the number of components in the ColorSpace entry. If present, then the mesh's vertex color data must be specified by single values rather than color

tuples, and the function(s) will be called with each interpolated color value to determine the actual color of each point. The value of function's **Domain** must be a superset of the range specified in the value of shading's **Decode** array for the color value. The **Function** key may not be used with unencoded vertex data, nor with the value of **ColorSpace** set to **Indexed** color space. If values returned by the functions are out of range for a given color component, the values will be adjusted to the nearest allowed value.

The stream provides a sequence of Bézier control points and color values that define the shape and colors of each patch. All of a patch's control points are given first, followed by the color values for its corners. Note that this differs from a triangle mesh (shading types 4 and 5), in which the coordinates and color of each vertex are given together. All control point coordinates are expressed in the current coordinate space at the time the shading is painted. If the total number of data bits required to define the patch is not divisible by 8, the last byte is padded at the end with extra bits, which are ignored.

As in free-form triangle meshes (type 4), each patch has an *edge flag* that tells which edge, if any, it shares with the previous patch. An edge flag of 0 begins a new patch, unconnected to any other. This must be followed by 12 pairs of coordinates, $x_1 y_1 x_2 y_2 \dots x_{12} y_{12}$, which specify the Bézier control points that define the four boundary curves. Figure 7.13 shows how these control points correspond to the Bézier curves $C_1, C_2, D_1,$ and D_2 identified earlier in Figure 7.12. Color values are then given for the four corners of the patch, in the same order as the control points corresponding to the corners. Thus, c_1 is the color at coordinates (x_1, y_1) , c_2 at (x_4, y_4) , c_3 at (x_7, y_7) , and c_4 at (x_{10}, y_{10}) , as shown in the figure.

Figure 7.12 Edge flags




Figure 7.13 also shows how nonzero values of the edge flag ($f = 1, 2,$ or 3) connect a new patch to one of the edges of the previous patch. In this case, some of the previous patch's control points serve implicitly as control points for the new patch as well, and so are not explicitly repeated in the data source. Table 7.65 summarizes the required data values for various values of the edge flag.

Figure 7.13 Coons patch




Table 7.65 Data values in a Coons patch mesh

Edge flag	Next set of data values
$f = 0$	$x_1 y_1 x_2 y_2 x_3 y_3 x_4 y_4 x_5 y_5 x_6 y_6 x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_1 c_2 c_3 c_4$
	New patch; no implicit values
$f = 1$	$x_5 y_5 x_6 y_6 x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_3 c_4$
	Implicit values:

$(x_1, y_1) = (x_4, y_4)$ previous $c_1 = c_2$ previous
 $(x_2, y_2) = (x_5, y_5)$ previous $c_2 = c_3$ previous
 $(x_3, y_3) = (x_6, y_6)$ previous
 $(x_4, y_4) = (x_7, y_7)$ previous

$f = 2$ $x_5 \ y_5 \ x_6 \ y_6 \ x_7 \ y_7 \ x_8 \ y_8 \ x_9 \ y_9 \ x_{10} \ y_{10} \ x_{11} \ y_{11} \ x_{12} \ y_{12}$
 $c_3 \ c_4$

Implicit values:

$(x_1, y_1) = (x_7, y_7)$ previous $c_1 = c_3$ previous
 $(x_2, y_2) = (x_8, y_8)$ previous $c_2 = c_4$ previous
 $(x_3, y_3) = (x_9, y_9)$ previous
 $(x_4, y_4) = (x_{10}, y_{10})$ previous

$f = 3$ $x_5 \ y_5 \ x_6 \ y_6 \ x_7 \ y_7 \ x_8 \ y_8 \ x_9 \ y_9 \ x_{10} \ y_{10} \ x_{11} \ y_{11} \ x_{12} \ y_{12}$
 $c_3 \ c_4$

Implicit values:

$(x_1, y_1) = (x_{10}, y_{10})$ previous $c_1 = c_4$ previous
 $(x_2, y_2) = (x_{11}, y_{11})$ previous $c_2 = c_1$ previous
 $(x_3, y_3) = (x_{12}, y_{12})$ previous
 $(x_4, y_4) = (x_1, y_1)$ previous

7.18.2.7 ShadingType 7: Tensor Product Patch Meshes

ShadingType 7 shadings are identical to **ShadingType 6** shadings, except that instead of a bicubic Coons patch defined by twelve control points, a bicubic tensor product patch defined by sixteen control points is used. Each set of twelve coordinate pairs in the stream is replaced by a set of sixteen coordinate pairs. The two cases are distinguished only by the **ShadingType** key and the data in the stream; there are no other differences, except in the way the color varies across the geometry.

As with the Coons patch surface, the tensor product surface is defined by a mathematical mapping from a square patch (u, v) to the patch coordinate system (x, y) .

$$S(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 P_{ij} \times B_i(u) \times B_j(v)$$

P_{ij} is the control point for the i, j row and column of the tensor. Since each $P_{ij} = (X_{ij}, Y_{ij})$, one can also express this as

$$x(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 X_{i,j} \times B_i(u) \times B_j(v)$$

$$y(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 Y_{i,j} \times B_i(u) \times B_j(v)$$

$B_i(u)$ and $B_j(v)$ are Bernstein polynomials, where

$$B_0(t) = (1-t)^3$$

$$B_1(t) = 3t(1-t)^2$$


$$B_2(t) = 3t(1-t)$$

$$B_3(t) = t^3$$

An illustration may help the mathematical description. P_{ij} are defined as follows:

P00	P01	P02	P03
P10	P11	P12	P13
P20	P21	P22	P23
P30	P31	P32	P33

Graphically, this is as represented as follows:



This is a convenient numbering for the mathematical description, but a better numbering for the language is:

0	11	10	9
1	12	15	8
2	13	14	7
3	4	5	6

This allows the Coons patch numbering to be a subset of the tensor product patch numbering.

The tensor product patch mapping, like the Coons patch mapping, is controlled by the location and shape of the four boundary curves. Unlike the Coons patch, however, the tensor product patch has four more “internal” control points to adjust the mapping. For example, imagine a cubic Bézier curve moving from the top boundary of the patch to the bottom. In effect, the control points of this curve start out as the top curve’s control points and end up as the bottom curve’s control points. The curve changes shape as it moves down the patch because the control points change position. Each control point follows a trajectory defined by the four column control points. In particular, each column defines its own cubic Bézier curve, and this is the trajectory each of the control points of the moving curve take. The same can be said for a curve moving from left to right, except the trajectory taken by the control points is defined by the curves specified by four rows of control points.

The tensor product patch gives more control over mapping than does the Coons patch. However, the Coons patch is easier to use and more concise because the internal control points are implicitly specified by the boundary control points.

7.18.3 Painting With a Pattern Dictionary

When a painting operation is performed with a gradient pattern dictionary as the current color in the graphics state, the pattern coordinate space is obtained in the same way as with **PatternType 1** patterns. However, instead of executing the marking operations in a **PatternType 1** contents stream, the shape and color transitions described by the **Shading** dictionary are interpreted relative to this coordinate space to create a logical “paint” with which graphical objects can be rendered. The **Shading** dictionary’s **BBox** attribute, if present, will clip the logical painting region, which may also be affected by the geometry of the shading. All color values are interpreted relative to the **Shading** dictionary’s **ColorSpace** attribute.

If a **Background** color is defined in the **Shading** dictionary, then that color is used first to fill the background of the object being painted. Logically, this is equivalent to executing a **fill** operation or other painting operation first with the background color and then with the gradient pattern. The convenience of the **Background** color specification is provided because performing this sequence of operations explicitly would be verbose, especially for text or imagemasks. This is only of interest for gradient patterns that do not cover the entire area of the object being painted.

7.19 Property lists

PDF 1.2

A Property List is a dictionary that provides information about a sequence of graphics objects or a particular place with the stream of marking operators. It may be written in-line, or it may appear in the **Properties** sub-dictionary of the current Resources dictionary. Property lists are used by the Marked Content operators (see [Section 8.10.3 on page 351](#)).

Property Lists provide an open-ended extension mechanism that allows a set of key-value pairs to be associated with particular marks in a content stream. With the exception of the **Subtype** key, no particular keys are defined by PDF. It is suggested, however, that any particular extension use keys in a consistent way. For example, the values associated with a given key should always be of the same type (or small set of types). The meanings of the extended keys are determined by the extensions that create and use those keys.

Table 7.66 *Property List attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Subtype	name	(Optional) Object subtype. Intended to indicate the application or extension that defines the property list.

7.20 Name tree

A name tree contains three kinds of nodes, a *root* node (one per tree), *intermediate* nodes, and *leaf* nodes. The root node contains only a **Kids** array or a **Names** array. An intermediate node contains a **Kids** array and a **Limits** array. A leaf node contains a **Limits** array and a **Names** array. The names in the tree are stored in the leaf nodes, sorted by binary value, where each string is interpreted as a sequence of unsigned octets (8-bit bytes). The names are sorted in ascending order, like other strings. Shorter names appear before longer names that begin with the same byte sequence. The encoding of the names is not important as long as it is consistent because comparisons for name equality are done byte by byte without any transformation of the bytes.

Note The keys in name trees are strings. The strings in name trees are sometimes displayed, however, as *Text*.

Table 7.67 *The root node in a name tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Kids	array	(Required if Names is not present) An array of indirect references to the immediate children of the root node. They may be intermediate nodes or leaf nodes.
Names	array	(Required if Kids is not present) An array of the form [name value name value ...] , where each name is a string, and the value is an indirect reference to the object associated with that name. The names are sorted in binary order as described above.

PDF 1.3

Table 7.68 *An intermediate node in a name tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Kids	array	An array of indirect references to the immediate children of this node. They may be intermediate nodes or leaf nodes.
Limits	array	An array of two strings, representing the (alphabetically) least and greatest names included in any leaf nodes that are descendants of this node.

Table 7.69 *A leaf node in a name tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Limits	array	An array of two strings, representing the (alphabetically) least and greatest names in the Names array.
Names	array	An array of the form <code>[name value name value ...]</code> , where each name is a string, and the value is an indirect reference to the object associated with that name. The names are sorted in binary order as described above.

As an example of a name tree, consider a document that has named destinations for all the chemical elements, where the name of the destination is the name of the element itself, from actinium to zirconium. An outline of the tree, showing object numbers and nodes, might look like this:

```
#3: root
  #4: intermediate: actinium to phosphorus
    #6: leaf: actinium = ...
    #7: leaf: ...
    #8: intermediate: ... to gadolinium
    #9: leaf: gallium = #14, germanium = #15, gold = #16
      #14: destination: ...
      #15: destination: [/FitR 0 0 100 100]
      #16: destination: ...
    #10: intermediate: hafnium to phosphorus
  #5: intermediate: platinum to zirconium
    #11: leaf: platinum = #17, plutonium = #18, polonium = #19
      #17: destination: ...
      #18: destination: ...
      #19: destination: [/XYZ 50 50 800]
    #12: intermediate: potassium to ...
    #13: intermediate: ... to zirconium
```

The representation of this tree in the PDF file would look like this:

```
1 0 obj <<
  /Type /Catalog
  /Names 2 0 R
```

```

>> endobj

2 0 obj <<                                     Dictionary of Name Trees
/Dests 3 0 R
>> endobj

3 0 obj <<                                     Root node
/Kids [4 0 R 5 0 R]
>> endobj

4 0 obj <<                                     Intermediate node
/Limits [(actinium) (phosphorus)]
/Kids [6 0 R 7 0 R 8 0 R 9 0 R 10 0 R]
>> endobj

5 0 obj <<                                     Intermediate node
/Limits [(platinum) (zirconium)]
/Kids [11 0 R 12 0 R 13 0 R]
>> endobj

9 0 obj <<                                     Leaf node
/Limits [(gallium) (gold)]
/Names [(gallium) 14 0 R (germanium) 15 0 R (gold)
16 0 R]
>> endobj

11 0 obj <<                                     Leaf node
/Limits [(platinum) (polonium)]
/Names [(platinum) 17 0 R (plutonium) 18 0 R
(polonium) 19 0 R]
>> endobj

15 0 obj <<                                     Destination
/D [FitR 0 0 100 100]
>> endobj

19 0 obj <<                                     Destination
/D [XYZ 50 50 800]
>> endobj

```

7.21 Number tree

PDF 1.3

A number tree has the same purpose as a dictionary, mapping keys to values, but it does so in a different manner, and the keys in a number tree are integers, not PDF name objects.

A number tree contains three kinds of nodes, a *root* node (one per tree), *intermediate* nodes, and *leaf* nodes. The root node contains only a **Kids** or **Nums** array. An intermediate node contains a **Kids** array and a **Limits** array. A leaf node contains a **Limits** array and a **Nums** array. The numbers in the tree are stored in the leaf nodes, sorted in ascending order.

Table 7.70 *The root node in a number tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Kids	array	(Required if Nums is not present) An array of indirect references to the immediate children of the root node. They may be intermediate nodes or leaf nodes.
Nums	array	(Required if Kids not present) An array of the form <code>[integer value integer value ...]</code> , where each <i>value</i> is an indirect reference to the object associated with the preceding <i>integer</i> key. The integers are sorted in ascending order.

Either **Kids** or **Nums** must be specified in the root node, but not both.

Table 7.71 *An intermediate node in a number tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Kids	array	An array of indirect references to the immediate children of this node. They may be intermediate nodes or leaf nodes.
Limits	array	An array of two integers, representing the least and greatest numbers included in any leaf nodes that are descendants of this node.

Note Intermediate nodes can have children that are also intermediate nodes. This allows an essentially unlimited number of integer/value pairs. By contrast, an array has a maximum length (see [Appendix D](#)).

Table 7.72 *A leaf node in a number tree*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
Limits	array	An array of two integers, representing the least and greatest numbers in the Nums array.
Nums	array	An array of the form <code>[integer value integer value ...]</code> , where each <i>value</i> is an indirect reference to the object associated with the preceding <i>integer</i> key. The integers are sorted in ascending order.

CHAPTER 8

Page Descriptions

This chapter describes the PDF operators that draw text, graphics, and images on the page and in other “marking contexts” such as forms and patterns. It completes the specification of PDF. The following chapters describe how to produce efficient PDF files.

Text, graphics, and images are drawn using the coordinate systems described in [Chapter 3](#). It may be useful to refer to that chapter when reading the description of various operators, to obtain a better understanding of the coordinate systems used in PDF documents and the relationships among them.

[Appendix B](#) contains a complete list of operators, arranged alphabetically.

Note Throughout this chapter, PDF operators are shown with a list of the operands they require. For operators that correspond to one or more PostScript language operators, the corresponding PostScript language operator appears in bold on the first line of the operator’s definition. An operand specified as “number” may be either an integer or a real number. Otherwise, numeric operands must be integers.



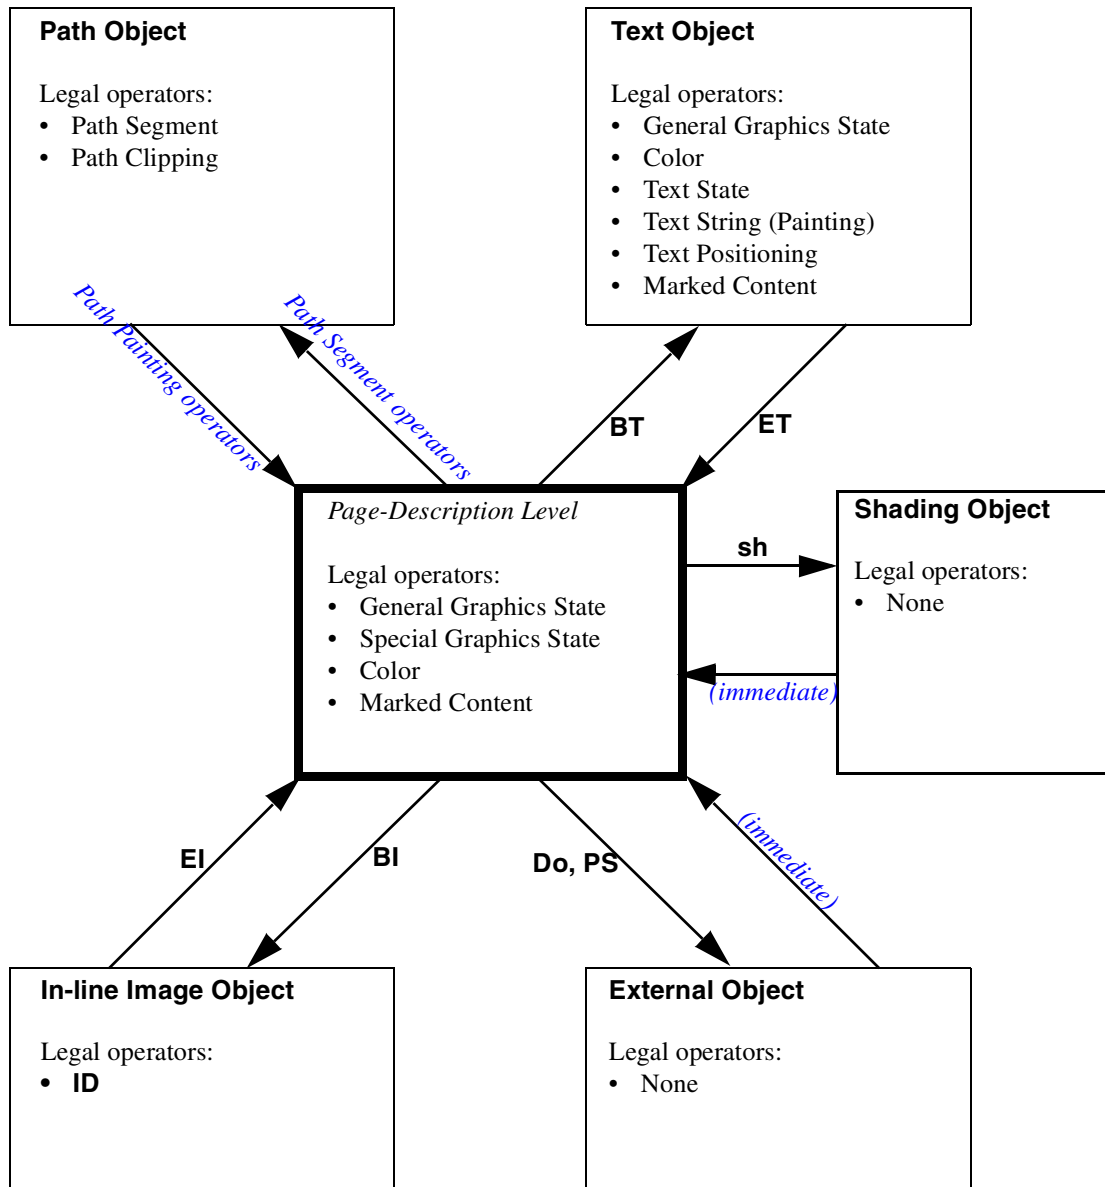
8.1 Overview

A PDF page description can be considered a sequence of graphics objects. These objects generate marks that are applied to the current page, obscuring any existing marks they may overlay.

PDF provides five types of graphics objects:

- A *path object* is an arbitrary shape made of straight lines, rectangles, and cubic curves. A path may intersect itself and may have disconnected sections and holes. A path object usually includes a painting operator that specifies whether the path is filled, stroked, and/or serves as a clipping path. A painting operator is not required, however; unpainted (“invisible”) path objects are sometimes used as placeholders or to denote text bounds.
- A *text object* consists of one or more character strings that can be placed anywhere on the page and in any orientation. Like a path, text can be stroked, filled, and/or serve as a clipping path.
- An *image object* consists of a set of samples using a specified color model. Images can be placed anywhere on a page and in any orientation.

Figure 8.1 Graphics Objects



General Graphics State operators: **d, gs, i, j, J, M, ri, w**
 Special Graphics State operators: **q, Q, cm**
 Color operators: **g, G, k, K, rg, RG, sc, SC, scn, SCN, cs, CS**
 Text State operators: **TC, Tf, TL, Tr, Ts, Tw, Tz**
 Text String (Painting) operators: **Tj, TJ, ', "**
 Text Positioning operators: **Td, TD, Tm, T***
 Path Segment operators: **c, h, l, m, re, v, y**
 Path Painting operators: **f, F, f*, n, s, S, b, b*, B, B***
 Path Clipping operators: **W, W***
 Marked Content operators: **BMC, BDC, EMC, MP, DP**
 Shading operator: **sh**

- An *External Object* (XObject) is an object defined outside the stream. The interpretation of an XObject depends on its type. PDF currently supports three types of XObjects: images, forms, and PostScript language fragments.
- *Ashading object* describes a smooth transition of colors across an area on the page.

As described in [Chapter 7](#), a PDF page description is not necessarily self-contained. It often contains references to resources such as fonts, patterns, forms, or images not found within the page description itself but located elsewhere in the PDF file.

[Figure 8.1](#) shows the ordering rules for the operations that define graphics objects. Some operations are permitted only in certain graphics objects or in the intervals between graphics objects, which is called the Page Description Level in the Figure. Every contents stream begins at the Page Description Level, where changes can be made to the graphics state, including colors and text-specific parameters, as explained in the following sections. The arrows indicate the operators that mark the beginning or end of each of the graphics objects. For example, any Path Segment operator such as **m** (moveto) marks the beginning of a Path object. Inside a Path object, additional Path Segment operators are permitted, as are Path Clipping operators, but not a General Graphics State operator such as **d** (setdash), for example. A Path Painting operator such as **f** (fill) marks the end of the Path Object and the return to the Page Description Level.

8.2 Graphics state

The exact effect of drawing a graphics object is determined by parameters such as the current line thickness, font, and leading. These parameters are part of the *graphics state*.

Although the contents of the PDF graphics state are similar to those of the graphics state in the PostScript language, there are several differences:

1. In PDF, the graphics state is divided into four distinct groups of parameters and operators. There are specific groups for text, for color, for “generic” marking operations, and for the graphics state itself. In this chapter, starting in [Figure 8.1](#), these are referred to as Text State, Color, General Graphics State, and Special Graphics State operators, respectively. The Text State, for example, enables the implementation of a more compact set of text operators.
2. The graphics state is extended to distinguish the parameters for fill operations from those for stroke operations. The use of separate fill and stroke colors in PDF is necessary to implement painting operators that both fill and stroke a path or text.
3. Finally, the graphics state in PDF 1.2 permits user extensions by means of the Marked Content operators. These have no effect on viewing or printing, but they preserve information that may be of use to plug-ins.

The graphics state is initialized at the beginning of each page, using the default values specified in each of the operator descriptions.

8.3 Special Graphics State

The Special Graphics State refers to parameters that apply to all four types of graphics objects: path, text, image, and external.

PDF provides a *graphics state stack* for saving and restoring the entire graphics state: the General Graphics State, the Color, and the Text State. PDF provides an operator, **q**, that saves a copy of the graphics state onto the graphics state stack. Another operator, **Q**, removes the most recently saved graphics state from the stack and makes it the current graphics state.

8.3.1 Special Graphics state parameters

8.3.1.1 Clipping path

The clipping path restricts the region to which paint can be applied on a page. Marks outside the region bounded by the clipping path are not painted. Clipping paths may be specified either by a path, or by using one of the clipping modes for text rendering. These are described in [Section 8.6.3, “Path clipping operators,”](#) and [Section 8.7.1.7, “Text rendering mode.”](#)

8.3.1.2 Current transformation matrix

The CTM is the matrix specifying the transformation from user space to device space. It is described in [Section 3.2, “User space.”](#)

8.3.1.3 Current point

All drawing on a page makes use of the *current point*. In an analogy to drawing on paper, the current point can be thought of as the location of the pen used for drawing.

The current point must be set before graphics can be drawn on a page. Several of the operators discussed in [Section 8.6.1, “Path segment operators,”](#) set the current point. As a path object is constructed, the current point is updated in the same way as a pen moves when drawing graphics on a piece of paper. After the path is painted using the operators described in [Section 8.6.2, “Path painting operators,”](#) the current point is undefined.

The current point also determines where text is drawn. Each time a text object begins, the current point is set to the origin of the page’s coordinate system. Several of the operators described in [Section 8.7.3, “Text positioning operators,”](#) change the current point. The current point is also updated as text is drawn using the operators described in [Section 8.7.5, “Text string operators”](#).

8.3.2 Special Graphics State operators

The operators in this section may be used only at the Page-Description Level; see [Figure 8.1](#). Adjacent to each PDF operator name is the PostScript language equivalent operator, if any.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
	q	Saves the current graphics state on the graphics state stack. (gsave)
	Q	Restores the graphics state to the most recently saved state. Removes the most recently saved state from the stack and makes it the current state. (grestore)
<i>a b c d e f</i>	cm	Modifies the CTM by concatenating the specified matrix. Although the operands specify a matrix, they are passed as six numbers, not as an array. (concat)


8.4 General Graphics state

8.4.1 Flatness

Flatness sets the maximum permitted distance in device pixels between the mathematically correct path and an approximation constructed from straight line segments, as shown in [Figure 8.2](#).

Note Flatness is inherently device-dependent, because it is measured in device pixels.

Figure 8.2 Flatness






<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>flatness</i>	i	Sets the flatness parameter in the graphics state. <i>flatness</i> is a number in the range 0 to 100, inclusive. The default value for <i>flatness</i> is 0, which means that the device's default flatness is used. (setflat)

8.4.2 Line cap style

The line cap style specifies the shape to be used at the ends of open subpaths when they are stroked. Allowed values are shown in [Figure 8.3](#).

Figure 8.3 *Line cap styles*

	Line cap style	Description
	0	Butt end caps—the stroke is squared off at the endpoint of the path.
	1	Round end caps—a semicircular arc with a diameter equal to the line width is drawn around the endpoint and filled in.
	2	Projecting square end caps—the stroke extends beyond the end of the line by a distance which is half the line width and is squared off.







<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
------------------	-----------------	------------------

<i>linecap</i>	J	Sets the line cap parameter in the graphics state. <i>linecap</i> has a default value of 0.
----------------	----------	---

8.4.3 Line dash pattern

The line dash pattern controls the pattern of dashes and gaps used to stroke paths. It is specified by an *array* and a *phase*. The array specifies the length of alternating dashes and gaps. The phase specifies the distance into the dash pattern to start the dash. Both the elements of the array and the phase are measured in user space units. Before beginning to stroke a path, the array is cycled through, adding up the lengths of dashes and gaps. When the sum of dashes and gaps equals the value specified by the phase, stroking of the path begins, using the array from the point that has been reached. [Figure 8.4](#) shows examples of line dash patterns. As can be seen from the figure, the command `[] 0 d` can be used to restore the dash pattern to a solid line. (**setlinecap**)

Figure 8.4 *Line dash pattern*

Dash pattern	Array and phase	Description
	[] 0	Turn dash off—solid line
	[3] 0	3 units on, 3 units off, ...
	[2] 1	1 on, 2 off, 2 on, 2 off, ...
	[2 1] 0	2 on, 1 off, 2 on, 1 off, ...
	[3 5] 6	2 off, 3 on, 5 off, 3 on, 5 off, ...
	[2 3] 11	1 on, 3 off, 2 on, 3 off, 2 on, ...

Dashed lines wrap around curves and corners just as solid stroked lines do. The ends of each dash are treated with the current line cap style, and corners within dashes are treated with the current line join style.

Arguments	Operator	Semantics
<i>[array] phase</i>	d	Sets the dash pattern parameter in the graphics state. If <i>array</i> is empty, the dash pattern is a solid, unbroken line; otherwise <i>array</i> is an array of numbers, all non-negative and at least one non-zero, that specify alternating distances in user space for the length of dashes and gaps. <i>phase</i> is a number that specifies a distance in user space into the dash pattern at which to begin marking the path. The default dash pattern is a solid line. (setdash)




The **S** (stroke) operator uses the contents of the array of dashes and gaps in a cyclical fashion; when it reaches the end of the array, it starts again at the beginning.

When a path consisting of several subpaths is stroked, each subpath is treated independently—in other words, the dash pattern is restarted and *phase* is applied to it at the beginning of each subpath.

8.4.4 Line join style

The line join style specifies the shape to be used at the corners of paths that are stroked. [Figure 8.5](#) shows the allowed values.

Figure 8.5 *Line join styles*

	Line join style	Description
	0	Miter joins—the outer edges of the strokes for the two segments are continued until they meet. If the extension projects too far, as determined by the miter limit, a bevel join is used instead.
	1	Round joins—a circular arc with a diameter equal to the line width is drawn around the point where the segments meet and filled in, producing a rounded corner.
	2	Bevel joins—the two path segments are drawn with butt end caps (see the discussion of line cap style), and the resulting notch beyond the ends of the segments is filled in with a triangle.

Arguments	Operator	Semantics
<i>linejoin</i>	j	Sets the line join parameter in the graphics state. <i>linejoin</i> has a default value of 0. (setlinejoin)

8.4.5 Line width

The line width specifies the thickness of the line used to stroke a path and is measured in user space units. A line width of 0 specifies the thinnest line that can be rendered on the output device.

Note A line width of 0 is an inherently device-dependent value. Its use is discouraged because the line may be nearly invisible when printing on high-resolution devices.

Arguments	Operator	Semantics
<i>linewidth</i>	w	Sets the line width parameter in the graphics state. <i>linewidth</i> is a number and has a default value of 1. (setlinewidth)

8.4.6 Miter limit


When two line segments meet at a sharp angle and mitered joins have been specified as the line join style, it is possible for the miter to extend far beyond the thickness of the line stroking the path. The miter limit imposes a maximum on the ratio of the miter length to the line width, as shown in [Figure 8.6](#). When the limit is exceeded, the join is converted from a miter to a bevel.

The ratio of miter length to line width is directly related to the angle ϕ between the segments in user space by the formula:

$$\frac{\text{miter length}}{\text{line width}} = \frac{1}{\sin\left(\frac{\phi}{2}\right)}$$

For example, a miter limit of 1.415 converts miters to bevels for ϕ less than 90 degrees, a limit of 2.0 converts miters to bevels for ϕ less than 60 degrees, and a limit of 10.0 converts miters to bevels for ϕ less than 11 degrees.

Figure 8.6 *Miter length*



Arguments	Operator	Semantics
<i>miterlimit</i>	M	Sets the miter limit parameter in the graphics state. <i>miterlimit</i> is a number that must be greater than or equal to 1, and has a default value of 10. (setmiterlimit)

8.4.7 Generic Graphics State operator

All the remaining parameters in the General Graphics State are set with the **gs** operator, whose operand is an “extended graphics state” dictionary. (See [page 272](#).) Each parameter uses a different keyword in this dictionary.

Note It is expected that any future extensions to the graphics state will also use the **gs** operator, with new keywords, rather than new operators.

Arguments	Operator	Semantics
<i>name</i>	gs	Sets the specified device-dependent parameters in the graphics state: stroke adjustment, overprint, black generation, undercolor removal, transfer function, halftone, and halftone phase. Parameters that are not specified are not changed. <i>name</i> is the name of an ExtGState dictionary in the current Resources dictionary.

PDF 1.2

In PDF 1.3, *all* the parameters of the General Graphics State, which can be set by operators described in Sections 8.4.1 through 8.4.6, have equivalent keys in the ExtGState dictionary. This is intended for use by Type 2 patterns (smooth shading), which do not have a contents stream.

PDF 1.3

8.4.8 Stroke adjustment

PDF 1.2

The stroke adjustment parameter controls whether the line width and the coordinates of a stroke are automatically adjusted as necessary to produce lines of uniform thickness. For details, see section 7.5.2, “Automatic Stroke Adjustment,” of the *PostScript Language Reference Manual, Third Edition* [1]. The keyword for stroke adjustment is **SA**. The default value is *true*.

8.4.9 Overprint

PDF 1.2

The overprint parameter is used only when producing separations. It specifies whether painting on one separation causes the corresponding areas of other separations to be erased (*false*) or left unchanged (*true*). See section 4.8.4, “Special Color Spaces,” of the *PostScript Language Reference Manual, Second Edition* [1].

Separate keywords are used for controlling overprint for fill and for stroke. The keyword **OP** specifies overprint for stroke; the keyword **op** specifies overprint for fill. In an ExtGState dictionary, if **OP** is specified but **op** is omitted, then the **OP** key applies to both fill and stroke. The default value is *false* for both **OP** and **op**.

PDF 1.3

The dictionary may also specify an *overprint mode* with the **OPM** key. The overprint mode affects the interpretation of the **k** and **K** operators, which set the color in a 4-component CMYK color space (**k** for fill, **K** for stroke). If the overprint mode is 0 (the default), then a color component of 0 erases (“paints white”) the current path on the corresponding separation. If the overprint mode is 1, then a color component of 0 leaves the corresponding separation unchanged. For example, if overprint mode is 1, then the operation **.2 .3 0 1 k** leaves the third (yellow) separation unchanged. It has a similar effect to **.2 .3 1 sc** in a **[/DeviceN [/Cyan /Magenta /Black] ...]** color space.

PDF 1.3

8.4.10 Black generation

PDF 1.2

The black-generation function computes the value of the black component during conversion from **DeviceRGB** color space to **DeviceCMYK**. For additional information, see section 6.2.3, “Conversion from **DeviceRGB** to **DeviceCMYK**,” of the *PostScript Language Reference Manual, Second Edition* [1]. The keyword for black generation is **BG**.

8.4.11 Undercolor removal

PDF 1.2

The undercolor removal function computes the amount to subtract from the cyan, magenta, and yellow components during conversion of color values from **DeviceRGB** color space to **DeviceCMYK**. See section 7.2.3, “Conversion from **DeviceRGB** to **DeviceCMYK**,” of the *PostScript Language Reference Manual, Third Edition* [1]. The keyword for undercolor removal is **UCR**.

8.4.12 Transfer function

PDF 1.2

The transfer function adjusts the values of the gray color component. It is also a part of some halftone screens. For complete details, see section 6.3, “Transfer Functions,” of the *PostScript Language Reference Manual, Second Edition* [1]. The keyword for transfer function is **TR**.

8.4.13 Halftone

PDF 1.2

The halftone parameter of the graphics state specifies how halftones should be produced. See 7.15, “Extended graphics states,” for details about halftones. For general information on halftones, see section 6.4.3, “Halftone Dictionaries,” of the *PostScript Language Reference Manual, Second Edition* [1]. The keyword for halftone is **HT**.

8.4.14 Halftone phase

PDF 1.2

The halftone phase parameters of the graphics state specifies the phase relationship of halftone cells to the coordinate axes. See section 7.3.3, “Halftone Phase,” of the *PostScript Language Reference Manual, Second Edition* [1]. The keyword for halftone phase is **HTP**.

8.4.15 Smoothness

PDF 1.3

This parameter controls the quality of smooth shading (Type 2 patterns and the **sh** operator), and thus indirectly controls the rendering performance. Smoothness is the allowable color error between shading approximated with piecewise linear interpolation and the true shading of a possibly nonlinear shading function. The error is measured for each color component, and the maximum error is used. The allowable error (or tolerance) is specified as a percentage of the range of the color component. This percentage is expressed as a value from 0 to 1. Thus, a smoothness parameter of 0.1 represents a tolerance of 10% in each color component.

Each device may have internal limits on the maximum and minimum tolerances attainable. For example, setting smoothness to 1 may result in an internal smoothness of 0.5 on a high-quality color device, and setting smoothness to 0 on the same device may result in an internal smoothness of 0.01 if an error of that magnitude is imperceptible on that device.

The smoothness parameter may also interact with the accuracy of color conversion. In the case of a color conversion defined by a sampled function, the conversion function is unknown. Thus, the error may be sampled at too low a

frequency, in which case, the accuracy defined by the smoothness parameter cannot be guaranteed. In most cases, however, where the conversion function is smooth and continuous, the accuracy should be within the specified tolerance.

The effect of the smoothness parameter is similar to that of the flatness parameter. Note, however, that flatness is measured in device-dependent units of pixel width, whereas smoothness is measured as a percentage of color component range.

The keyword for smoothness is **SM**.

8.5 Color

8.5.1 Color parameters

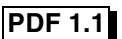
8.5.1.1 Fill color

The fill color is used to paint the interior of paths and text characters that are filled. Filling is described in [Section 8.6.2, “Path painting operators.”](#)

8.5.1.2 Stroke color

The stroke color is used to paint the border of paths and text that are stroked. Stroking is described in [Section 8.6.2, “Path painting operators.”](#)

8.5.1.3 Fill color space



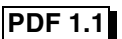
The fill color space is the color space in which the fill color is specified.

8.5.1.4 Stroke color space



The stroke color space is the color space in which the stroke color is specified.

8.5.1.5 Rendering intent

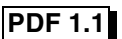


The rendering intent is a name that indicates the style of color rendering that should occur. See [Section 7.13, “XObjects,”](#) and especially [Table 7.36, “Color rendering intents,”](#) for further detail.

8.5.2 Color operators



The operators that set colors and color spaces fall into two classes. Operators in the first class, which were defined in PDF 1.0, set the color and color space at the same time, and they include only device-dependent color spaces. Operators in the second class, which are defined in PDF 1.1, set colors and color spaces separately, and they apply to all color spaces.



The default color space is **DeviceGray**, and the default fill and stroke colors are both black.

Color operators may appear inside Text Objects or at the Page-Description Level. See [Figure 8.1](#).

8.5.2.1 Device-dependent color space operators

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>gray</i>	g	Sets the color space to DeviceGray (or the DefaultGray color space, see Section 7.12.12 on page 245), and sets the gray tint to use for filling paths. <i>gray</i> is a number between 0 (black) and 1 (white). (setgray (fill))
<i>gray</i>	G	Sets the color space to DeviceGray (or the DefaultGray color space, see Section 7.12.12 on page 245), and sets the gray tint to use for stroking paths. <i>gray</i> is a number between 0 (black) and 1 (white). (setgray (stroke))
<i>r g b</i>	rg	Sets the color space to DeviceRGB (or the DefaultRGB color space, see Section 7.12.12 on page 245), and sets the color to use for filling paths. Each operand must be a number between 0 (minimum intensity) and 1 (maximum intensity). (setrgbcolor (fill))
<i>r g b</i>	RG	Sets the color space to DeviceRGB (or the DefaultCMYK color space, see Section 7.12.12 on page 245), and sets the color to use for stroking paths. Each operand must be a number between 0 (minimum intensity) and 1 (maximum intensity). (setrgbcolor (stroke))
<i>c m y k</i>	k	Sets the color space to DeviceCMYK (or the DefaultCMYK color space, see Section 7.12.12 on page 245), and sets the color to use for filling paths. Each operand must be a number between 0 (no ink) and 1 (maximum ink). The behavior of the k operator is affected by the overprint mode; see Section 8.4.9 on page 328 . (setcmykcolor (fill))
<i>c m y k</i>	K	Sets the color space to DeviceCMYK (or the DefaultRGB color space, see Section 7.12.12 on page 245), and sets the color to use for stroking paths. Each operand must be a number between 0 (no ink) and 1 (maximum ink). The behavior of the K operator is affected by the overprint mode; see Section 8.4.9 on page 328 . (setcmykcolor (stroke))

8.5.2.2 Generic color space operators

PDF 1.1

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>colorspace</i>	cs	Sets the color space to use for filling paths. <i>colorspace</i> must be a name. If the color space is specified by a name (the device-dependent color spaces DeviceGray , DeviceRGB , and DeviceCMYK ; or the Pattern color space for colored tiling patterns or shading patterns), then that name may be used. If it is specified by an array (all other color spaces), then <i>colorspace</i> must be a name defined in the current Resources dictionary. (setcolorspace (fill))

For example, the following expression is illegal:

```
[/CalGray dict] cs
```

Instead, one would write

```
/CS42 cs
```

and the Resources dictionary would contain

```
/CS42 [/CalGray dict]
```


The **cs** operator also sets the current fill-color to its initial value, which depends on the color space. For the device-dependent, calibrated, and **ICCBased** color spaces, the initial color is black. For a **Lab** color space, the initial value is specified by the minimum **Range** values. For an **Indexed** color space, the initial value is 0. The initial value in a **Separation** color space is 1, and the initial color value in a **Pattern** color space is a pattern that has an empty stream of marking operators, thus producing no marks on the page.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>colorspace</i>	CS	Same as cs , but for strokes. (setcolorspace (stroke))
$c_1 c_2 c_3 c_4$	sc	Sets the color to use for filling paths. The number of operands required and their interpretation is based on the current fill color space. For DeviceGray , CalGray , and Indexed color spaces, one operand is required. For DeviceRGB , CalRGB , and Lab color spaces, three operands are required. For DeviceCMYK and CalCMYK , four operands are required. (setcolor (fill))
$c_1 c_2 c_3 c_4$	SC	Same as sc , but for stroking paths. (setcolor (stroke))
$c_1 \dots c_n$ $c_1 \dots c_n name$	scn scn	scn accepts the same parameters, and has the same effect, as sc . In addition, it supports Pattern , Separation , ICCBased , and DeviceN colors. (setcolor (fill for patterns))

If the current fill color space is a **Pattern** color space, then **scn** sets the pattern to use for filling paths. *name* is the name of a **Pattern** resource in the current Resources dictionary. If the pattern is uncolored (if **PatternType** is 1 and **PaintType** is 2), then the color is determined by the component values $c_1 \dots c_n$ in the underlying color space. If the pattern is colored (if **PatternType** is 1 and **PaintType** is 1), or if it is a shading pattern (if **PatternType** is 2), then the component values must not be specified.

If the current fill color space is a **Separation** color space, then **scn** sets the tint for filling paths to c_1 , which is a number in the range 0 to 1 that represents the amount of colorant to be applied.

If the current fill color space is an **ICCBased** color space, then **scn** sets the color values for filling paths to $c_1 \dots c_n$, which are numbers in the range 0 to 1.



If the current fill color space is a **DeviceN** color space, then **scn** sets the tints for filling paths to $c_1 \dots c_n$, which are numbers in the range 0 to 1 that represents the amount of each colorant to be applied. (**setcolor** (fill))

PDF 1.3

$c_1 \dots c_n$	SCN	
$c_1 \dots c_n$ <i>name</i>	SCN	Same as scn , but for strokes. (setcolor (stroke))

PDF 1.2

8.5.2.3 Color rendering intent

PDF 1.1


<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>intent</i>	ri	Sets the color rendering intent in the graphics state.

intent is a name of a color rendering intent, which indicates the style of color rendering that should occur, as described in [Table 7.36 on page 252](#).

8.6 Paths

Paths are used to represent lines, curves, and regions. A path consists of a series of path segment operators describing where marks are to appear on the page, followed by a path painting operator, which actually marks the path in one of several ways. A path may be composed of one or more disconnected sections, referred to as *subpaths*. An example of a path with two subpaths is a path containing two parallel line segments.

Path segments may be straight lines or curves. Curves in PDF files are represented as cubic Bézier curves. A cubic Bézier curve is specified by the *x*- and *y*-coordinates of four points: the two endpoints of the curve (the current point, P_0 , and the final point, P_3) and two *control points* (points P_1 and P_2), as shown in [Figure 8.7](#).

Figure 8.7 *Bézier curve*

Once these four points are specified, the cubic Bézier curve $R(t)$ is generated by varying the parameter t from 0 to 1 in the following equation:

$$R(t) = (1-t)^3 P_0 + 3t(1-t)^2 P_1 + 3t^2(1-t) P_2 + t^3 P_3$$

In this equation, P_0 is the current point before the curve is drawn. When the parameter t has the value 0, $R(t) = P_0$ (the current point). When $t = 1$, $R(t) = P_3$. The curve does not, in general, pass through the two control points P_1 and P_2 .

Bézier curves have two desirable properties. First, the curve is contained within the convex hull of the control points. The convex hull is most easily visualized as the polygon obtained by stretching a rubber band around the outside of the four points defining the curve. This property allows rapid testing of whether the curve is completely outside the visible region, and so does not have to be rendered. Second, Bézier curves can be very quickly split into smaller pieces for rapid rendering.

Note In the remainder of this book, the term Bézier curve means cubic Bézier curve.

Paths are subject to and may also be used for clipping. Path clipping operators replace the current clipping path with the intersection of the current clipping path and the current path.

```
<path> ::= <subpath>+
          [path clipping operator]
          <path painting operator>
```


```
<subpath> ::= m <path segment operator except m and re>* |
            re
```

8.6.1 Path segment operators


All operands are numbers that are coordinates in user space.

Arguments	Operator	Semantics
$x\ y$	m	Moves the current point to (x, y) , omitting any connecting line segment. (moveto)
$x\ y$	l	(operator is lowercase L) Appends a straight line segment from the current point to (x, y) . The new current point is (x, y) . (lineto)
$x_1\ y_1\ x_2\ y_2\ x_3\ y_3$	c	Appends a Bézier curve to the path. The curve extends from the current point to (x_3, y_3) using (x_1, y_1) and (x_2, y_2) as the Bézier control points, as shown in Figure 8.7 . The new current point is (x_3, y_3) . (curveto)
$x_2\ y_2\ x_3\ y_3$	v	Appends a Bézier curve to the current path between the current point and the point (x_3, y_3) using the current point and (x_2, y_2) as the Bézier control points, as shown in Figure 8.8 . The new current point is (x_3, y_3) . (curveto (first control point coincides with initial point on curve))

Figure 8.8 *v* operator



$x_1\ y_1\ x_3\ y_3$	y	Appends a Bézier curve to the current path between the current point and the point (x_3, y_3) using (x_1, y_1) and (x_3, y_3) as the Bézier control points, as shown in Figure 8.9 . The new current point is (x_3, y_3) . (curveto (second control point coincides with final point on curve))
----------------------	----------	---

Figure 8.9 *y* operator*x y width height***re** Adds a rectangle to the current path.*width* and *height* are distances in user space. The operation*x y width height re*

is defined to have the same effect as the sequence

```

x y m
x+width y l
x+width y+height l
x y+height l
h

```

h Closes the current subpath by appending a straight line segment from the current point to the starting point of the subpath. (**closepath**)

8.6.2 Path painting operators

Paths may be stroked and/or filled. As in the PostScript language, painting completely obscures any marks already on the page under the region that is painted.


Stroking draws a line along the path, using the line width, dash pattern, miter limit, line cap style, line join style, stroke color, stroke color space, and stroke adjustment from the graphics state. The line drawn when a path is stroked is centered on the path. If a path consists of multiple subpaths, each is treated separately.

The process of filling a path paints the entire region enclosed by the path, using the fill color and fill color space. If a path consists of several disconnected subpaths, each is filled separately. Any open subpaths are implicitly closed before being filled. Closing is accomplished by adding a segment between the first and last

points on the path. For a simple path, it is clear what lies inside the path and should be painted by a fill. For more complicated paths, it is not so obvious. One of two rules is used to determine which points lie inside a path.


The *non-zero winding number rule* uses the following test to determine whether a given point is inside a path and should be painted. Conceptually, a ray is drawn in any direction from the point in question to infinity, and the points where the ray crosses path segments are examined. Starting from a count of zero, add one to the count each time a path segment crosses the ray from left to right, and subtract one from the count each time a path segment crosses the ray from right to left. If the ray encounters a path segment that coincides with it, the result is undefined. In this case, a ray in another direction can be picked, since all rays are equivalent. After counting all the crossings, if the result is zero then the point is outside the path. The effect of using this rule on various paths is illustrated in [Figure 8.10](#). The non-zero winding number rule is used by the PostScript language **fill** operator.

Figure 8.10 *Non-zero winding number rule*



The *even-odd rule* uses a slightly different strategy. The same calculation is made as for the non-zero winding number rule, but instead of testing for a result of zero, a test is made as to whether the result is even or odd. If the result is odd, the point is inside the path; if the result is even, the point is outside. The result of applying this rule to various paths is illustrated in [Figure 8.11](#). The even-odd rule is used by the PostScript language **eofill** operator.

Figure 8.11 *Even-odd rule*



<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
	n	Ends the path without filling or stroking it. This is a “path painting no-op,” primarily used with a path clipping operator (see Section 8.6.3, “Path clipping operators), but like the other path painting operators, it terminates a Path Object. (newpath)
	S	Strokes the path. (stroke)
	s	Similar to the S operator, but closes the path before stroking it. (closepath and stroke) s has the same effect as h S .
	f	Fills the path, using the non-zero winding number rule to determine the region to fill. (fill)
	F	Same as the f operator. Included only for compatibility. Although applications that read PDF files must be able to accept this operator, applications that generate PDF files should use the f operator instead. (fill)
	f*	Fills the path, using the even–odd rule to determine the region to fill. (eofill)
	B	fill and stroke . B has the same effect as q f Q S .
	b	closepath , fill , and stroke . b has the same effect as h B .
	B*	eofill and stroke . B* has the same effect as q f* Q S .
	b*	closepath , eofill , and stroke . b* has the same effect as h B* .
<i>name</i>	sh	When a path that is to be filled with a gradient (see Section 7.17 on page 287) has the same geometry as the gradient itself, it is not necessary to define the gradient as a Type 2 pattern, define the path separately, and then use the f (fill) operator. Instead, one may use the sh operator to paint the same area. (shfill)
		<i>name</i> is the name of a Shading dictionary in the current Resources dictionary. (If its ShadingType is greater than 3, this dictionary is part of a stream.) All coordinates in the Shading dictionary are interpreted relative to the current user space. (When a Shading dictionary is used in a pattern, the coordinates are expressed in pattern space.)

PDF 1.3

8.6.3 Path clipping operators

Path clipping operators cause the current clipping path to be replaced with the intersection of the current clipping path and the path. A path is made into a clipping path by inserting a path clipping operator (**W** or **W***) between the last path segment operator and the path painting operator.

Although the path clipping operator appears before the path painting operator, the path clipping operator does not alter the clipping path at the point it appears. Rather, it modifies the effect of the path painting operator. After the path is filled,

stroked, or ended by the path painting operator, it is set to be the current clipping path. If the path is both filled and stroked, the painting is done in that order before making the path the current clipping path.

The definition of the clipping path and all subsequent operations it is to affect should be contained between a pair of **q** and **Q** operators. Execution of the **Q** operator causes the clipping path to revert to that saved by the **q** operator, before the clipping path was modified.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
	W	Uses the non-zero winding number rule to determine which regions are inside the clipping path. (clip)
	W*	Uses the even-odd rule to determine which regions are inside the clipping path. (eoclip)

8.7 Text state

The text state is composed of those graphics state parameters that affect only text.

8.7.1 Text State parameters and operators

There are nine parameters in the text state, each of which can be set individually:

1. T_c is the character spacing parameter.
2. T_w is the word spacing parameter.
3. T_h is the horizontal spacing parameter.
4. T_l is the “leading” parameter.
5. T_f is the text font.
6. T_{fs} is the text font size.
7. T_m is the text matrix.
8. T_{mode} is the rendering mode.
9. T_{rise} is the “text rise”.

There are two additional parameters of the text state:

1. T_{LM} is the matrix for the current text line.
2. T_{RM} is the rendering matrix.

Each of the items in the text state is described in the following sections.

Note [Section 8.7.4, “Text rendering,”](#) describes how these parameters are used, and their exact effects on the text state.

Note These operators can appear outside of text objects, and the values they set are retained across text objects on a single page. Like other graphics state parameters, the values are initialized to the default values at the beginning of each page.

8.7.1.1 Character spacing

The character spacing parameter, T_c , is a number specified in text space units. It is added to the displacement between the origin of one character and the origin of the next. See [Figure 7.3 on page 209](#) for examples of character origins and displacements. In the default coordinate system, the positive direction of the x -axis points to the right, and the positive direction of the y -axis points upward. So for horizontal writing, a positive value of T_c has the effect of expanding the space between characters; see [Figure 8.12](#). For vertical writing, however, a *negative* value of T_c has the effect of expanding the space between characters.

Figure 8.12 *Character spacing for horizontal writing*

Character	0 (default)
C h a r a c t e r	0.25

Character spacing is applied to each glyph in the string, regardless of the number of bytes used for that glyph’s character code. Therefore character spacing is used even with fonts that have multi-byte encodings.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>charSpace</i>	Tc	Set character spacing
Sets T_c to <i>charSpace</i> . Character spacing is used, together with word spacing, by the Tj , TJ , and ' operators. <i>charSpace</i> is a number expressed in text space units and has an initial value of 0.		

8.7.1.2 Word spacing

The word spacing parameter, T_w , is a number specified in text space units. It works in the same way as character spacing, but applies only to the space character, <20>. T_w is added to the displacement between the origin of the space character and the origin of the following character. For horizontal writing, a positive value for T_w has the effect of increasing the spacing between words. For vertical writing, a positive

value for T_w *decreases* the space between words, since the positive direction of the y -axis points upward; therefore a negative value will increase the space between words. [Figure 8.13](#) illustrates the effect of word spacing in horizontal writing.

Figure 8.13 *Effect of word spacing in horizontal writing*

Word Space	0 (default)
Word Space	2.5

Word spacing is applied to every instance of the single byte <20> in a string. Therefore word spacing is not used with fonts that have only multi-byte encodings or with fonts whose encodings do not use the single byte <20> as the space character.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<code>wordSpace</code>	Tw	Set word spacing

Sets T_w to `wordSpace`. Word spacing is used by the **Tj**, **TJ**, and ' operators. `wordSpace` is a number expressed in text space units and has an initial value of 0.

8.7.1.3 Horizontal scaling

The horizontal scaling parameter, T_h , adjusts the width of characters by stretching or shrinking them in the horizontal direction. The scaling is specified as a percent of the normal width of the characters, with 100 being the normal width. [Figure 8.14](#) shows the effect of horizontal scaling. The scaling always applies to the x coordinate, independent of the writing mode.

Figure 8.14 *Horizontal scaling*

Word	100 (default)
WordWord	50


<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<code>scale</code>	Tz	Set horizontal scaling

Sets T_h to $(scale \div 100)$. $scale$ is a number expressed in percent of the normal scaling and has an initial value of 100.

8.7.1.4 Leading

The leading parameter, T_l is measured in text space units. It specifies the vertical distance between the baselines of adjacent lines of text, as shown in [Figure 8.15](#). The leading parameter is used by the **TD**, **T***, **'**, and **"** operators; it is independent of the writing mode.

Figure 8.15 *Leading*



<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>leading</i>	TL	Set text leading

Sets T_l to *leading*. The **TL** operator need not be used in a PDF file unless the **T***, **'**, or **"** operators are used. *leading* has an initial value of 0.

8.7.1.5 Text font and size

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>fontname size</i>	Tf	Set font and size

Sets T_f to *fontname* and T_{fs} to *size*. There is no initial value for either *fontname* or *size*; they must be specified using **Tf** before drawing any text. *fontname* is the name of a Font in the current Resources dictionary. *size* is a number expressed in text space units.

8.7.1.6 Text matrix

The text matrix specifies the transformation from text space (see [Section 3.3](#), “Text space”) to user space. The text matrix is set with the **Tm** operator (see [page 345](#)).

8.7.1.7 Text rendering mode








Determines whether text is stroked, filled, or used as a clipping path.

Note The text rendering mode has no effect on text displayed using a Type 3 font.

The rendering modes are shown in [Figure 8.16](#). In the figure, a stroke color of black and a fill color of light gray are used. After one of the clipping modes is used for text rendering, the text object must be ended using the **ET** operator before changing the text rendering mode.

Note For the clipping modes (4–7), a series of lines has been drawn through the characters in [Figure 8.16](#) to show where the clipping occurs.

Figure 8.16 Text rendering modes

	Rendering mode	Description
	0	Fill text
	1	Stroke text
	2	Fill then stroke text
	3	Text with no fill and no stroke (invisible)
	4	Fill text and add it to the clipping path
	5	Stroke text and add it to the clipping path
	6	Fill then stroke text and add it to the clipping path
	7	Add text to the clipping path

Arguments	Operator	Semantics
-----------	----------	-----------

<i>render</i>	Tr	Set the text rendering mode.
---------------	-----------	------------------------------

Sets T_{mode} to *render*, which is an integer and has an initial value of 0.

8.7.1.8 Text rise

Text rise specifies the amount, in text space units, to move the baseline up or down from its default location. Positive values of text rise move the baseline up. Adjustments to the baseline are useful for drawing superscripts or subscripts. The default location of the baseline can be restored by setting the text rise to 0. [Figure 8.17](#) illustrates the effect of the text rise, which is set using the **Ts** operator. Text rise always applies to the y coordinate, regardless of the writing mode.

Figure 8.17 *Text rise*

This text is ^{superscripted}	(This text is) Tj 5 Ts (superscripted) Tj
This text is _{subscripted}	(This text is) Tj -5 Ts (subscripted) Tj
This text ^{moves} around	(This) Tj -5 Ts (text) Tj 5 Ts (moves) Tj 0 Ts (around) Tj

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
------------------	-----------------	------------------

<i>rise</i>	Ts	Set text rise.
-------------	-----------	----------------

Sets T_{rise} to *rise*, which is a number expressed in text space units and has an initial value of 0.

8.7.2 Text Object operators

A PDF text object consists of operators that specify character strings, movement of the current point, and text state. A text object begins with the **BT** operator and ends with the **ET** operator. See [Figure 8.1 on page 320](#).

```
<text object> ::= BT
                    <text operator or graphics state operator>*
                    ET
```

When **BT** is encountered, the text matrix is initialized to the identity matrix. When **ET** is encountered, the text matrix is discarded. Text objects cannot be nested—a second **BT** cannot appear before an **ET**.

Note If a page does not contain any text, no text operators (including operators that merely set the text state) may be present in the page description.

Arguments	Operator	Semantics
	BT	Begins a Text Object. Initializes the text matrix, T_m , and the line matrix, T_{LM} , to the identity matrix.
	ET	Ends a Text Object. Discards the text matrix.

8.7.3 Text positioning operators

A text object keeps track of the current point and the start of the current line. The text string operators move the current point as the various forms of the PostScript language **show** operator do. Operators that move the start of the current line move the current point as well.

Note These operators may appear only within text objects. See [Figure 8.1 on page 320](#).

Arguments	Operator	Semantics
$t_x t_y$	Td	Moves to the start of the next line, offset from the start of the current line by (t_x, t_y) . t_x and t_y are numbers expressed in text space units. More precisely, Td performs the following assignments:
		$T_m = T_{LM} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1 \end{bmatrix} \times T_{LM}$
$t_x t_y$	TD	Moves to the start of the next line, offset from the start of the current line by (t_x, t_y) . As a side effect, this sets the leading parameter in the text state.
		$t_x t_y \mathbf{TD}$ is defined to have the same effect as $-t_y \mathbf{TL} t_x t_y \mathbf{Td}$
$a b c d x y$	Tm	Sets the text matrix, T_m , and the text line matrix, T_{LM} . It also sets the current point and line start position to the origin. Tm performs the following assignments:

$$T_m = T_{LM} = \begin{bmatrix} a & b & 0 \\ c & d & 0 \\ x & y & 1 \end{bmatrix}$$

The operands are all numbers, and the initial value for T_m and T_{LM} is the identity matrix, $\begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}$. Although the operands specify a matrix, they are passed to **Tm** as six numbers, not as an array.

The matrix specified by the operands passed to the **Tm** operator is not concatenated onto the current text matrix, but replaces it.

T* Moves to the start of the next line.

T* is defined to have the same effect as **0 T_l Td**

where T_l is the leading parameter of the text state.

8.7.4 Text rendering

Before text is rendered by the **Tj** or **TJ** operator, it is placed and transformed according to the parameters in the text state. The *rendering matrix* for the text is computed as follows:

The current text matrix, T_m , is translated by the text rise, T_{rise} . Next, that is scaled by the font size, T_{fs} , and the horizontal text scale, T_h . Finally, that is concatenated to the current transformation matrix in the graphics state (*CTM*) to produce the rendering matrix, T_{RM} :

$$T_{RM} = \begin{bmatrix} T_{fs} \times T_h & 0 & 0 \\ 0 & T_{fs} & 0 \\ 0 & T_{rise} & 1 \end{bmatrix} \times T_m \times CTM$$

This calculation occurs, in effect, whenever any of the text parameters change, before **Tj** or **TJ** occur. When text is rendered, the text line matrix, T_{LM} , is unaffected, but the text matrix, T_m , is translated by the origin-displacement of the text, which affects subsequent rendering operations, as shown above. For horizontal-mode writing, the origin-displacement is along the x axis; for vertical writing (see [Section 7.7.8 on page 208](#)), the displacement is along the y axis.

8.7.5 Text string operators

These operators draw text on the page. Although it is possible to pass individual characters to the text string operators, text searching performs significantly better if the text is grouped by word and paragraph.

PDF supports the same conventions as the PostScript language for specifying non-printable ASCII characters. That is, a character can be represented by an escape sequence, as described in [Table 4.1 on page 38](#).

Note The default current point is at the page origin. Therefore, unless some prior operation in the same text object changes the current point, the text will appear at the origin. It is suggested that a **Tm** operation be used to establish the initial current point in a text object at the position in text space where initial text is to appear. Subsequent text operations may change the current point.

Arguments	Operator	Semantics
<i>string</i>	Tj	Shows text string, using the character and word spacing parameters from the text state. (show)
<i>string</i>	'	Moves to next line and shows text string, using the character and word spacing parameters from the text state. (show) <i>string</i> ' is defined to have the same effect as T* string Tj
<i>a_w a_c string</i>	"	Moves to next line and shows text string. <i>a_w</i> and <i>a_c</i> are numbers expressed in text space units. <i>a_w</i> specifies the additional space width, and <i>a_c</i> specifies the additional space between characters. (show) <i>a_w a_c string</i> " is defined to have the same effect as a_w Tw a_c Tc string '
	Note	The values specified by <i>a_w</i> and <i>a_c</i> remain the word and character spacings after the " operator is executed.
<i>[number or string ...]</i>	TJ	Shows text string, allowing individual character positioning, and using the character and word spacing parameters from the text state. (show with displacements)



For each element of the array, if the element is a string, **TJ** shows the string. If it is a number, it is expressed in thousandths of an em. (An *em* is a typographic unit of measurement equal to the size of a font. For example, in a 12-point font, an em is 12 points.) **TJ** *subtracts* this amount from the current *x* coordinate in horizontal writing mode, or from the current *y* coordinate in vertical writing mode. In the normal case of horizontal writing in the default coordinate system, this has the effect of moving the current point to the *left* by the given amount.

Each character is first justified according to any character and word spacing settings made with the **Tc**, **Tw**, or " operators, and then any numeric offset present in the array passed to the **TJ** operator is applied. An example of the use of **TJ** is shown in [Figure 8.18](#).

Figure 8.18 Operation of **TJ** operator in horizontal writing

AWAY again	[(AWAY again)] TJ
AWAY again	[(A) 120 (W) 120 (A) 95 (Y again)] TJ

8.7.6 Text strings in multi-byte fonts



The text string operators can be used with any string. For strings that use multi-byte encodings, the high-order byte of a character code must appear first. The strings must conform to the syntax for string objects. Therefore care must be taken when including multi-byte character codes. These codes may contain single-byte values that are the same as the ASCII characters for left parenthesis (<28>), right parenthesis (<29>), and backslash (<5C>). When a string is written by enclosing the data in parentheses, these bytes must be preceded by the backslash character. All other byte values between <00> and <FF> may be used in a string object.


8.8 External objects (XObject)

PDF defines three types of XObjects: Image XObjects, Form XObjects, and PostScript XObjects.

8.8.1 XObject operators

The **Do** operator permits the execution of an arbitrary object whose data is encapsulated within a PDF object. The currently defined XObjects are images and PostScript language forms, discussed in [Section 7.13, “XObject.”](#)

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>xobject</i>	Do	Executes the specified XObject. <i>xobject</i> must be the name of an Image, Form, or PostScript XObject in the current Resources dictionary. See Section 7.13, “XObject.”
<i>string</i>	PS	The PS operator provides an in-line equivalent to a PostScript XObject. The PS operator has one argument, a string. When a PS operator is encountered while a document is being printed to a PostScript printer, the contents of the string are placed into the PostScript output as the argument of an instance of the PostScript operator exec . This string is copied without interpretation and may include PostScript comments. In any other case, the PS operator has no other effect. See Section 7.13.8 on page 257 for additional information.




8.9 In-line image objects

In addition to the Image XObject described in [Section 7.13, “XObject,”](#) PDF supports in-line images. An in-line Image Object consists of the operator **BI**, followed by Image XObject key–value pairs, followed by the operator **ID**, followed by the image data, followed by **EI**:

```
<in-line image> ::=
    BI
    <Image XObject key–value pairs>
```

ID
 <lines of data>*
EI

Note If an in-line image does not use **ASCIIHexDecode** or **ASCII85Decode** as one of its filters, **ID** should be followed by a single space. The character following the space is interpreted as the first byte of image data.

Image data may be encoded using any of the standard PDF filters. The key–value pairs provided in an in-line image should not include keys specific to resources: **Type**, **Subtype**, and **Name**. Within in-line images, the standard key names may be replaced by the shorter names listed in [Table 8.1](#). These abbreviations may not be used in Image XObjects, however.

Table 8.1 *Abbreviations for in-line image names*

<i>Name</i>	<i>Abbreviated name</i>
ASCIIHexDecode	AHx
ASCII85Decode	A85
BitsPerComponent	BPC
CCITTFaxDecode	CCF
ColorSpace	CS
DCTDecode	DCT
Decode	D
DecodeParms	DP
DeviceCMYK	CMYK
DeviceGray	G
DeviceRGB	RGB
Filter	F
FlateDecode	FI
Height	H
ImageMask	IM
Indexed	I
Intent	<i>no abbreviation</i>
Interpolate	I
LZWDecode	LZW
RunLengthDecode	RL
Width	W

PDF 1.2

PDF 1.1

Note The in-line format should be used only for small images (4K or less) because viewer applications have less flexibility when managing in-line image data.

In-line images, like Image XObjects, are one unit wide and one unit high in user space and drawn at the origin. Images are sized and positioned by transforming user space using the **cm** operator.

Arguments	Operator	Semantics
	BI	Begins image
	ID	Begins image data
	EI	Ends image

The value of the **CS** or **ColorSpace** key may be a device-dependent color space (**DeviceGray**, **DeviceRGB**, or **DeviceCMYK**, or its abbreviation from the preceding table). The value may not be a device-independent color space or a special color space, with the exception of a limited form of the **Indexed** color space, which may be written as

[/Indexed base hival lookup]

where *base* is a device-dependent color space and *lookup* is a string; see [Section 7.12.10](#), “Indexed color spaces.” The name **/Indexed** may be abbreviated as **/I**.

In PDF 1.2, the value may also be the name of a color space in the current Resources dictionary. In this case, any color space that may be used with an Image XObject may be used for the in-line image (see [Section 7.13.1](#), “Images”).

[Example 8.1](#) shows a 17×17 sample in-line image. The image has 8 bits per component; it is an RGB image that has been LZW and ASCII85 encoded. The **cm** operator has been used to scale the image to render at a size of 17×17 user space units and to be located at an *x*-coordinate of 298 and a *y*-coordinate of 388. The **q** and **Q** operators limit the scope of the **cm** operator’s effect to resizing the image.

Example 8.1 *In-line image*

```

q
17 0 0 17 298 388 cm
BI
/W 17
/H 17
/BPC 8
/CS /RGB
/F [/A85 /LZW]
ID
J1/gKA> .]AN&J?]-<HW]aRVcg*bb.\eKAdVV%/PcZ
... omitted data ...

```

R . s (4KE3&d&7hb*7 [%Ct2HCqC~>
 EI
 Q

8.10 Other operators

8.10.1 Type 3 font operators

Type 3 font operators can be used only within the character definitions inside a Type 3 font. Each Type 3 font definition must begin with either a **d0** or **d1** operator. See Section 5.7 of the *PostScript Language Reference Manual, Third Edition* [1] for details.

Arguments	Operator	Semantics
$w_x w_y$	d0	(d zero) setcharwidth The operands are both numbers.
$w_x w_y ll_x ll_y ur_x ur_y$	d1	(d one) setcachedevice The operands are all numbers.

8.10.2 Compatibility operators

PDF 1.1

PDF does not specify a viewer's behavior when it encounters an undefined page description operator. However, [Appendix G](#) does describe the behavior of the Adobe Acrobat viewers. An Acrobat viewer usually alerts the user when it encounters an undefined page description operator. The operators below modify this behavior.

Arguments	Operator	Semantics
	BX	This operator directs a viewer to not report any undefined operators until a matching EX is encountered. (BX-EX pairs may nest.)
	EX	This operator ends a section of page description in which undefined operators should not be reported.



8.10.3 Marked Content operators

PDF 1.2

The Marked Content operators are used in page descriptions such as the **Contents** stream of a page to indicate a part of the stream that may be significant to an application other than a strict PDF consumer, such as a PDF Viewer. The content

that is marked is not a sequence of bytes in the stream, but a sequence of graphics objects. Each graphics object is fully qualified by the graphics state in which it is rendered.

For example, a graphics application might use these operators to indicate that a certain set of objects constitute a “group.” A text-processing application might use them to maintain a connection between a footnote number in the running text and the footnote itself at the bottom of the page.

There are two kinds of marks, those that bracket a sequence of objects, and those that mark a place in the stream. Bracketed sequences begin with either **BMC** or **BDC**, and they end with **EMC**. **BDC** has the same effect as **BMC** but includes a property list as additional information. Places are marked with either **MP** or **DP**. **DP** has the same effect as **MP** but, like **BDC**, includes a property list.

These operators may appear only *between* graphics objects; they may not occur within a graphics object nor between a graphics state operator and its operands. See [Figure 8.1 on page 320](#).

Bracketed sequences may be nested within each other. A bracketed sequence must be entirely contained within a single stream; it may not cross page boundaries, for example. (The **Contents** key of a Page object is permitted to be either a stream or an array of streams; such an array is considered to be a single stream.)

When Marked Content is used with text, the begin-end Marked Content operators (**BMC/BDC** and **EMC**) and the begin-end text operators (**BT** and **ET**) must be properly (separately) nested. That is, the sequence **BMC BT ... EMC ET** is illegal, as is **BT BMC ... ET EMC**. The sequence **BMC BT ... ET EMC** is legal, as is the sequence **BT BMC ... EMC ET**.

The **BMC** and **MP** operators have only one operand, a *tag* which indicates the role of the operator. The **BDC** and **DP** operators have an additional operand, a list of *properties* that are associated with the mark and whose interpretation is relative to the tag. The properties are represented by a dictionary. This dictionary may be written inline in the content stream if all its values are direct objects. If any value is an indirect object (referring to an object outside the stream), then the list is specified by the name of a Property List in the current Resources dictionary. (See [page 313](#).)

With the exception of the **Subtype** key, PDF makes no assumptions about the properties; interpretation of this dictionary is up to the application or PDF extension that placed the content markers in the stream. It is suggested, however, that any particular extension use keys in a consistent way and always use the same type (or small set of types) for the values of a particular key.

The *tags* that are associated with marks must be registered (see [Appendix F](#)) to prevent conflicting usage when more than one application may be marking a particular content stream. The components of the name, including the registered prefix, must be separated by a single period, and the tag may not begin with a period.

<i>Arguments</i>	<i>Operator</i>	<i>Semantics</i>
<i>tag</i>	BMC	Begin marked content. BMC indicates the beginning of a sequence of graphics objects that is “marked” in some way. <i>tag</i> must be a name; it should indicate the role of the content that is marked.
<i>tag properties</i>	BDC	Begin marked content with a property list. BDC indicates the beginning of a sequence of graphics objects that is “marked” in some way. <i>tag</i> must be a name; it should indicate the role of the content that is marked. <i>properties</i> is either an inline dictionary, that is, a direct object dictionary in the content stream, or it is the name of a property list in the current Resources dictionary.
	EMC	End marked content. EMC indicates the end of a marked sequence of graphics objects. Sequences may be nested.
<i>tag</i>	MP	Mark a point in the content. MP indicates a place within the sequence of graphics objects that is “marked.” MP is not intended for use when some subsequence of the content is being marked: BMC and EMC should be used when the beginning and end of a subsequence is to be indicated. <i>tag</i> must be a name; it should indicate the role of the place that is marked.
<i>tag properties</i>	DP	Mark a point in the content and include a property list. DP is similar to MP , but includes a property list, as BDC does.

Marked content and clipping

When Marked Content is used to bracket a path or text clip object, then additional restrictions apply. A path object may or may not include a clip, and it may or may not be painted; the same is true of text objects. A “clip object” is either an *unpainted, clipped path object* (defined by a sequence including a path segment operator, a clip operator, followed by the **n** operator) or an *unpainted, clipped text object* (defined by a sequence in which text is painted in text-rendering mode 7). If a Marked Content includes *only* clip objects, then the Marked Content applies to those objects. Otherwise, Marked Content does not apply to clip objects.

Nesting of clip objects within Marked Content is allowed. For example, if multiple lines of text are used to mask an image, each line of text may be bracketed by Marked Content, and, the lines of bracketed text may be bracketed by an outer Marked Content. An empty Marked Content within a clip Marked Content is considered to be a nested within the clip Marked Content. An additional restriction is that the save and restore operators (**q** and **Q**) may not occur within Marked Content that is used to bracket clip objects.

The precise rules for determining whether a Marked Content applies to a clip object are as follows:

1. If the only objects within a Marked Content are clip objects, then the Marked Content applies to those clip objects.
2. A Marked Content that contains only clip objects is a clip object.

3. An empty Marked Content that is contained by a clip Marked Content is part of the clip Marked Content. A Marked Place (denoted by **MP** and **DP**) is treated the same as an empty Marked Content.
4. If both clip and marking objects occur between Marked Content delimiters, then the clip objects are not marked by the enclosing Marked Content. That is, any Marked Content attributes do not apply to the clip objects.
5. The largest tree of nested Marked Content operators that contains only empty Marked Content and clip Marked Content is a clip Marked Content.
6. The save and restore operators (**q** and **Q**) may not occur within a Marked Content that is used to bracket clip objects.
7. Marked Content must nest within **BT/ET**, and **BT/ET** must nest within Marked Content.
8. Invisible graphic objects inside Marked Content are treated as rendered objects. They are not clip objects.

Examples

Example 1:

```

/Clip BMC
  100 100 10 10 re W n           clip path
  (Clip me) Tj                  object that is clipped
EMC

```

Example 2:

```

/Clip BMC
  /PointText <<...>> BDC
  BT
    7 Tr                          begin text clip mode
  /Pgf BMC
    (Line 1) Tj
  EMC
  /Pgf BMC
    (Line) Tj ( 2) Tj
  EMC
  ET                              set current text clip
  EMC
  100 100 10 10 re f             filled path
EMC

```

Example 3:

```

/G1 BMC
/G2 BMC

```

```

    /G3 BMC
      0 0 m
      100 100 1
      0 100 1 W n
      0 0 m
      200 200 1
      0 100 1 f
    EMC
    /G4 BMC
      0 0 m
      300 300 1
      0 100 1 W n
    EMC
EMC
100 100 10 10 re f
EMC

```

clip path 1

filled path

clip path 2

filled path

Example 3 shows how nested clip Marked Content is handled. **G3** does not apply to clip path 1 because **G3** also includes a filled path. **G4** *does* apply to clip path 2, but **G2** does *not* apply to clip path 2.

Example 4:

```

    /1 BMC
      <clip path>
    /2 BMC
      /3 BMC
        EMC
      DP
    EMC
EMC

```

BMC 2 contains only an empty Marked Content, and a DP. However, they are all considered clip objects because they are all nested in BMC 1 which is clip-only and has a real clip object in it. BMC 2 is not empty, but it only contains empty Marked Content, so it is a clip object. The same rule applies to BMC 3.

Example 5:

```

    /1 BMC
      /2 BMC
        /3 BMC
          EMC
        EMC
      /4 BMC
        <clip path>
      EMC
EMC

```

Here, BMC 1 becomes clip-only due to the nested clip path, and BMC 2 and MC 3 become clip objects due to their containment in BMC 1.

CHAPTER 9**Linearized PDF**

9.1 Introduction

A linearized PDF file is one that has been organized in a special way to enable efficient incremental access in a network environment. The file is valid PDF in all respects, and it is compatible with all existing viewers and other PDF applications. Enhanced viewers can recognize that a PDF file has been linearized and can take advantage of that organization to enhance viewing performance.

PDF 1.2

The primary goal of the linearized PDF organization is to achieve the following behavior:

1. When a document is opened, display the first page as quickly as possible. The first page to be viewed can be an arbitrary page of the document, not necessarily page 0 (though opening at page 0 is most common).
2. When the user requests another page of an open document (either by going to the next page or by following a link to an arbitrary page), display that page as quickly as possible.
3. When data for a page is delivered over a slow channel, display the page incrementally as it arrives. Insofar as is possible, the most useful data should be displayed first.
4. Permit user interaction, such as following a link, to be performed even before the entire page has been received and displayed.

The above behavior should be achieved for documents of arbitrary size. The total number of pages in the document should have little or no effect on the user-perceived performance of viewing any particular page.

The primary focus of linearized PDF is optimized viewing of read-only PDF documents. It is intended that the linearized PDF will be generated once and read many times. Incremental update is still permitted, but the resulting PDF is no longer linearized and subsequently will be treated as ordinary PDF. Re-linearizing it requires reprocessing the entire file.

Linearized PDF requires two additions to the PDF specification:

1. Rules for ordering of objects in the PDF file.

2. Additional data structures called *hint tables* that enable efficient navigation within the document.

Both of these additions are relatively simple to describe. However, using them effectively requires a deeper understanding of their purpose. Consequently, the following presentation goes considerably beyond a simple specification of PDF extensions; it includes background, motivation, and strategies.

[Section 9.2, “Background and Assumptions,”](#) provides background about the properties of the World Wide Web that are relevant to the design of linearized PDF. [Section 9.3, “Linearized PDF document structure specification,”](#) specifies the file format and object-ordering requirements of linearized PDF. [Section 9.4, “Hint Tables,”](#) specifies the detailed representation of the hint tables. [Section 9.5, “Access Strategies,”](#) outlines strategies for accessing linearized PDF over a network, which in turn determine the optimal way in which to organize the PDF file itself.

The reader is assumed to be familiar with PDF document structure and with the basic architecture of the World Wide Web, and is assumed not to be intimidated by terms such as URL, HTTP, and MIME.

9.2 Background and Assumptions

The principal problem addressed by the linearized PDF design is accessing PDF documents through the World Wide Web. This environment has the following important properties:

1. The access protocol (HTTP) is a transaction consisting of a request and a response. The client presents a request in the form of a URL, and the server sends a response consisting of one or more MIME-tagged data blocks.
2. After a transaction has completed, obtaining more data requires a new request-response transaction. The connection between client and server does not ordinarily persist beyond the end of a transaction, although some implementations may attempt to cache the open connection in order to expedite subsequent transactions with the same server.
3. Round-trip delay can be significant. A request-response transaction can take up to several seconds, independent of the amount of data requested.
4. The data rate may be limited. A typical bottleneck is a 14.4K or 28.8K bit/sec modem link between the client and the Internet service provider.

The above properties are generally shared by other wide-area network architectures aside from the World Wide Web. Additionally, CD-ROMs share some of these properties, since they have relatively slow seek times and limited data rates compared to magnetic media. In the remainder of this chapter, we concentrate on the World Wide Web exclusively.

There are some additional properties of the HTTP protocol that are relevant to the problem of accessing PDF files efficiently. These properties may not all be shared by other protocols or network environments.

5. When a PDF file is initially accessed (say, by following a URL hyperlink from some other document), the file type is not known to the client. Therefore, the client initiates a transaction to retrieve the entire document, then inspects the MIME tag of the response as it arrives. Only at that point is the document known to be PDF. Additionally, the length of the document becomes known at that time.
6. The client can abort a response while it is still in progress, if it decides that the remainder of the data is not of any immediate interest. How quickly the abort takes effect depends on round-trip time and server responsiveness. In HTTP, aborting the transaction requires closing the connection, which will interfere with the strategy of caching the open connection between transactions.
7. The client can request retrieval of portions of a document by specifying one or more byte ranges (*offset, count*) as part of the URL. Each range can be relative to either the beginning or the end of the file. The client can specify as many ranges as it wants in the request, and the response will consist of multiple blocks, each properly tagged.
8. The client can initiate multiple concurrent transactions in an attempt to obtain multiple responses in parallel. This is commonly done, for instance, to retrieve in-line images referenced from a HTML document. This strategy isn't always reliable and may backfire if the transactions interfere with each other by competing for scarce resources in the server or the communication channel.

Note Extensive experimentation has determined that multiple concurrent transactions don't work very well for PDF in some important environments. Therefore, linearized PDF is designed to enable good performance to be achieved using only one transaction at a time. In particular, this means that the client must have sufficient information to determine the byte ranges for all the objects required to display a given page of the PDF file so that it can specify all those byte ranges in a single request.

Finally, we make some additional assumptions about the PDF viewer and its local environment.

9. The viewer has plenty of local temporary storage available. It should rarely need to retrieve a given portion of a PDF document more than once from the server.
10. The viewer is able to display PDF data quickly once it has been received. The performance bottleneck is assumed to be in the transport system (throughput or round-trip delay), not in the processing of data after it arrives.

The consequence of these assumptions is that it may be advantageous for the client to do considerable extra work in order to minimize delays due to communications. Such work includes maintaining local caches and reordering actions according to when the needed data becomes available.

9.3 Linearized PDF document structure specification

Except as noted below, all elements of a linearized PDF file are as specified in [Chapter 5](#).

9.3.1 File structure

Except as noted, all indirect objects in the PDF file are numbered sequentially in two groups, based on their order of appearance in the file.

- The first group consists of the Catalog, certain other document-level objects, and all objects belonging to the first page of the document. These are numbered sequentially, starting at the first object number after the second group. (The stream containing the hint tables may be numbered out of sequence; see [Section 9.3.5, “Hint Streams.”](#))
- The second group consists of all remaining objects in the document, including all pages after the first, all shared objects, etc. These are numbered sequentially, starting at 1.

These groups of objects are indexed by precisely two cross-reference table sections, located as shown below. The composition of these groups is discussed in more detail in the sections that follow. All objects have a generation number of 0.

Example 9.1 Outline of a linearized PDF file

Part 1: Header

```
%PDF-1.1
% binary stuff
```

Part 2: Linearization parameters

```
43 0 obj
<<
  /Linearized 1                               version
  /L 54567                                     file length
  /H [475 598] Primary Hint Stream offset and length (Part 5)
  /O 45    object number of first page's Page object (Part 6)
  /E 5437   offset of end of first page
  /N 11     number of pages in document
  /T 52786  offset of first entry in main xref table (Part 11)
>>
endobj
```

Part 3: First Page xref table and trailer

```
xref
43 14
0000000052 00000 n
0000000392 00000 n
0000001073 00000 n
...cross-reference entries for remaining objects in the first page...
0000000475 00000 n
```



```

trailer
<<
  /Size 57      total number of xref table entries in document
  /Prev 52776  offset of main xref table (Part 12)
  /Root 44 0 R  indirect reference to Catalog (Part 4)
  ...any other attributes, e.g., Info, Encrypt... (Part 9)
>>
startxref
0          dummy xref table offset
%%EOF

```

Part 4: Catalog and other required document-level objects

```

44 0 obj
<<
  /Type /Catalog
  /Pages 42 0 R
>>
endobj
...other objects...

```

Part 5: Primary Hint Stream (Note: Parts 5 and 6 may be placed in the opposite order)

```

56 0 obj
<<
  /Length 457
  ...possibly other stream attributes, e.g., Filter...
  /P 0          position of Page Offset hint table
  /S 221       position of Shared Objects hint table
  ...possibly entries for other hint tables...
>>
stream
Page Offset hint table
Shared Object hint table
...possibly other hint tables...
endstream
endobj

```

Part 6: First Page's objects

```

45 0 obj
<<
  /Type Page
  ...
>>
  Outlines tree (if the PageMode in the Catalog is UseOutlines)
  ...objects for first page, including both shared and non-shared resources...

```

Part 7: All remaining pages (each Page object is followed by the non-shared objects for that page)

```

1 0 obj
<<
  /Type /Page

```

*...other Page attributes, such as **MediaBox**, **Parent**, and **Contents**...*

>>

...objects for that page, including only non-shared resources...

...other pages and their non-shared objects...

last page

...objects for last page, including only non-shared resources...

Part 8: Shared objects for all pages except the first

...shared objects...

Part 9: Other objects not associated with pages, if any

...other objects...

Part 10: Overflow Hint Stream (optional)

Overflow Hint Stream

Part 11: Main xref table and trailer

xref

0 43

0000000000 65535 f

...cross-reference entries for all except First Page's objects...

trailer

<<

/Size 43

*This trailer does not need to contain any other attributes;
in particular, it should not have a **Prev** attribute.*

>>

startxref

257

offset of First Page xref table (Part 3)

%%EOF

9.3.2 Header and linearization information

The file begins with the standard header line (see [Section 5.13 on page 56](#)). Linearization is independent of PDF version number and can be applied to any PDF file version 1.1 or greater.

The *binary stuff* following the percent sign on the second line is some text that includes characters with codes 128 or greater, as recommended in [Section 5.12 on page 55](#).

Following this, the first object in the body of the file (Part 2) must be an indirect dictionary object containing the parameters listed in [Table 9.1](#). All values in this dictionary must be direct objects. Note that there are no references to this dictionary anywhere in the document. (However, there is a normal entry for it in the First Page Cross-Reference Table, Part 3.)



Table 9.1 *Linearization parameters*

<i>Parameter</i>	<i>Type</i>	<i>Semantics</i>
Linearized	number	<i>(Required)</i> Linearized format version identification. As usual, a change in the integer part indicates an incompatible change in the linearized format. A change in the fractional part indicates an upward-compatible change. The current version is 1.0.
L (file Length)	integer	<i>(Required)</i> Length of entire file in bytes. This must be exactly equal to the actual length of the PDF file. A mismatch indicates that the PDF is not linearized and must be treated as ordinary PDF, ignoring linearization information.
H (Hints)	array	<i>(Required)</i> Array of two or four integers, [<i>offset1 length1</i>] or [<i>offset1 length1 offset2 length2</i>]. <i>Offset1</i> is the offset from beginning of the file of the Primary Hint Stream. (This is the beginning of the stream object, not the beginning of the stream data.) <i>Length1</i> is the length of this stream (including stream object overhead). If there is an Overflow Hint Stream, <i>offset2</i> and <i>length2</i> specify its offset and length.
O (Object number)	integer	<i>(Required)</i> Object number of the first page's Page object.
E (End of first page)	integer	<i>(Required)</i> Offset of the end of the first page (i.e., the end of Part 6), relative to the beginning of the file.
N (Number of pages)	integer	<i>(Required)</i> Number of pages in the document.
T (main xref Table)	integer	<i>(Required)</i> Location of the first entry of the main cross-reference table (the entry for object number 0). Note that this differs from the Prev attribute of the First Page Trailer, which gives the location of the xref line that precedes the table.
P (First Page number)	integer	<i>(Optional)</i> Page number of the first page (see Section 9.3.6, "First Page's objects"). The default value is 0.

The **Linearized** dictionary must be entirely contained within the first 1024 bytes of the PDF file. This limits the amount of data a viewer must read before deciding that the file is not linearized.

9.3.3 First Page Cross-Reference and Trailer

(Part 3) This is the cross-reference table for all the first page's objects (discussed in [Section 9.3.6, "First Page's objects"](#)), as well as for the Catalog and document-level objects appearing before the first page (discussed in [Section 9.3.4, "Catalog and document-level objects"](#)). Additionally, it contains entries for the **Linearized** dictionary (at the beginning) and the Primary Hint Stream (at the end).

It is a valid cross-reference table section as defined in [Section 5.15 on page 57](#), although its position in the file is rather unconventional. The table consists of a single cross-reference subsection, with no free entries.

The **startxref** line at the end of the file gives the offset of the First Page Cross-Reference Table. The First Page Trailer's **Prev** entry gives the offset of the main cross-reference table near the end of the file. Once again, this is valid PDF, though the trailers are linked in an unusual order. A PDF viewer that is unaware of linearization interprets the First Page Cross-Reference Table as an "update" to an "original" document that is indexed by the main cross-reference table.

The First Page Trailer must contain valid **Size** and **Root** attributes, as well as any other attributes needed to display the document. The **Size** must be the combined number of entries in both the First Page Cross-Reference Table and the main cross-reference table.

This trailer may optionally end with **startxref**, an integer, and **%%EOF**, just the same as an ordinary trailer. This information is ignored.

9.3.4 Catalog and document-level objects

(Part 4) After the First Page Cross-Reference Table must appear the Catalog dictionary and other objects that are required when opening the document. These objects include:

- The **Catalog** object.
- If the Catalog contains **PageMode** or **OpenAction** entries, those objects must be located here also, with the exception of the Outlines tree: if **PageMode** is **UseOutlines**, the entire Outlines tree is located in Part 6; otherwise it is located in Part 9. See [Section 9.3.9 on page 368](#) for details.
- The Encryption dictionary referenced from the **Encrypt** attribute, if any, of the First Page Trailer. All attribute values in this dictionary must be located here also.
- The **Threads** array in the Catalog, if any, along with all the thread dictionaries that it refers to. This does not include the threads' **Info** dictionaries or the individual beads of the threads.
- The **AcroForm** dictionary in the Catalog, if any. Only the top-level dictionary is needed, not the objects that it refers to.

Objects that are not ordinarily needed when opening the document should not be located here but instead should be at the end of the file; see [Section 9.3.9, "Other objects."](#) This includes **Info**, **Pages**, and **Dests**.

Note that the objects located here are indexed by the First Page Cross-Reference Table, even though they are not logically part of the first page.

9.3.5 Hint Streams

(Part 5) The core of the linearization information is stored in data structures known as *hint tables*, whose format is described in [Section 9.4, "Hint Tables"](#). They provide indexing information that enables the client to construct a single request

for all the objects that are needed to display any page of the document or to retrieve certain other information efficiently. The hint tables may contain additional information to optimize access by plug-ins to application-specific information.

The hint tables are not logically part of the information content of the document; they can be derived from the document. Any action that changes the document—for instance, appending an incremental update—will invalidate the hint tables. The document is still a valid PDF file that just isn't linearized any more.

The hint tables are binary data structures that are enclosed in a stream object. Syntactically, this stream is a normal PDF indirect object. However, there are no references to this stream anywhere in the document, so it is not logically part of the document; any operation that regenerates the document will remove the stream.

Usually, all the hint tables are contained in a single stream, known as the *Primary Hint Stream*. Optionally, there may be an additional stream, containing more hints, known as the *Overflow Hint Stream*. The contents of the two hint streams are to be concatenated and treated as if they were a single unbroken stream.

The Primary Hint Stream, which is required, is shown as Part 5 in [Example 9.1](#), and the First Page section is shown as Part 6. The order of these two parts may be reversed. See [Section 9.5, “Access Strategies,”](#) for considerations on the choice of placement.

The Overflow Hint Stream, Part 10, is optional.

The location and length of the Primary Hint Stream, and of the Overflow Hint Stream if present, are given in the **Linearized** dictionary at the beginning of the file.

The hint streams are assigned the last object numbers in the file, i.e., after the object number for the last object in the first page. Their cross-reference table entries are at the end of the First Page Cross-Reference Table. This object number assignment is independent of the physical locations of the hint streams in the file. (This convention keeps their object numbers out of the way of the numbering of the linearized objects.)

All attributes in the hint streams' dictionaries must be direct objects. The streams may have **Filter** and **DecodeParms** attributes.

In addition to the standard stream attributes, the dictionary of the Primary Hint Stream contains attributes giving the position of the beginning of each hint table in the stream. These positions are in bytes relative to the beginning of the stream data (after applying decoding filters, if any), and with the Overflow Hint Stream concatenated if present. The dictionary of the Overflow Hint Stream should not contain these attributes. The standard hint tables are:

Table 9.2 *Standard Hint Tables*

<i>Key</i>	<i>Hint Table</i>
P	(<i>Required</i>) Page Offset hint table
S	(<i>Required</i>) Shared Objects hint table

T	<i>(Only if thumbnails exist)</i> Thumbnails hint table
O	<i>(Only if outlines exist)</i> Outline hint table
A	<i>(Only if threads exist)</i> Thread info hint table
E	<i>(Only if named destinations exist)</i> Dests hint table
V	<i>(Only if AcroForm dictionary exists)</i> Forms hint table
I	<i>(Only if Info dictionary exists)</i> Info dictionary hint table
C	<i>(Only if logical structure exists)</i> Structure hint table

New keys may be registered for additional hint tables required for new PDF features or for application-specific data accessed by plug-ins.

[Section 9.4, “Hint Tables,”](#) documents the format of the standard hint tables that are enclosed in this stream.

9.3.6 First Page’s objects

(Part 6) As mentioned above, this section may either precede or follow the Primary Hint Stream. The starting file offset and length of this section may be determined from the hint tables. Additionally, the **E** attribute in the **Linearized** dictionary specifies the end of the first page, and the **O** attribute gives the first page’s Page object number.

This part of the file contains all the objects needed to display the first page of the document. Ordinarily, the “first page” is page 0, i.e., the leftmost leaf Page object in the Pages tree. However, if the Catalog contains an **OpenAction** that specifies opening at some page other than page 0, then that page is the “first page” and should be located here. The page number of the first page is given in the **Linearized** dictionary at the beginning of the file.

Implementation note Acrobat always treats page 0 as the first page for linearization, regardless of OpenAction.

The objects contained here should include:

1. **Page** object for the first page. This must be the first object in this part of the file. Its object number is given in the **Linearized** dictionary. This Page object must explicitly specify all required attributes, such as **Resources** and **MediaBox**; the attributes cannot be inherited from ancestor **Pages** objects.
2. The entire Outlines tree, if the **PageMode** key in the Catalog is **UseOutlines**. (If the **PageMode** key is omitted or has some other value and the document has an Outlines tree, then it appears in Part 9. See [Section 9.3.9 on page 368](#) for details.)
3. All objects that the Page object refers to, to arbitrary depth. This includes **Contents**, **Resources**, **Annots**, and **B** (Beads), but excludes **Thumb**.

The order of objects referenced from the Page object should facilitate early user interaction and incremental display of the page data as it arrives. The following order is recommended:

1. The **Annots** array and all annotation objects, to a depth sufficient to allow those annotations to be activated. Information required to draw the annotation can be deferred until later, since annotations are always drawn on top of (hence after) the Contents.
2. The **B** (Beads) array and all bead dictionaries, if any, for this page. If any beads exist for this page, the **B** array is required to be present in the Page dictionary. Additionally, each bead in the thread (not just the first) must contain a **T** attribute referring to the associated thread dictionary.
3. The **Resources** dictionary, but not the resource objects contained in the dictionary.
4. Resource objects, other than the types listed below, in the order that they are first referenced (directly or indirectly) from the Contents stream. If Contents is represented as an array of streams, each resource object should precede the stream in which it is first referenced. Note that Font, FontDescriptor, and Encoding resources should be included here, but not substitutable FontFiles referenced from FontDescriptors (see below).
5. **Contents**. If it is large, it should be represented as an array of indirect references to streams, which in turn are interleaved with the resources that they require. If it is small, the entire Contents should be a single stream preceding the resources.
6. Image XObjects, in the order that they are first referenced. Images are assumed to be large and slow to transfer, so the viewer defers rendering images until all the other Contents have been displayed.

FontFile streams, which contain the actual definitions of embedded fonts. These are assumed to be large and slow to transfer, so the viewer draws substitute fonts until the real ones have arrived. Only those fonts for which substitution is possible can be deferred in this way (Currently, this includes any Type 1 or TrueType font that has a **FontDescriptor** and whose Flags bit 6 is set (indicating Adobe standard Roman character set). The base 14 Type 1 fonts cannot be deferred, although it is unlikely that they would be embedded.

See [Section 9.5, "Access Strategies,"](#) for additional discussion about object order and incremental drawing strategies.

9.3.7 Objects contained in remaining pages

(Part 7) This part of the file contains the non-shared objects for all remaining pages of the file, with the objects for each page grouped together. The pages are contiguous and are ordered by page number. (If the first page of the file is not page 0, this section starts with page 0 and skips over the first page when its position in the sequence is reached.)

For each page, the objects required to display that page are grouped together, except for resources and other objects that are shared with other pages. Shared objects are located in the Shared Objects section, described below. The starting file offset and length of any page can be determined from the hint tables.

The recommended order of objects within a page is essentially the same as in the first page. In particular, the Page object must be the first object in each section.

In most cases, there will be little benefit from interleaving contents with resources. This is because most resources other than images—fonts in particular—are shared among multiple pages and therefore reside in the Shared Objects section. Image XObjects usually are not shared, but they should appear at the end of the page, since rendering of images is deferred.

9.3.8 Shared objects

(Part 8) This portion of the file contains objects, primarily named resources, that are referenced from more than one page and that are not referenced (directly or indirectly) from the first page. The hint tables contain an index of these objects.

The order of these objects is essentially arbitrary. However, wherever a resource consists of a multiple-level structure, all components of the structure should be grouped together. If only the top-level object is referenced from outside the group, the entire group can be described by a single entry in the Shared Object hint table. This helps to minimize the size of the Shared Object hint table and the number of individual references from entries in the Page Offset hint table.

Implementation note Acrobat does not generate Shared Object groups containing more than one object.

9.3.9 Other objects

(Part 9) Following the shared objects are any other objects that are part of the document but aren't required for displaying pages. These objects are divided into functional categories. Objects within each of these categories should be grouped together. The relative order of the categories is unimportant.

- Pages tree. This can be located here, since the linearized PDF viewer never needs to consult it. Note that all Resources and other inheritable attributes of the Pages objects must be pushed down and replicated in each of the leaf Page objects (but they may contain indirect references to shared objects).
- Thumbnails. These should simply be ordered by page number. Note that the thumbnail for page 0 should be first, even if the first page of the linearized PDF is some page other than 0. Each thumbnail consists of one or more objects. These objects may refer to objects in the Thumbnail Shared Objects section (see the next item).
- Thumbnails Shared Objects. These are objects that are shared among some or all thumbnail objects and are not referenced from any other objects.
- Outline tree, if not located in Part 6. The order of objects should be the same as the order in which they are displayed by the viewer. This is a preorder traversal of the tree, skipping over any subtree that is closed (i.e., whose parent's **Count** is negative). Following that should be the subtrees that were skipped over, in the order that they would have appeared if they were all open.

- Thread info dictionaries, referenced from the **I** (Info) attributes of thread dictionaries. Note that the thread dictionaries themselves are co-located with the Catalog, and the beads with the individual pages.
- Named destinations. These objects include the **Dests** or **Names** attribute of the Catalog and all the destination objects that it refers to. See [Section 9.5.2, “Opening at an arbitrary page.”](#)
- Info dictionary and the objects contained within it.
- AcroForm tree. This does not include the top-level AcroForm dictionary, which is co-located with the Catalog.
- Other entries in the Catalog that aren’t referenced from any page.
- Logical structure tree.

PDF 1.3

9.3.10 Main cross-reference and trailer

(Part 11) This is the cross-reference table for all objects in the PDF file except those listed in the First Page Cross-Reference Table (Part 3). As indicated earlier, this cross-reference table plays the role of the “original” cross-reference table for the file (prior to appending any “updates”). It must conform to the PDF rules for this table:

- It consists of a single cross-reference subsection, beginning at object number 0.
- The first entry (for object number 0) must be a free entry.
- The remaining entries are for in-use objects, which are numbered consecutively starting at 1.

As indicated earlier, the **startxref** line gives the offset of the First Page Cross-Reference Table. The **Prev** entry of the First Page Trailer gives the offset of the main cross-reference table. The main trailer has no **Prev** entry, and in fact does not need to contain any entries other than **Size**.

9.4 Hint Tables

There are two or more hint tables, as indicated by the attributes of the Primary Hint Stream (see [Section 9.3.5, “Hint Streams”](#)). The format of the standard hint tables is described below.

There can be additional hint tables for application-specific data accessed by plugins. A generic format for such hint tables is defined; see [Section 9.4.4, “Generic hint tables.”](#) Alternatively, the format of a hint table can be private to the application.

Each hint table consists of a portion of the stream, beginning at the position in the stream indicated by the corresponding stream attribute. (If there is an Overflow Hint Stream, its contents are to be appended seamlessly to the primary Hint

Stream. Hint table positions are relative to the beginning of this combined stream.) In general, this byte stream is treated as a bit stream, high-order bit first, which is then subdivided into fields of arbitrary width without regard to byte boundaries. However, each hint table begins at a byte boundary.

The hint tables are designed to encode the required information as compactly as possible. Interpreting the hint tables requires reading them sequentially; they are not designed for random access. The client is expected to read and decode the tables once and retain the information for as long as the document remains open.

A hint table encodes the positions of various objects in the file. The representation is either explicit (an offset from the beginning of the file) or implicit (accumulated lengths of preceding objects). Regardless of the representation, the resulting positions must be interpreted as if the Primary Hint Stream itself were not present. That is, a position greater than the *hint stream offset* must have the *hint stream length* added to it in order to determine the actual offset relative to the beginning of the file. (The *hint stream offset* and *hint stream length* are the values *offset1* and *length1* in the **H** array in the **Linearized** dictionary at the beginning of the file.)

The reason for this rule is that the length of the Primary Hint Stream depends on the information contained within the hint tables, and this is not known until after they have been generated. Any information that gets put into the hint tables must not depend on knowing the Primary Hint Stream's length in advance.

Note that this rule applies only to offsets given in the hint tables, and not to offsets given in the cross-reference tables or **Linearized** dictionary. Also, the offset and length of the Overflow Hint Stream, if present, need not be taken into account, since this object follows all other objects in the file.

9.4.1 Page Offset hint table

The Page Offset hint table gives information required to locate each page. Additionally, for each page except the first, the page entry enumerates all shared objects that the page references, directly or indirectly.

This table consists of a header section, described in [Table 9.3](#), followed by one or more per-page entries, described in [Table 9.4](#).

Note All the “bits needed” items, such as item 3 in the following table, may have values in the range 0 to 32. Although that range requires only 6 bits, 16-bit numbers are used.

Table 9.3 Page Offset hint table, header section

<i>Item</i>	<i>Size (bits)</i>	<i>Description</i>
1	32	Least number of objects in a page (including the Page object itself).
2	32	Location of first page's Page object.
3	16	Bits needed to represent the greatest number of objects in a page.

4	32	Least length of a page, in bytes. This is the least length from the beginning of a Page object to the last byte of the last object used by that page.
5	16	Bits needed to represent the greatest page length.
6	32	Least start of Contents offset.
7	16	Bits needed to represent the greatest start of Contents offset.
<i>Implementation note</i> In Acrobat, item 7 is unused and always set to zero.		
8	32	Least Contents length.
9	16	Bits needed to represent the greatest Contents length.
10	16	Bits needed to represent the greatest number of Shared Object references.
11	16	Bits needed to identify a Shared Object.
12	16	Bits needed to represent numerator of fraction (see Table 9.4).
13	16	Denominator used to divide page Contents into fractions. For each shared object referenced from a page, there is an indication of where in the page's Contents the object is first referenced. That position is given as the numerator of a fraction, whose denominator is specified once for the entire document. The fraction is explained in more detail below.

Table 9.4 Page Offset hint table, per-page entry

<i>Item</i>	<i>Size (bits)</i>	<i>Description</i>
1	(see Table 9.3 , item 3)	This value, when added to the least number of objects in a page (Table 9.3 , item 1), gives the number of objects in the page. The first object of the first page has an object number that is the value of the O attribute in the Linearized dictionary at the beginning of the file. The first object of the second page has an object number of 1. Objects numbers for subsequent pages can be determined by accumulating the number of objects in all previous pages.
<i>Note</i> The entries for all the pages are listed in ascending order. That is, the table contains item 1 for the first page, item 1 for the second page, item 1 for the third page, and so on, and then continues with item 2 for the first page, item 2 for the second page, and so on.		
2	(Table 9.3 , item 5)	This value, when added to the least page length (Table 9.3 , item 4), gives the total length of the page in bytes. The location of the first object of the first page can be determined from the cross-reference table entry for that object (see above). The locations of subsequent pages can be determined by accumulating the lengths of all previous pages. Note that one must skip over the Primary Hint Stream, wherever it is located.

Note The entries for all the pages are listed in ascending order.

- 3 (Table 9.3, item 10) Number of shared objects referenced from page. Note that this must be 0 in the first page's entry.
- Note* The entries for all the pages are listed in ascending order.
- 4 (multiple) variable Shared Object references, each consisting of a Shared Object identifier (see Table 9.3, item 11) followed by the numerator.
- A Shared Object identifier is an index into the Shared Object hint table, described in Section 9.4.2, "Shared Object hint table." Note that a single entry in the Shared Object hint table can designate a group of shared objects, only one of which is referenced from outside the group. That is, Shared Object identifiers are not directly related to object numbers.
- Note* The entries for all the pages are listed in ascending order, beginning with the second page (since the number of shared objects for the first page is 0).
- 5 (multiple) variable The numerator of the fractional position (see Table 9.3, item 12) for each Shared Object reference, in the same order as the preceding item.
- The fraction indicates where in the page's Contents the shared object is first referenced. It is interpreted as the numerator of a fraction whose denominator is specified once for the entire document.
- If the denominator is d , a numerator of 0 indicates that the first reference lies in the interval $0/d$ to $1/d$ of the Contents. Similarly, a numerator of $d-1$ indicates that the first reference lies in the interval $(d-1)/d$ to the end of the Contents.
- The numerator can take on two (or more) additional values, which indicate that the shared object is not referenced from the Contents but is needed by annotations or other objects that are drawn after the Contents. The value d indicates that the shared object is needed before Image XObjects and other non-shared objects that are at the end of the page. The value $d+1$ or greater indicates that the shared object is needed after those objects.
- This method of dividing the page into fractions is only approximate. Determining the first reference to a shared object entails inspecting the unencoded Contents stream. The relationship between positions in the unencoded and encoded streams is not necessarily linear.
- 6 (Table 9.3, item 7) This value, when added to the least start of Contents offset (Table 9.3, item 6), gives the offset in bytes of the start of the Contents stream, relative to the beginning of the page. This is the offset of the stream object, not the stream data.
- Implementation note* In Acrobat, item 7 of Table 9.3 is assumed to be zero.
- 7 (Table 9.3, item 9) This value, when added to the least Contents length (Table 9.3, item 8), gives the length of the Contents stream in bytes. This includes object overhead preceding and following the stream data.

9.4.2 Shared Object hint table

The Shared Object hint table gives information required to locate shared objects (see [Section 9.3.8, “Shared objects”](#)). Shared objects can be physically located in either of two places. Objects that are referenced from the first page are co-located with the First Page objects (Part 6). All other shared objects are located in the Shared Objects section (Part 8).

A single entry in the Shared Object hint table can actually describe a group of adjacent objects, under the following condition: Only the first object in the group is referenced from outside the group. The remaining objects in the group are referenced only from other objects in the same group.

The Page Offset hint table refers to an entry in the Shared Object hint table by a simple index that is its sequence in the table, counting from 0.

This table consists of a header section, described in [Table 9.5](#), followed by one or more Shared Object Group entries, described in [Table 9.6](#). There are two sequences of Shared Object Group entries: the ones for objects located in the first page, followed by the ones for objects located in the Shared Objects section. The entries have the same format in both cases.

For convenience of representation, the first page is treated as if it consisted entirely of shared objects. That is, the first entry refers to the beginning of the first page and has an object count and length that span all the initial non-shared objects. The next entry refers to a group of shared objects. Subsequent entries span additional groups of either shared or non-shared objects consecutively, until all shared objects in the first page have been enumerated. (Obviously, the entries that refer to non-shared objects will never be used.)

Note If there are no shared objects, all of the entries in this table should be zero.

Table 9.5 Shared Object hint table, header section

<i>item</i>	<i>Size (bits)</i>	<i>Description</i>
1	32	Object number of first object in Shared Objects section (part 8).
2	32	Location of first object in Shared Objects section.
3	32	Number of Shared Object entries for first page.
4	32	Number of Shared Object entries for Shared Objects section.
5	16	Bits needed to represent the greatest number of objects in a Shared Object Group.
<i>Implementation note</i>		<i>Acrobat always sets item 5 to zero.</i>
6	32	Least length of a Shared Object Group in bytes.
7	16	Bits needed to represent the greatest length of a Shared Object Group.

Table 9.6 Shared Object hint table, Shared Object Group entry

Item	Size (bits)	Description
1	(Table 9.5, item 7)	This value, when added to the least Shared Object Group length (Table 9.5, item 6), gives the total length of the object group in bytes. The location of the first object of the first page is given in the Page Offset hint table, header section (Table 9.3, item 4). The locations of subsequent object groups can be determined by accumulating the lengths of all previous object groups until all shared objects in the first page have been enumerated. Following that, the location of the first object in the Shared Objects section can be obtained from the Shared Objects hint table, header section (Table 9.5, item 2).
		<i>Note</i> The entries for all the Shared Object Groups are listed, in order. That is, the table contains item 1 for the first group, item 1 for the second group, item 1 for the third group, and so on, and then continues with item 2 for the first group, item 2 for the second group, and so on.
2	1	A flag indicating whether the MD5 Signature (see item 4) is present (1 if present, 0 if absent).
		<i>Note</i> The entries for all the Shared Object Groups are listed, in order.
		<i>Implementation note</i> MD5 Signatures are not implemented in Acrobat. This flag must be 0.
3	(see Table 9.5, item 5)	This value plus 1 gives the number of objects in the group. The first object of the first page is the one whose object number is given by the O attribute in the Linearized dictionary at the beginning of the file. Object numbers for subsequent entries can be determined by accumulating the number of objects in all previous entries, until all shared objects in the first page have been enumerated. Following that, the first object in the Shared Objects section has a number that can be obtained from the Shared Object hint table header (Table 9.5, item 1).
		<i>Note</i> The entries for all the Shared Object Groups are listed, in order.
		<i>Implementation note</i> Acrobat does not support more than one shared object in a group, so the value of this item should always be 0.
4	128	(Optional) MD5 Signature. The MD5 Signature is a 16-byte MD5 hash that uniquely identifies the resource that the group of objects represents. This is intended to enable the client to substitute a locally cached copy of the resource instead of reading it from the PDF file. Note that this signature is unrelated to the signature form field defined in Section 6.14.12, “Signature field.”

9.4.3 Thumbnails hint table

The Thumbnail hint table consists of a header section, described in Table 9.7, followed by one or more per-page entries, described in Table 9.8.

Each entry describes the thumbnail for a single page. The pages are considered in page number order, starting at page 0 (even if page 0 is not the first page of the file). Thumbnails can exist for some but not all pages.

Table 9.7 *Thumbnails hint table, header section*

<i>Item</i>	<i>Size (bits)</i>	<i>Description</i>
1	32	Object number of first object in Thumbnail section.
2	32	Location of first object in Thumbnail section.
3	32	Number of entries in thumbnail table.
4	16	Bits needed to represent the greatest number of consecutive pages that have no thumbnails. This can be zero, meaning that all pages have thumbnails.
5	32	Least length of thumbnail in bytes.
6	16	Bits needed to represent the greatest length of a thumbnail.
7	32	Least number of objects in a thumbnail.
8	16	Bits needed to represent the greatest number of objects in a thumbnail.
9	32	Object number of first object in Thumbnail Shared Objects section (a subsection of Part 8). Those are objects that are referenced from some or all thumbnail objects and are not referenced from any other objects. The Thumbnail Shared Objects are undifferentiated; there is no indication of which shared objects are referenced from any given page's thumbnail.
10	32	Location of the first object in Thumbnail Shared Objects section.
11	32	Number of Thumbnail Shared Objects.

Table 9.8 *Thumbnails hint table, per-page entry*

<i>Item</i>	<i>Size (bits)</i>	<i>Description</i>
1	(see Table 9.7 , item 4)	(<i>Optional</i>) Count of preceding pages lacking thumbnails. This indicates how many pages without thumbnails lie between the previous entry's page and this one.
2	(Table 9.7 , item 8)	This value, when added to the least number of objects in a thumbnail (Table 9.7 , item 7), gives the number of objects in this thumbnail.
3	(Table 9.7 , item 6)	This value, when added to the least length of a thumbnail (Table 9.7 , item 5), gives the length of this thumbnail in bytes.

9.4.4 Generic hint tables

Certain categories of objects are associated with the document as a whole rather than with individual pages (see [Section 9.3.9, “Other objects”](#)). It is sometimes useful to provide hints for accessing those objects efficiently. For each category of hints, there is a separate entry in the Primary Hint Stream giving the starting position of the table within the stream (see [Section 9.3.5, “Hint Streams”](#)).

There is a generic representation for such hints, specified below. This representation is useful for some standard categories of objects, such as outlines, threads, and named destinations. It may also be useful for application-specific objects accessed by plug-ins. It is considerably more convenient for a plug-in to use the generic hint representation than to specify custom hints.

A generic hint table describes a single group of objects that are located together in the PDF file. See [Table 9.9](#).

Table 9.9 *Generic Hint Table*

<i>item</i>	<i>Size (bits)</i>	<i>Description</i>
1	32	Object number of first object in group.
2	32	Location of first object in group.
3	32	Number of objects in group.
4	32	Length of object group in bytes.

9.4.5 Outline, Thread Info, Dests, and Info hint tables

These tables use the generic hint table representation; see [Section 9.4.4, “Generic hint tables.”](#) The objects that they refer to are grouped together as described in [Section 9.3.9, “Other objects.”](#)

9.4.6 Forms and Logical Structure hint tables

If an AcroForm or StructTrees dictionary is present, these tables refer to the contents of those dictionaries; see [Section 9.3.9, “Other objects.”](#) Forms and structure can refer to objects that are also shared with other parts of the document. The table lists those shared objects.

A Form or Structure hint table begins with a generic hint table, described in [Section 9.4.4, “Generic hint tables.”](#) It then continues as described in [Table 9.10](#).

Table 9.10 *Forms and Structure hint tables, continued*

<i>Item</i>	<i>Size (bits)</i>	<i>Description</i>
5	32	Number of Shared Object references.

6	16	Bits needed for each Shared Object reference.
7...	(Table 9.3, item 11)	Shared Object references, each consisting of a Shared Object identifier. See Section 9.4.2, “Shared Object hint table.”

9.5 Access Strategies

This section outlines how the client can take advantage of the structure of a linearized PDF file in order to retrieve and display it efficiently. This material is not formally a part of the linearized PDF specification, but it may help to explain the rationale for the organization.

9.5.1 Opening at the first page

As indicated earlier, when a document is initially accessed, a request is issued to retrieve the entire file, starting at the beginning. Consequently, linearized PDF is organized so that all the data required to display the first page is at the beginning of the file. This includes all resources that are referenced from the first page, whether or not they are also referenced from other pages.

The first page is usually but not necessarily page 0. If the Catalog contains an **OpenAction** that specifies opening at some page other than page 0, that page will be the one physically located at the beginning of the document. Thus, opening a document at the default place (rather than a specific destination) requires simply waiting for the first page data to arrive; no additional transactions are required.

In an ordinary PDF viewer, opening a document requires first positioning to the end to obtain the **startxref** line. Since a linearized PDF file has the first page’s cross-reference table at the beginning, reading the **startxref** line is not necessary. All that is required is to verify that the file length given in the **Linearized** dictionary at the beginning of the file matches the actual length of the file, indicating that no updates have been appended to the PDF file.

The Primary Hint Stream is located either before or after the First Page objects. This means that it will also be retrieved as part of the initial sequential read of the file. The client is expected to interpret and retain all the information in the hint tables. They are reasonably compact and are not designed to be obtained from the file in random pieces.

The client must now decide whether to continue reading the remainder of the document sequentially or to abort the initial transaction and access subsequent pages using separate transactions requesting byte ranges. This decision is a function of the size of the file, the data rate of the channel, and the overhead cost of a transaction.

9.5.2 Opening at an arbitrary page

The viewer may be requested to open a PDF file at an arbitrary page. The page can be specified in one of three ways:

- by page number (**GoToR** link action, integer page specifier)
- by named destination (**GoToR** link action, name or string page specifier)
- by thread (**Thread** link action).

Additionally, an indexed search results in opening a document by page number. Handling this case efficiently is considered especially important.

As indicated above, when the document is initially opened, it is retrieved sequentially starting at the beginning. As soon as the hint tables have been received, the client has sufficient information to request retrieval of any page of the document given its page number. Therefore, it can abort the initial transaction and issue a new transaction for the target page, as described in [Section 9.5.3, “Going to another page of an open document.”](#)

The position of the Primary Hint Stream (Part 5) with respect to the First Page objects (Part 6) determines how quickly this can be done. If the Primary Hint Stream precedes the First Page objects, the initial transaction can be aborted very quickly. However, this is at the cost of increased delay when opening the document at the first page. On the other hand, if the Primary Hint Stream follows the First Page objects, displaying the first page is quicker (since the hint tables are not needed for that), but opening at an arbitrary page is delayed by the time required to receive the first page.

At the time a PDF file is linearized, one must decide whether to favor opening at the first page or opening at an arbitrary page.

If an Overflow Hint Stream exists, obtaining it requires issuing an additional transaction. For this reason, inclusion of an Overflow Hint Stream in linearized PDF, although permitted, is not recommended. The feature exists to allow the linearizer to write the PDF file with space reserved for a Primary Hint Stream of an estimated size, then go back and fill in the hint tables. If the estimate is too small, the linearizer can append an overflow stream containing the remaining hint table data. This allows writing the PDF file in one pass, which may be an advantage if the performance of writing PDF is considered important.

Opening at a named destination requires the viewer first to read the entire **Dests** or **Names** dictionary, for which a hint is present. Using this information, one can determine the page containing the specific destination identified by the name.

Opening at a thread requires the viewer first to read the entire Threads array, which is located with the Catalog at the beginning of the document. Using this information, one can determine the page containing the first bead of any thread. Opening at other than the first bead of a thread requires chaining through all the beads until the desired one is reached; there are no hints to accelerate this.

9.5.3 Going to another page of an open document

Given the information in the hint tables, it is now straightforward for the client to construct a single request to retrieve any arbitrary page of the document. The request should include:

1. The objects of the page itself, whose byte range can be determined from the entry in the Page Offset hint table.
2. The portion of the main cross-reference table referring to those objects. This can be computed from main cross-reference table location (**T** attribute in the **Linearized** dictionary) and the cumulative object number in the Page Offset hint table.
3. The shared objects referenced from the page, whose byte ranges can be determined from information in the Shared Object hint table.
4. The portion of the main cross-reference table referring to those objects, as in (2) above.

The purpose of the fractions in the Page Offset hint table is to enable the client to schedule retrieval of the page in a way that allows incremental display of the data as it arrives. It accomplishes this by constructing a request that interleaves pieces of the page contents with the shared resources that the contents refer to. This serves much the same purpose as the physical interleaving that is done for the first page.

9.5.4 Drawing a page incrementally

The ordering of objects in pages and the organization of the hint tables is intended to allow progressive update of the display and early opportunities for user interaction when the data is arriving slowly. The viewer must recognize instances in which the targets of indirect object references haven't arrived yet and, where possible, rearrange the order in which it acts on the objects in the page.

The following sequence of actions is recommended:

1. Activate the annotations, but don't draw them yet. Also activate the cursor feedback for any article threads in the page.
2. Begin drawing the Contents. Whenever there is a reference to an Image XObject that hasn't arrived yet, skip over it. Whenever there is a reference to a font whose definition is an embedded **FontFile** that hasn't arrived yet, draw the text using a substitute font (if that is possible).
3. Draw the annotations.
4. Draw the images, together with anything that overlaps them.
5. Once the embedded font definitions have arrived, redraw the text using the correct fonts, together with anything that overlaps the text.

(The last two steps should be done using an off-screen buffer, if possible, to avoid objectionable flashing during the redraw process.)

9.5.5 Following an article thread

As indicated earlier, the bead objects for any article thread that visits a given page are located with that page. This enables the bead rectangles to be activated and proper cursor feedback to be shown.

If the user follows a thread, the viewer can obtain the object number from the **N** (Next) or **P** (Prev) attribute of the bead. This identifies a target bead object, which is co-located with the page to which it belongs. Given this object number, the viewer can perform a binary search in the Page Offset hint table to determine which page. It can then go to that page, as discussed in [Section 9.5.3, “Going to another page of an open document.”](#)

9.5.6 Accessing an updated file

As stated earlier, if a linearized PDF file subsequently has an incremental update appended to it, the linearization and hints are no longer valid. Actually, this is not necessarily true, but the viewer must do some additional work to validate them.

When the viewer sees that the file is longer than the length given in the **Linearized** dictionary, it must issue an additional transaction to read everything that was appended. It must then analyze the objects in that update to see if any of them modifies objects that are in the first page or that are the targets of hints. If so, it must augment its internal data structures as necessary to take the updates into account.

For a PDF file that has received only a small update, this approach may be worthwhile. Accessing the file this way will be quicker than accessing it without hints or retrieving the entire file before displaying any of it.



Section II

Optimizing PDF Files

:

March 11, 1999

CHAPTER 10**General Techniques for
Optimizing PDF Files**

The first section of this book describes the syntax allowed in a PDF file. In many cases there is more than one way to represent a particular construct, and the previous chapters do not indicate which alternative is preferred. This section describes techniques to optimize PDF files. Most optimizations reduce the size of a PDF file, reduce the amount of memory needed to display pages, or improve the speed with which pages are drawn. Some optimizations, such as sharing of resources, allow a viewer application to display a document when it may not have otherwise been possible in low memory situations. A few optimizations improve the appearance of pages.

This chapter contains techniques that can be generally applied to PDF files. Following chapters discuss optimizations specifically for text, graphics, and images.

While it may not be possible to take advantage of all the techniques described here, it is worth taking more time producing a PDF file to improve its viewing performance. A PDF file will be produced only once but may be viewed many times.

File size is a good gauge of the level of optimization, but of course the most accurate measure is the time it takes to view and print the pages of a document.

10.1 Use short names

Names in PDF files specify resources, including fonts, forms, images, and other objects. Whenever a name is used, it should contain as few characters as possible. This minimizes the space needed to store references to the object.

Instead of specifying a name as:

```
/FirstFontInPage4  
/SecondImageInPage8
```

use names such as:

```
/F1  
/Im8
```

Resource names need not be unique throughout a document. The names of resource objects must be unique within a given resource type within a single marking context. For example, the names of all fonts on a page must be unique.

10.2 Use direct and indirect objects appropriately

As mentioned in [Chapter 4](#), objects contained in composite objects such as arrays and dictionaries may either be specified directly in the composite object or be referred to indirectly. Using indirect objects frequently improves performance and reduces the size of a PDF file. In addition, programs that produce PDF files sometimes must write a reference to an object into a PDF file before the object's value is known. Indirect objects are useful in this situation.

10.2.1 Minimizing object size

Although PDF allows random access to objects in a file, it does not permit random access to the substructure that may be present in a single object, such as the individual key–value pairs in a dictionary object. If a PDF viewer application needs to access a particular piece of information contained in an object, it reads the entire object. However, if it encounters an indirect object reference, it will not read the indirect object until needed. Using indirect objects minimizes the amount of extra data a PDF viewer application must read before locating the desired information.

As an example, if a PDF viewer application needs to obtain the PostScript language name of a font, it must search the appropriate Font dictionary object. If (in that dictionary object) the **Widths** array is specified directly, the application must read the entire array. If the **Widths** array is specified by an indirect reference, the application only needs to read the few bytes that specify the indirect reference and can avoid reading the **Widths** array itself.

In general, using indirect references improves the performance of a PDF file. However, there is some overhead associated with locating an indirect object, and an indirect object takes up more space than a direct object in a PDF file. Because of this, small objects should not be specified indirectly. A rough rule of thumb is that arrays with more than five elements and dictionaries with more than three key–value pairs should be stored as indirect objects.

10.2.2 Sharing objects

Indirect objects can be referred to from more than one location in a file. Because of this, using indirect objects can decrease the size of a PDF file by defining an object only once and making multiple references to it.

As an example, suppose each page in a document requires the same ProcSets. Each page's Resources dictionary can refer to the same ProcSet array indirectly instead of duplicating the array.

10.2.3 Placeholder for an unknown value

Indirect objects can also be used when an object must be written at one location in a file, but its value will not be known until later in the process of writing the file. The best example of this situation is the **Length** key in the dictionary of a Stream object. The dictionary must be placed in the file ahead of the stream data itself, and must include the **Length** key, which specifies the length of the stream that follows. It may not be possible to know the length of the stream until after the data has been written, however. By specifying the value of the **Length** key as an indirect object, the length of the stream can be written after the stream.

10.3 Take advantage of combined operators

PDF provides several operators that combine the function of two or more other operators. For example, PDF defines operators that close (**h**) and stroke (**S**) a path, but also provides an operator that performs both operations (**s**). These combined operators should be used whenever possible. Table 10.1 lists the combined operators provided by PDF. Some operators in the table require one or more operands; the operands have been omitted from the table.

Table 10.1 *Optimized operator combinations*

<i>Use...</i>	<i>Instead of...</i>
s	h S
b	h B
TD	Td TL
TJ	Repeated series of Tj and Td operators
'	Td Tj or T* Tj
"	Tc Tw Td Tj

Note To both fill and stroke a path, the combination operators must be used. Using the fill operator followed by the stroke operator does not work. The fill operator ends the path, leaving nothing for the stroke operator to stroke. Unlike the PostScript language, PDF does not allow you to save the path, fill it, restore the path, and stroke it, because the current path is not part of the PDF graphics state.

10.4 Remove unnecessary clipping paths

Whenever anything is drawn on a page, all marks are made inside the current clipping path. When a clipping path other than the default (the crop box) is specified, rendering speed is reduced. If a portion of a page requires the use of a clipping path other than the default, the default clipping path should be restored as soon as possible. Text, graphics, and images are all clipped to the current clipping path, so it is important for the performance of all three to not use unnecessary clipping paths.

Restoration of a clipping path can be accomplished by saving the graphics state (which includes the clipping path) using the **q** operator before setting the new clipping path, and subsequently using the **Q** operator to restore the previous clipping path as soon as the new clipping path is no longer needed.

*Note Remember that the **Q** operator restores more than just the clipping path. See [Section 8.3.2, “Special Graphics State operators”](#) for a list of the graphics state parameters restored by the **Q** operator.*

10.5 Omit unnecessary spaces

Spaces are unnecessary before **(**, after **)**, and before and after **[** and **]**. For instance an array should be written **[1 2 3]** instead of **[1 2 3]**. This slightly reduces the size of files.

10.6 Omit default values

A number of the parameters that affect drawing have default values that are initialized at the start of every page. (See [Chapter 8](#).) For example, the default stroke and fill colors are both black. When drawing, do not explicitly set a drawing parameter unless the default value is not the desired value.

Similarly, many PDF objects are represented by dictionaries and some of the keys in these dictionaries have default values. Omit any keys whose default value is the desired value.

Omitting unnecessary key–value pairs and graphics and text state operators reduces the size of a PDF file and the time needed to process it.

10.7 Take advantage of Form XObjects

PDF files may contain Form XObjects, which are arbitrary collections of PDF operators that draw text, graphics, or images. The structure of a Form XObject is discussed in [Section 7.13, “XObjects.”](#) A Form XObject may be used to draw the same marks in one or more locations on one or more pages.

Form XObjects can be used, for example, to draw a logo, a heading for stationery, or a traditional form. The location and appearance of a form is controlled by the CTM in effect when the Form XObject is drawn.

The use of Form XObjects can reduce the size of a PDF file. In addition, Form XObjects that contain an XUID can be cached by PDF viewer applications and PostScript printers, improving rendering speed if the Form XObject is used multiple times.

10.8 Limit the precision of real numbers

The pixel size on most monitors is 1/72 of an inch, or 1 unit in default user space. The dot size on printers and imagesetters generally ranges from 1/300 of an inch (.24 units) to 1/2400 of an inch (.03 units). For this range of devices, it suffices to store coordinates to two digits to the right of the decimal point. However, because coordinates can be scaled, they should be written using more than two digits, but generally not more than five. Acrobat stores numbers in a fixed format that allows 16 bits for a fraction, which is equivalent to four or five decimal places.

Most monitors and printers cannot produce more than 256 shades of a given color component. Color component values, which are numbers between 0 and 1, should not be written using more than four decimal places.

10.9 Write parameters only when they change

Graphics state operators should be written only when the corresponding graphics state parameters change. Changes to graphics state parameters typically occur both when the application explicitly changes them and when the graphics state is restored using the **Q** operator.

When explicit changes are made to the value of a graphics state parameter, new and old values of the parameter should be compared with the precision with which they will be written, not their internal precision.

A pair of **q** and **Q** operators is commonly used to bracket a sequence of operators that uses a non-default clipping path. The **q** operator saves the default clipping path, and the **Q** operator discards the clipping path when it is no longer needed. However, the **q** and **Q** operators save and restore the entire graphics state, not just the clipping path. To avoid unnecessarily setting all graphics state parameters to achieve a known state after a **Q** operator, an application that produces PDF files may wish to maintain its own graphics state stack, mimicking the PDF graphics state stack. This enables the application to determine the values of all graphics state parameters at all times, and only write operators to change graphics state parameters that do not have the desired value after the **Q** operator.

10.10 Don't draw outside the crop box

Objects entirely outside the crop box do not appear on screen or on the final printout. Nevertheless, if such objects are present in a PDF file, each time the page is drawn, time is spent determining if any portion of them is visible. Simply omit any objects that are entirely outside of the crop box, instead of relying on clipping to keep them from being drawn.

10.11 Consider target device resolution

When producing a PDF file, it is extremely important to consider the device that is the primary target of the document contained in the file. A number of decisions may be made differently depending on whether the document will be primarily viewed on a low-resolution device such as a computer screen or printed to an extremely high-resolution device such as an imagesetter.

If the primary target of the document is a computer screen, users are generally most interested in small file sizes and fast display, and are willing to accept somewhat reduced resolution in exchange for those. If, on the other hand, the primary target is a 1200-dpi imagesetter, file size and drawing time are not as important as obtaining the highest quality possible.

PDF, like the PostScript language, allows graphics objects to be drawn at an arbitrary size and scaled to the desired size. It is often convenient to design objects at a standard size and scale them for a particular situation. Greatly reducing the size of an object, however, can result in unnecessary detail and slow drawing. Choose a level of detail that is appropriate for both monitors and common printer resolutions. In some cases it may be appropriate to replace a complex element of a page with an equivalent image.

Decisions related to the target device primarily affect text, images, and blends. They are discussed further in the following chapters.

10.12 Share resources

Typically, many pages of a document share the same set of fonts. A PDF file will be smaller, display faster, and use less memory if the page's Resources dictionaries refer to the same Font objects. Similarly, if multiple fonts use the same custom encoding, one Encoding object should be shared. The same holds true for ProcSets—if multiple pages require the same combination of ProcSets, they should refer to the same ProcSet array.

10.13 Store common Page attributes in the Pages object

Several Page attributes need not be specified directly in the Page object, but can be inherited from a parent Pages object. Attributes that are the same for all pages in a document may be written once in the root Pages object. If a particular page has a different value, it can directly specify that value and override its parent's value. For example, all pages except one in a document might have the same media box. This value can be stored in the root Pages object, and the media box for the odd-size page can be specified directly in its Page object.

10.14 Use strings for named destinations

PDF 1.2

Named destinations (see [page 185](#)) can be represented with names stored in the Catalog's **Dests** dictionary, or with strings stored in the **Dests** name tree in the Catalog's **Names** dictionary. For a small number of named destinations, either method is acceptable, and for compatibility with PDF 1.0 and 1.1, names are required. But applications that generate many named destinations should use the PDF 1.2 feature of storing these as strings, where there are essentially no implementation limits.

CHAPTER 11

Optimizing Text

Most text optimizations relate to using appropriate operators and taking advantage of the automatic line, character, and word spacing operators supported by PDF. A few optimizations relate to searching.

11.1 Don't produce unnecessary text objects

A PDF viewer application initializes the text environment at the beginning of each text object, and this initialization takes some time. Minimizing the number of text objects used reduces this overhead and reduces file size.

It is not necessary to end one text object and begin another whenever the text matrix is changed using the **Tm** operator. Instead, the text matrix can be changed inside the text object. For example, to create a text object containing several lines of text at various rotations, the following text object could be used:

Example 11.1 *Changing the text matrix inside a text object*

```
BT
/F13 24 Tf
200 100 Td
(Horizontal text) Tj
0.866 0.5 -0.5 0.866 186 150 Tm
(Text rotated 30 degrees counterclockwise) Tj
0.5 0.866 -0.866 0.5 150 186 Tm
(Text rotated 60 degrees counterclockwise) Tj
0 1 -1 0 100 200 Tm
(Text rotated 90 degrees counterclockwise) Tj
ET
```

This sequence draws the text in the font whose name in the current Resources dictionary is **F13**, at a size of 24 points. Keep in mind that the matrix specified using the **Tm** operator replaces the text matrix; it is not concatenated onto the text matrix.

Similarly, font and most other graphics state parameters can change inside a text object. There is one exception—if one of the clipping text-rendering modes is used, the text object must end before changing the text-rendering mode again.

11.2 Use automatic leading

Several of the text string operators make use of the text leading setting to position the drawing point at the start of the next line of text. This makes generating multiple lines of text easy and compact. Use automatic leading whenever possible. The ' and " operators automatically move to the next line of text, as defined by the leading, and the **T*** operator can be used to manually move to the next line of text. Define leading using either the **TD** or **TL** operators.

*Note Do not use the **TD** or **TL** operator unless you use a text operator that has automatic leading.*

For example, the text object in [Example 11.2](#) can be more efficiently written using automatic leading and the ' operator as in [Example 11.3](#).

Example 11.2 *Multiple lines of text without automatic leading*

```
BT
/F13 12 Tf
200 400 Td
(First line of text) Tj
0 -14 Td
(Second line of text) Tj
0 -14 Td
(Third line of text) Tj
0 -14 Td
(Fourth line of text) Tj
ET
```

Example 11.3 *Multiple lines of text using automatic leading*

```
BT
/F13 12 Tf
200 414 Td
14 TL
(First line of text) '
(Second line of text) '
(Third line of text) '
(Fourth line of text) '
ET
```

Note in [Example 11.3](#) that the initial point has been offset vertically by one line. This is because the ' operator moves to the next line before drawing the text.

If it is not possible to use either the ' or " operators to draw a line of text (for example, because the **TJ** operator is used to adjust spacing between particular characters within the line), you can still use the **T*** operator, which advances the point to the beginning of the next line, using the current leading. For example, the text object in [Example 11.4](#) can be more efficiently written using automatic leading and the **T*** operator, as in [Example 11.5](#).

Example 11.4 *TJ operator without automatic leading*

```
BT
/F13 12 Tf
200 700 Td
[(First line) 100 ( of text)] TJ
0 -14 Td
[(Second line) 50 ( of text)] TJ
0 -14 Td
[(Third line) 40 ( of text)] TJ
0 -14 Td
[(Fourth line) 50 ( of text)] TJ
ET
```

Example 11.5 *Use of the T* operator*

```
BT
/F13 12 Tf
200 700 Td
14 TL
[(First line) 100 ( of text)] TJ
T*
[(Second line) 50 ( of text)] TJ
T*
[(Third line) 40 ( of text)] TJ
T*
[(Fourth line) 50 ( of text)] TJ
ET
```

Finally, you can set the leading in either of two ways. The **TL** operator sets the leading directly, while the **TD** operator sets the leading as a side effect of moving the line start position. The methods shown in [Example 11.6](#) and [Example 11.7](#) are equivalent.

Example 11.6 *Using the TL operator to set leading*

```
BT
/F13 12 Tf
200 500 Td
14 TL
[(First line) 100 ( of text)] TJ
T*
[(Second line) 50 ( of text)] TJ
T*
[(Third line) 40 ( of text)] TJ
T*
[(Fourth line) 50 ( of text)] TJ
ET
```

Example 11.7 *Using the TD operator to set leading*

```

BT
/F13 12 Tf
200 500 Td
[(First line) 100 ( of text)] TJ
0 -14 TD
[(Second line) 50 ( of text)] TJ
T*
[(Third line) 40 ( of text)] TJ
T*
[(Fourth line) 50 ( of text)] TJ
ET

```

When using the **TD** operator to set the leading, keep in mind that any horizontal component supplied as an operand to **TD** affects the movement of the drawing point, but not the leading. As a result, the commands

```
0 -14 TD
```

and

```
10 -14 TD
```

both set the leading to 14, although in the latter case the drawing point is ten units to the right of where it is in the former case.

11.3 Take advantage of text spacing operators

The **Tc** and **Tw** operators adjust the spacing between characters and the spacing between words, respectively. Use these operators to make constant adjustments on one or more lines of text. [Example 11.8](#) shows a text object in which one half unit of space has been added between characters on a line and two units between words.

Example 11.8 *Character and word spacing using the Tc and Tw operators*

```

BT
/F13 12 Tf
200 514 Td
14 TL
.5 Tc
2 Tw
(Line of text) '
(Line of text) '
ET

```

Equivalently, the same two lines of text could be produced using the " operator instead of the **Tc**, **Tw**, and ' operators, as shown in [Example 11.9](#).

Example 11.9 *Character and word spacing using the " operator*

```
BT
/F13 12 Tf
200 514 Td
14 TL
2 .5 (Line of text) "
(Line of text) '
ET
```

Using the " operator is preferable if entire lines of text are being written, because it is more compact. If more than one text string operator is used to produce a line of text, the " operator can be used to position the first string of the line and **Tj** or **TJ** for subsequent strings. Remember that the " operator changes the character and word spacing settings for subsequent **Tj**, **TJ**, and ' operators.

11.4 Don't replace spaces between words

When deciding how to represent a line of text in a PDF file, keep in mind that text can be searched. In order to search text accurately, breaks between words must be found. For this reason, it is best to leave spaces in strings, instead of replacing them with an operator that moves the drawing point.

For example, text containing three words could be drawn by:

```
(A few words) Tj
```

Or, replacing the spaces between words with movements of the drawing point:

```
[(A) -300 (few) -300 (words)] TJ
```

The first method is preferred.

11.5 Use the appropriate operator to draw text

In most cases, a line of text can be represented in several ways. When deciding among the various methods, try to draw the line using as few operations as possible. [Table 11.1](#) provides guidelines for selecting the appropriate text string operator.

Table 11.1 *Comparison of text string operators*

<i>Use...</i>	<i>When...</i>
'	Complete line of text can be drawn together No need for individual character spacings

"	Complete line of text can be drawn together Non-zero character or word spacings on each line No need for individual character spacings
Tj	Multiple text operators per line of text No need for individual character spacings
TJ	Individual character spacings needed

When laying out a line of text with non-default character spacings, such as kerned text, use the **TJ** operator rather than a series of pairs of **Tj** and **Td** operators. For example, both of the following lines produce the same output for the Helvetica Bold Oblique font at a size of 12 points:

```
(A f) Tj 15.64 0 Td (ew w) Tj 28.08 0 Td (ords) Tj
[(A f) 30 (ew w) 50 (ords)] TJ
```

The second method is preferred because it minimizes the size of the file and the number of text operators.

11.6 Use the appropriate operator to position text

The **TD**, **Td**, **Tm**, and **T*** operators each change the location at which subsequent text is drawn. Use each of these operators under different circumstances. [Table 11.2](#) provides guidelines for selecting the appropriate text positioning operator.

Table 11.2 *Comparison of text positioning operators*

<i>Use...</i>	<i>When...</i>
Td	Changing only the text location
TD	Changing text location and leading
Tm	Rotating, scaling, or skewing text
T*	Moving to start of next line of text, as defined by the leading

11.7 Remove text clipping

After text has been used as a clipping path through one of the clipping text-rendering modes (4–7), the original clipping path must be restored. Restoration of the original clipping path is accomplished using the **q** and **Q** operators to save and subsequently restore the clipping path, respectively.

Neither **q** nor **Q** may appear inside a text object. Save the original clipping path using the **q** operator before beginning the text object in which a new clipping path is set. When you want to restore the original clipping path, the text object must be

ended using the **ET** operator. Then, use the **Q** operator to restore the original clipping path. Following this, another text object can be entered if more text is to be drawn.

Example 11.10 illustrates the proper way to save and restore a clipping path when using one of the clipping text-rendering modes.

Example 11.10 *Restoring clipping path after using text as clipping path*

```


q
BT
/F13 48 Tf
200 414 Td
Set clip path

0.25 w
5 Tr
(Clip) Tj
ET
BT
200 450 Td
/F13 6 Tf
0 Tr
6 TL
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
(ClipClipClipClipClipClipClipClip) '
ET
Q

BT
/F13 12 Tf
175 395 Td
(Default Clipping Restored) Tj
ET

```

Figure 11.1 shows the output produced by this example when F13 is Helvetica Bold Oblique. The presence of the words “Default Clipping Restored” at the bottom of the figure demonstrates that the clipping path has been restored to its previous value.

Figure 11.1 *Restoring clipping path after clipping to text*

11.8 Consider target device resolution

Although text in a PDF file is resolution-independent (unless a document contains bitmapped Type 3 fonts), there are still reasons to consider the resolution of the target device. Text positioning, in particular, may depend on the primary target device.

It is possible to individually position each character in a string using, for example, the **TJ** operator. This allows precise layout of text. However, adjusting the location of each character increases the size of a PDF file because the positioning must be specified by numbers that are otherwise not needed. In addition, drawing text is slower when each character is individually positioned. As mentioned in [Section 10.11](#), “[Consider target device resolution](#),” if the primary target is a low-resolution device such as a computer screen, producing a small file and one that draws quickly is generally more important than having extremely precise positioning. If the primary target is an imagesetter, extremely precise positioning is generally the primary concern.

As an example of the choices that can be made, suppose the positions of each character on a 60-character line are adjusted from their normal positions by an amount corresponding to 0.01 pixels on a 72 pixel per inch computer screen. The total adjustment across the entire line is just over half a pixel on the screen. If the document is primarily intended to be viewed on a computer screen, omitting the adjustments would make sense because such a small adjustment is invisible. The result would be a smaller file that can be drawn more quickly. On the other hand, the same adjustment corresponds to 10 pixels on a 1200 pixel per inch imagesetter. If the primary target is such an imagesetter, it may be worthwhile retaining the individual position adjustment. Note that precise text positioning is most important for justified text, where positioning errors are easily detected by users.

CHAPTER 12

Optimizing Graphics

12.1 Use the appropriate color-setting operator

Use `0 g` to set the fill color to black, rather than the equivalent, but longer, `0 0 0 rg` or `0 0 0 1 k`. Similarly, `0 G` should be used to set the stroke color to black instead of `0 0 0 RG` or `0 0 0 1 K`. In general, if a color contains equal color components, use either `g` or `G`, as appropriate. For example, use `.8 G` instead of `.8 .8 .8 RG`.

12.2 Defer path painting until necessary

When representing graphics in a PDF file, each path segment can be treated as a separate path, or a number of segments can be grouped together into a single path. Wherever possible, group segments together into a single path. This reduces the size of the file and improves drawing speed. However, a path should not contain more than approximately 1500 segments. For further information, see Appendix B of the *PostScript Language Reference Manual, Third Edition* [1].

Because a path can only be filled with a single color and stroked with a single color, line width, miter limit, and line cap style, a new path must be started whenever one or more of these values is changed.

As an illustration, [Example 12.1](#) and [Example 12.2](#) produce identical output, but the technique shown in [Example 12.2](#) is preferred. Note that [Example 12.2](#) still contains two paths. These paths cannot be combined, because they have different stroke colors.

Example 12.1 *Each path segment as a separate path*

```
.5 0 1 RG
100 100 m
100 200 l
S
100 200 m
200 200 l
S
200 200 m
200 100 l
S
200 100 m
100 100 l
S
0 .2 .4 RG
300 300 m
400 300 l
S
```

Example 12.2 *Grouping path segments into a single path*

```
.5 0 1 RG
100 100 m
100 200 l
200 200 l
200 100 l
S
0 .2 .4 RG
300 300 m
400 300 l
S
```

12.3 Take advantage of the closepath operator

The **h** (**closepath**) operator closes the current subpath by drawing a straight segment from the endpoint of the last segment drawn to the first point in the subpath. When the last segment in a path is straight, use the **h** operator to draw the final segment and close the path.

Two inefficient ways of closing a path commonly occur. The first, shown in [Example 12.3](#), uses the **l** operator to draw the final segment, followed by the **h** operator to close the path.

Example 12.3 *Using redundant **l** and **h** operators to close a path inefficiently*

```
100 100 m
100 200 l
200 200 l
200 100 l
```



```
100 100 1
h
```

The second, shown in [Example 12.4](#), uses the **l** operator to draw the final segment of the path.

Example 12.4 *Using the l operator to close a path inefficiently*

```
100 100 m
100 200 1
200 200 1
200 100 1
100 100 1
```

[Example 12.5](#) shows the correct way of closing a path with a straight segment, using the **h** operator.

Example 12.5 *Taking advantage of the h operator to close a path*

```
100 100 m
100 200 1
200 200 1
200 100 1
h
```

If the **h** operator is not used, the appropriate line join will not occur at the juncture of the path's initial and final point.

12.4 Don't close a path more than once

Close a path only once. Don't use the **h** operator before a path painting operator that implicitly closes the path: the **n**, **b**, **f**, **f*** and **s** operators. In addition, the **h** operator should not be used with the **re** operator, because the **re** operator produces a path that is already closed.

For example, do not use a sequence as in [Example 12.6](#), because the **s** operator automatically closes the path before stroking it.

Example 12.6 *Improperly closing a path: multiple path closing operators*

```
150 240.7 m
253.2 200 1
180.4 150 1
75.4 134.5 1
h
s
```

Instead, use the sequence:

Example 12.7 *Properly closing a path: single path closing operator*

```
150 240.7 m
253.2 200 l
180.4 150 l
75.4 134.5 l
s
```

12.5 Don't draw zero-length lines

When generating graphics from a computer program, it is not uncommon to produce line segments of zero length. Such line segments produce no useful output and should be eliminated before the PDF file is written.

Line segments of zero length may arise when straight line segments are used to approximate a curve. Generally, the programmer wants to make sure that the approximation is close to the actual curve, and so takes small steps in approximating the curve. Occasionally the steps are small enough that they produce segments of zero length after the coordinates have been converted to the format in which they are written to the file. (See [Section 10.8, “Limit the precision of real numbers.”](#))

Zero-length line segments may also be generated when making a two-dimensional projection of a three-dimensional object. Lines in the three-dimensional object that go directly into the page have zero length in the two-dimensional projection.

Strictly speaking, a zero-length line with a round end-cap draws a circle whose diameter is the current linewidth, which may be useful to do in certain circumstances, but is usually not what you want if you're drawing line segments to approximate a curve.

12.6 Make sure drawing is needed

When generating graphics from a computer program, test before writing the graphics to a PDF file to ensure that the graphics actually make new marks on the page and do not simply draw over marks already made.

Redundant graphics typically arise when making a two-dimensional projection of a three-dimensional object. It is possible to end up with several images that lie on top of one another after being projected.

12.7 Take advantage of rectangle and curve operators

Use the **re** operator to draw a rectangle, instead of the corresponding sequence of **m** and **l** operators.

Curves can be drawn in one of two ways; either by approximating the curve with a sequence of straight segments or by using the curve operators present in PDF. Although approximating curves using straight segments is easy, it typically results in a very large amount of data. Use the curve operators (**c**, **v**, **y**) to represent curves in PDF files. Doing so results in a smaller file that can be rendered more quickly.

An algorithm for automatically fitting an arbitrary set of points with a cubic Bézier curve, like those used by PDF, can be found in the series of books called *Graphics Gems*. The algorithm described in *Graphics Gems* (see [16], [28], and [31] in the Bibliography) begins by assuming the points supplied can be fit by a single cubic Bézier curve, with the two endpoints of the Bézier curve being the first and last data points, and the Bézier control points calculated from the approximate tangents at the endpoints of the supplied data. The algorithm minimizes the sum of the squares of the distances between the data points and the curve being fit by moving the control points. If a satisfactory fit cannot be obtained, the data points are separated into two groups at the point with the greatest distance between the curve being fit and the actual data point, and two separate Bézier curves are fit to the two sets of points. This fitting and splitting is repeated until a satisfactory fit is obtained.

12.8 Coalesce operations

Graphics generated by a computer program occasionally contain a group of operations that can be combined into a single operation. These can arise, for example, when a curve is approximated by a series of short straight segments. Significant sections of the curve being approximated may be effectively straight, but the approximation program typically does not realize this and continues to approximate the curve as a sequence of small line segments, instead of combining the collinear segments.

For example, the sequence shown in [Example 12.8](#) contains a number of segments that should be combined. Specifically, the first four **l** operators simply draw one straight line segment and should be combined.

Example 12.8 *Portion of a path before coalescing operations*

```
100 100 m
100 101 l
100 102 l
100 103 l
100 104 l
101 105 l
```

The entire sequence can be replaced by the equivalent and more efficient sequence in [Example 12.9](#).

Example 12.9 *Portion of a path after coalescing operations*

```
100 100 m
100 104 l
101 105 l
```


CHAPTER 13**Optimizing Images**

Sampled images typically require more memory and take more time to process and draw than any other graphics object element of a page. By carefully choosing an appropriate resolution, number of bits per color component, and compression filter, it is possible to significantly enhance image performance.

13.1 Preprocess images

PDF provides operators that transform and clip images. These operators should be used with care. For example, performance often improves if rotation and skewing of an image is performed *before* the image is placed in the PDF file, rather than *by* the PDF viewer application. Similarly, if an image is clipped, it is best to reduce the image to the smallest dimensions possible before placing the image in the PDF file, perhaps eliminating the need for clipping.

13.2 Match image resolution to target device resolution

If a grayscale or color image will primarily be viewed on computer screens (which typically have resolutions between 70–100 pixels per inch) or printed on typical color and monochrome printers (which have resolutions of 300 dpi and default halftone screens of approximately 60 lines per inch), there is no point in producing the image at 300 samples per inch. Most of the information in the higher resolution image will never be seen, the image will contain at least nine times as much data as it needs to (90,000 samples per square inch versus a maximum of 10,000 samples per square inch), and will draw more slowly.

Monochrome images can be stored at higher resolutions of 200 to 300 dpi. This resolution can be achieved on typical printers.

13.3 Use the minimum number of bits per color component

The amount of data needed to represent an image increases as the number of bits per color component increases. This is very important to consider when deciding how many bits per component to use for an image.

If an image requires continuous colors, it might very well need to use 8 bits per color component. However, many graphs, plots, and other types of drawings do not require continuous tone reproduction and are completely satisfactory with a small number of bits per color component.

13.4 Take advantage of indexed color spaces

If an image contains a relatively small number of colors, indexed color spaces can be used to reduce the amount of data needed to represent the image. In an indexed color space, the number of bits needed to represent each sample in an image is determined by the total number of colors in the image rather than by the precision needed to specify a single color.

Most computers currently have displays that support a limited number of colors. For example, it is very common for color displays on the Macintosh computer to provide no more than 256 colors, and many computers running the Microsoft Windows environment provide only 16 colors. On such devices, little loss of image quality will occur if 24-bit color images are replaced by 8-bit indexed color images.

As an example of the compression possible using indexed color spaces, suppose an image contains 256 different colors. Each pixel's color can then be encoded using only 8 bits, regardless of whether the colors in the image are 8-bit grayscale, 24-bit RGB, or 32-bit CMYK. If the colors are 24-bit RGB, using an indexed color space instead of the RGB values would reduce the amount of data needed to represent the image by approximately a factor of three: 24 bits per pixel using an RGB color space versus 8 bits per pixel using an indexed color space. The reduction is not exactly three because the use of an indexed color space requires that a lookup table, containing the list of colors used in the image, be written to the file. For a large image, the size of this lookup table is insignificant compared to the image and can be ignored. For a small image, the size of the lookup table must be included in the calculation. The size of the lookup table can be calculated from the description of indexed color spaces in [Section 7.12.10 on page 240](#).

13.5 Use the DeviceGray color space for monochrome images

For a bitmap (monochrome) image, use the **DeviceGray** color space instead of **DeviceRGB**, **DeviceCMYK**, or **Indexed** color space. In addition, the **BitsPerComponent** attribute for bitmap images should be 1. These settings significantly reduce the amount of data used to represent the image.

Using a different color space or a larger **BitsPerComponent** greatly increases the amount of image data. As an extreme example, a bitmap image represented using a **DeviceCMYK** color space with 8 bits per component contains 32 times as much data as necessary: four color components with 8 bits per component, instead of a single color component with 1 bit per component.

13.6 Use in-line images appropriately

In-line images occupy less disk space and memory than image resources. However, image resources give viewer applications more flexibility in managing memory—the data of an image resource can be read on demand, while an in-line image must be kept in memory together with the rest of a page’s contents.

Implementation note *PDF Writer and the Acrobat Distiller application represent images with less than 4K of data as in-line images until a total of 32K of in-line data are present on a page. Once this limit is reached, subsequent images on that page are represented in-line only if they are smaller than 1K.*

13.7 Don’t compress in-line images unnecessarily

In-line images should not always be compressed and converted to ASCII. Specifically, in-line images should not be compressed if the Contents stream of the page on which the in-line image appears is itself compressed.

Because an in-line image is located completely within the Contents stream of the page, it is automatically passed through the compression and ASCII conversion filters specified for the page’s Contents stream. The specification of an additional compression or ASCII conversion filter in the in-line image itself under these circumstances results in the in-line image being compressed and converted to ASCII twice. This does not result in additional compression and slows down the decoding of the image.

13.8 Choose the appropriate filters

The selection of filters for image streams can be confusing, although a few relatively simple rules can greatly simplify the task.

The order of filters specified when data is decoded must be the opposite of the order in which the filters were applied when the data was encoded. For example, if data is encoded first using LZW and then by ASCII base-85, during decoding the ASCII base-85 filter must be used before the LZW decoding filter. In a stream object, the decoding filters and the order in which they are applied are specified by the **Filter** key. The example would appear as:

```
/Filter [/ASCII85Decode /LZWDecode]
```


As discussed in [Section 2.3.2 on page 22](#), it may be necessary to ensure that a PDF file contains only 7-bit ASCII characters, so the binary data in the PDF file must be encoded with one of the two binary-to-ASCII conversion filters supported by PDF: ASCII hexadecimal and ASCII base-85. Between these two, the ASCII base-85 encoding, which is decoded by the **ASCII85Decode** filter, is preferred because it produces a much smaller expansion in the amount of data than ASCII hexadecimal encoding does.

PDF supports several compression filters that reduce the size of data written into a PDF file. The compression filters can be broken down into two classes: lossless and lossy. A lossless filter is one in which the process of encoding and decoding results in no loss of information: the decoded data is bit-by-bit identical to the original data. For a lossy filter, the process of encoding and decoding results in a loss of information: the decoded data is not bit-by-bit identical to the original data. Lossy filters can be used when the resulting loss of information is not visually significant. The JPEG filter supported by PDF is a lossy filter.

JPEG encoding, which can be decoded by the **DCTDecode** filter, provides very significant compression of color and grayscale images, but because it is a lossy compression it is not appropriate in all circumstances. Screenshots, in particular, are often unacceptable when JPEG encoded. This happens because each pixel in a screenshot is usually significant, and the loss or alteration of just a few pixels can drastically alter the appearance of the screenshot.

Figure 13.1 shows the effect of JPEG encoding on screenshots. The images shown in the figure are magnified by a factor of two to show the changes due to the compression. The 8×8 pixel blocks used in JPEG encoding appear clearly in the pattern inside the icon encoded using a high JPEG compression. The definitions of high, medium, and low JPEG compression are those used by the Acrobat Distiller program. The amount of data in the image from which the figure is taken is: 153,078 bytes with no JPEG encoding, 28,396 bytes with low compression JPEG encoding, 16,944 bytes with medium compression JPEG encoding, and 10,679 bytes with high compression JPEG encoding. All these sizes are for the data after it has been ASCII base-85 encoded.

Figure 13.1 *Effect of JPEG encoding on a screenshot*



Unlike screenshots, the effect of JPEG encoding on continuous-tone images is typically acceptable, particularly when high compression is not demanded. Figure 13.2 illustrates the effect. The image shown in the figure has been magnified by a

factor of two to make the effect of JPEG encoding more readily observable. The version obtained using high compression clearly shows the 8×8 pixel blocks used in JPEG encoding. As in the previous example, the definitions of high, medium, and low JPEG compression are those used by the Acrobat Distiller program, and the sizes shown are for the data after it has been ASCII base-85 encoded.

Figure 13.2 *Effect of JPEG encoding on a continuous-tone image*



No JPEG compression; 20,707 bytes *Low JPEG compression; 7,717 bytes*



Medium JPEG compression; 3,470 bytes

High JPEG compression; 1,997 bytes

In addition to JPEG, PDF supports several lossless compression filters that may be used for images. Guidelines for selecting among them are summarized in [Table 13.1](#).

Table 13.1 *Comparison of compression filters for images*

<i>Use...</i>	<i>When...</i>
DCTDecode	Image is grayscale or color Decompressed image doesn't need to be bit-by-bit identical to original image

CCITTFaxDecode	Image is monochrome (bitmap) Group 4 encoding should be used unless the application generating the file does not support Group 4 encoding
RunLengthDecode	Image contains many groups of identical bytes, such as an 8-bit grayscale image with many areas of same gray level. Should rarely be used
LZWDecode	Images that cannot use DCTDecode and that do not compress well using either CCITT or run length encoding

13.9 Use predefined spot functions

PDF 1.2

A Type-1 halftone screen (see [page 276](#)) includes a spot function. In PDF, this is represented with a Function resource; in PDF 1.2, there are *named* spot functions and *sampled* spot functions. The named spot functions are predefined, so an application can implement them directly in code. When printing a PDF file to a PostScript printer or a Postscript language file, an application can insert the predefined PostScript code for each named function.

Spot functions that do not correspond to any of the predefined set are sampled, and linear interpolation is used to approximate the function. For most functions, this is quite adequate, and the error introduced by sampling and interpolating is quite small.

Spot functions, however, are very sensitive to small errors. They are not used for their values, but rather for their *relative* values. In a 5-by-5 halftone cell, for example, the coordinates of each cell are passed to the spot function. The cell with the smallest spot value is darkened first, and the cell with the greatest value is darkened last. If the values for two cells are interchanged because of an error in the approximation, no matter how small, then they darken in reverse order. This can produce visible differences in the spot. This is an issue only for high-resolution, professional publishing applications.

Sampled functions require interpolation, which takes more time than executing a function directly. The sample values themselves consume space in the PDF file and the PostScript file that is produced by the viewer. The PostScript file will also include an interpolation function.

PDF 1.2 predefines several spot functions, intended to cover all of the common cases of spot functions found in desktop publishing and graphics applications. An application that uses a spot function that is very close to one of these may benefit from using the predefined function instead.

Implementation note

The Acrobat Distiller samples spot functions over a 33-by-33 grid, producing about 1K of data (before compression). The Distiller attempts to match the PostScript language source code of a spot function with the code for the predefined functions. If it fails to find a match with the source code, it samples the function

March 11, 1999

and compares the results with sampled values of the predefined functions, so that a spot function that computes the same results as, say, `InvertedDoubleDot`, but uses different code to do so, will be treated as if the source code had matched.

CHAPTER 14

Clipping

Clipping restricts the areas on a page where marks can be made. It is similar to using a stencil when painting or airbrushing. A stencil with one or more holes in it is placed on a page. As long as the stencil remains in place, paint only reaches the page through the holes in the stencil. After the stencil is removed, paint can again be applied anywhere on the page. More than one stencil may be used in the production of a single page, and if a second stencil is added before the first one is removed, paint only reaches the portion of the page where there are holes in *both* stencils.

Similarly, in producing a PDF page, one or more clipping paths may be used. If a clipping path is not removed before a second clipping path is applied, the resulting clipping path is the intersection of the two paths.

Clipping paths may be specified in two distinct ways: paths and text. These provide clipping that affects all subsequent marking operations until the clipping path is explicitly changed. An example of each type of clipping is provided in the following sections.


Note Whenever a clipping path is no longer needed, the default clipping path should be restored, as described in [Section 10.4](#), “Remove unnecessary clipping paths.”

Image masks do not provide clipping as paths and text do, but they can be thought of as specifying a bitmap clipping template that is placed on the page, painted with a color, and then immediately removed. The differences between images and image masks are discussed below.

14.1 Clipping to a path

As described in [Section 8.6.3](#), “Path clipping operators,” the **W** and **W*** operators can make any path a clipping path. To do this, insert the operator between the path segment operators and one of the path painting operators described in [Section 8.6.2](#), “Path painting operators.”

[Figure 14.1](#) shows the effect of clipping to a region in the shape of a four-pointed star. In the figure, the graphics are shown with and without the star as a clipping path. To draw the figure, the star is first stroked and set to be the current clipping path. A series of lines is then drawn through the star, and the points of the star are filled using arcs.

Figure 14.1 *Clipping to a path**Without clipping to star**With clipping to star*

Note When a path is stroked and used as the current clipping path, remember that the stroke extends half the line width on each side of the path, while subsequent drawing is clipped to the path itself. Because of this, subsequent clipped drawing operations can draw over the “inner half” of the stroke.

The PDF operations needed to produce this output are shown in [Example 14.1](#). The star is first drawn using a series of **I** operators. It is set to be a clipping path using the **W** operator and stroked using the **s** operator. Next, a series of lines is drawn across the star using the **m** and **l** operators. The lines have different gray levels (set by the **G** operator) and line widths (set by the **w** operator). Because each line has a different width and color, each must be stroked (using the **S** operator) individually. To generate the non-clipped portion of the figure, the only change made to the PDF files was to remove the **W** operator.

Example 14.1 *Clipping to a path*

```

391 392 m                               Draw the outline of a star
370 450 l
311 472 l
370 494 l
391 552 l
412 494 l
471 472 l
412 450 l
W
s
.6 G 2 w 311 502 m 471 502 l S           Draw lines
.5 G 3 w 311 492 m 471 492 l S
.4 G 4 w 311 482 m 471 482 l S
.3 G 5 w 311 472 m 471 472 l S
.4 G 4 w 311 462 m 471 462 l S
.5 G 3 w 311 452 m 471 452 l S
.6 G 2 w 311 442 m 471 442 l S
0.6 g                                     Draw and fill circles on the endpoints
340 443 m

```


357 460 357 486 341 502 c
 311 472 l
 f
 421 422 m
 405 438 379 438 362 421 c
 391 392 l
 f
 442 501 m
 425 484 425 458 441 442 c
 471 472 l
 f
 361 522 m
 377 506 403 506 420 523 c
 391 552 l
 f

14.2 Clipping to text

Several of the text rendering modes described in [Section 8.7.1.7, “Text rendering mode”](#) allow text to be used as a clipping path. In particular, modes 4 through 7 can be used to clip subsequent drawing to the shapes of one or more characters.

[Figure 14.2](#) shows the word “and” used as a clipping path. The word is first drawn as stroked and clipped text. Following this, a series of lines containing various ampersands is drawn on top of the word. Only those ampersands contained inside the clipping path defined by the word are visible. The font used for the word “and” is Poetica® Chancery III. The font used for the ampersands is Poetica Ampersands.

Figure 14.2 *Using text as a clipping path*



[Example 14.2](#) shows the page description used to produce [Figure 14.2](#). In the example, the font named **F6** is Poetica Ampersands and the font named **F24** is Poetica Chancery III.

Example 14.2 *Using text as a clipping path*

BT

100 500 Td

Draw the word "and", stroke it, and use it as a clipping path.

/F24 144 Tf

0.25 w

5 Tr

(and) Tj

ET

BT

/F6 6 Tf

Select the Poetica Ampersands font

100 615 Td

0 Tr

6 TL

Draw lines of ampersands

(aAbBcCdDeEfFgGhHiIjJkKlLmMnNoOpPqQrRsStTuU) '
(vVwWxXyYzZ123456aAbBcCdDeEfFgGhHiIjJkKlLmMnNoO) '
(pPqQrRsStTuUvVwWxXyYzZ123456aAbBcCdDeEfFgGhH) '
(jJkKlLmMnNoOpPqQrRsStTuUvVwWxXyYzZ123456) '
(aAbBcCdDeEfFgGhHiIjJkKlLmMnNoOpPqQrRsStTuU) '
(vVwWxXyYzZ123456aAbBcCdDeEfFgGhHiIjJkKlLmMnNoO) '
(pPqQrRsStTuUvVwWxXyYzZ123456aAbBcCdDeEfFgGhH) '
(jJkKlLmMnNoOpPqQrRsStTuUvVwWxXyYzZ123456) '
(aAbBcCdDeEfFgGhHiIjJkKlLmMnNoOpPqQrRsStTuU) '
(vVwWxXyYzZ123456aAbBcCdDeEfFgGhHiIjJkKlLmMnNoO) '
(pPqQrRsStTuUvVwWxXyYzZ123456aAbBcCdDeEfFgGhH) '
(jJkKlLmMnNoOpPqQrRsStTuUvVwWxXyYzZ123456) '
(aAbBcCdDeEfFgGhHiIjJkKlLmMnNoOpPqQrRsStTuU) '
(vVwWxXyYzZ123456aAbBcCdDeEfFgGhHiIjJkKlLmMnNoO) '
(pPqQrRsStTuUvVwWxXyYzZ123456aAbBcCdDeEfFgGhH) '
(jJkKlLmMnNoOpPqQrRsStTuUvVwWxXyYzZ123456) '
(aAbBcCdDeEfFgGhHiIjJkKlLmMnNoOpPqQrRsStTuU) '
(vVwWxXyYzZ123456aAbBcCdDeEfFgGhHiIjJkKlLmMnNoO) '
(pPqQrRsStTuUvVwWxXyYzZ123456aAbBcCdDeEfFgGhH) '
(jJkKlLmMnNoOpPqQrRsStTuUvVwWxXyYzZ123456) '

ET

After beginning a text object by using the **BT** operator, the point at which text will be drawn is set using the **Td** operator. Following this, the font (named **F24**) and the size (144 points) are set using the **Tf** operator, the linewidth for the stroke is set to 0.25 units using the **w** operator, and the stroked clipping text rendering mode (mode 5) is selected using the **Tr** operator. The word "and" is then drawn using the **Tj** operator. Next, the text object is ended using the **ET** operator. This is necessary in order to draw text using a different rendering mode. Following this, another text object is started, the ampersand font (named **F6**) and the size (6 points) are set, the position where text will be drawn is moved, the filled text rendering mode (mode 0) is selected, and the line leading is set to 6 points using the **TL** operator. Finally, the ampersands are drawn by a series of ' operators, and the text object ends.

14.3 Image masks


Although image masks do not provide clipping as described above, they can be thought of as operating as follows: a bitmap image defines the clipping path, where 1s and 0s are considered to be holes and masks. The rectangle containing the bitmap is painted with the current fill color. Immediately following this, the bitmap-derived clipping path is removed.

Image masks differ from images in two ways. First, when an image is drawn, all pixels in the rectangle of the image are painted. In an image mask, only the pixels under holes in the mask are painted; all other pixels are left unchanged. Second, the colors in which an image is painted are encoded inside the image itself, while an image mask is painted using the current fill color at the time the image mask is drawn. Because of this, an image mask may appear in different colors each time it is drawn.


As described in [Section 7.13, “XObjects,”](#) the structure of an image mask differs from that of an image in several ways. First, an image mask must have only one bit per color component. Second, an image mask must not contain a color space specification, while an image must. Third, the image mask dictionary must contain the **ImageMask** key with a value of *true*. For both images and image masks, the array specified as the value of the **Decode** key in the image can be used to choose whether bits containing 1s or bits containing 0s are considered to be set.

[Figure 14.3](#) shows examples of images and image masks. The examples also illustrate how the decode array can be used to invert the image.

Figure 14.3 *Images and image masks*



Image



Inverted image




Image mask



Inverted image mask

Example 14.3 shows the relevant sections from the PDF file used to produce the figure. Because the only difference between the PDF files used to draw each of the four examples is in the image object itself, all the drawing operations are common. The **0.6 g** operation appearing just before the image or image mask is drawn has an effect only when the object being drawn is an image mask, not an image. The example shows the operations used to draw the image mask portion of the figure. To produce the image portion of the figure, the line

```
/ImageMask true
```

was replaced with the line

```
/ColorSpace /DeviceGray
```

For the inverted image and inverted image mask, the line

```
/Decode [1 0]
```

was added to the dictionary of the image or image mask.

Example 14.3 *Images and image masks*

```
3 0 obj <<  
  /Type /Page  
  /Parent 4 0 R  
  /MediaBox [53 470 198 616]  
  /Resources <<  
    /XObject << /Im0 60 0 R >>  
    /ProcSet [ /PDF /ImageC ] >>  
  /Contents 23 0 R  
>> endobj
```

```
23 0 obj <<  
  /Length 205  
>> stream  
0.8 g  
126 472 m  
165 472 197 504 197 543 c  
197 582 165 614 126 614 c  
87 614 55 582 55 543 c  
55 504 87 472 126 472 c  
f
```

Draw a circle and fill it

```
q  
100 0 0 100 76 493.2 cm  
.6 g  
/Im0 Do  
Q  
endstream  
endobj
```

Draw image or mask

March 11, 1999

```
60 0 obj <<  
/Type /XObject  
/Subtype /Image  
/Width 24  
/Height 23  
/BitsPerComponent 1  
/Filter /ASCIIHexDecode  
/Length 162  
/ImageMask true  
>> stream  
003b00 002700 002480 0e4940  
114920 14b220 3cb650 75fe88  
17ff8c 175f14 1c07e2 3803c4  
703182 f8edfc b2bbc2 bb6f84  
31bfc2 18ea3c 0e3e00 07fc00  
03f800 1e1800 1ff800>  
endstream  
endobj
```


APPENDIX A**Example PDF Files**

A.1 Minimal PDF file

Although the PDF file shown in this example does not draw anything, it is almost the minimum PDF file possible. It is not strictly the minimum acceptable file because it contains an Outlines object, a Contents object, and a Resources dictionary with a ProcSet resource. These objects were included to make this file useful as a starting point for developing test files. The objects present in this file are listed in [Table A.1](#).

Note When using this file as a starting point for creating other files, remember to update the ProcSet resource as needed (see [Section 7.6](#), “ProcSets”). Also, remember that the cross-reference table entries may need to have a trailing blank (see [Section 5.15](#), “Cross-reference table.”)

Table A.1 *Objects in empty example*

<i>Object number</i>	<i>Object type</i>
1	Catalog
2	Outlines
3	Pages
4	Page
5	Contents
6	ProcSet array

Example A.1 *Minimal PDF file*

```
%PDF-1.0
1 0 obj
<<
/Type /Catalog
/Pages 3 0 R
/Outlines 2 0 R
>>
endobj

2 0 obj
<<
/Type /Outlines
/Count 0
>>
endobj

3 0 obj
<<
/Type /Pages
/Count 1
/Kids [4 0 R]
>>
endobj

4 0 obj
<<
/Type /Page
/Parent 3 0 R
/Resources << /ProcSet 6 0 R >>
/MediaBox [0 0 612 792]
/Contents 5 0 R
>>
endobj

5 0 obj
<< /Length 35 >>
stream
                                     Place page marking operators here
endstream
endobj

6 0 obj
[/PDF]
endobj

xref
0 7
```

```

0000000000 65535 f
0000000009 00000 n
0000000074 00000 n
0000000120 00000 n
0000000179 00000 n
0000000300 00000 n
0000000384 00000 n
trailer
<<
/Size 7
/Root 1 0 R
>>
startxref
408
%%EOF

```

A.2 Simple text string

This PDF file is the classic “Hello World.” It displays a single line of text containing that string. The string is displayed in 24-point Helvetica. Because Helvetica is one of the base 14 fonts, no font descriptor is needed. This example illustrates the use of fonts and several text-related PDF operators. The objects contained in the file are listed in [Table A.2](#).

Table A.2 *Objects in “Hello World” example*

<i>Object number</i>	<i>Object type</i>
1	Catalog
2	Outlines
3	Pages
4	Page
5	Contents
6	ProcSet array
7	Font (Type 1 font)

Example A.2 *PDF file for simple text example*

```

%PDF-1.0
1 0 obj
<<
/Type /Catalog
/Pages 3 0 R
/Outlines 2 0 R
>>
endobj

```

```
2 0 obj
<<
/Type /Outlines
/Count 0
>>
endobj

3 0 obj
<<
/Type /Pages
/Count 1
/Kids [4 0 R]
>>
endobj

4 0 obj
<<
/Type /Page
/Parent 3 0 R
/Resources << /Font << /F1 7 0 R >> /ProcSet 6 0 R
>>
/MediaBox [0 0 612 792]
/Contents 5 0 R
>>
endobj

5 0 obj
<< /Length 44 >>
stream
BT
/F1 24 Tf
100 100 Td (Hello World) Tj
ET
endstream
endobj

6 0 obj
[/PDF /Text]
endobj

7 0 obj
<<
/Type /Font
/Subtype /Type1
/Name /F1
/BaseFont /Helvetica
/Encoding /MacRomanEncoding
>>
```



```

endobj

xref
0 8
0000000000 65535 f
0000000009 00000 n
0000000074 00000 n
0000000120 00000 n
0000000179 00000 n
0000000322 00000 n
0000000415 00000 n
0000000445 00000 n
trailer
<<
/Size 8
/Root 1 0 R
>>
startxref
553
%%EOF

```

A.3 Simple graphics

This PDF file draws a thin black line segment, a thick black dashed line segment, a filled and stroked rectangle, and a filled and stroked Bézier curve. The file contains comments showing the various operations. The objects present in this file are listed in [Table A.3](#).

Note The byte-addresses in the cross-reference table are not necessarily accurate.

Table A.3 Objects in graphics example

Object number	Object type
1	Catalog
2	Outlines
3	Pages
4	Page
5	Contents
6	ProcSets

Example A.3 PDF file for simple graphics example

```

%PDF-1.0
1 0 obj
<<
/Type /Catalog

```

```

/Pages 3 0 R
/Outlines 2 0 R
>>
endobj

2 0 obj
<<
/Type /Outlines
/Count 0
>>
endobj

3 0 obj
<<
/Type /Pages
/Count 1
/Kids [ 4 0 R ]
>>
endobj

4 0 obj
<<
/Type /Page
/Parent 3 0 R
/Resources << /ProcSet 6 0 R >>
/MediaBox [ 0 0 612 792 ]
/Contents 5 0 R
>>
endobj

5 0 obj
<< /Length 612 >>
stream
    Draw a black line segment, using the default line width.
    150 250 m
    150 350 l
    S
    Draw thicker, dashed line segment.
    4 w
    Set a linewidth of 4 points.
    [4 6] 0 d
    Set a dash pattern with 4 units on, 6 units off.
    150 250 m
    400 250 l
    S
    [ ] 0 d
    Reset dash pattern to a solid line
    1 w
    Reset linewidth to 1 unit
    Draw a rectangle, 1 unit red border, filled with light blue.
    .5 .75 1 rg
    Light blue for fill color.

```

March 11, 1999

```
1 0 0 RG
200 300 50 75 re
B
```

Red for stroke color.


*Draw a curve using a Bézier curve,
filled with gray and with a colored border.*

```
.5 .1 .2 RG
0.7 g
300 300 m
300 400 400 400 400 300 c
b
endstream
endobj
```

```
6 0 obj
[
/PDF
]
endobj
```

```
xref
0 7
0000000000 65535 f
0000000016 00000 n
0000000086 00000 n
0000000136 00000 n
0000000200 00000 n
0000000328 00000 n
0000000992 00000 n
trailer
<<
/Size 8
/Root 1 0 R
>>
startxref
1110
%%EOF
```


Figure 14.4 Visual representation of Example A.3



A.4 Pages tree

This example is a fragment of a PDF file, illustrating the structure of the Pages tree for a large document. It contains the Pages objects for a 62-page file. The structure of the Pages tree for this example is shown in Figure A.1. In the figure, the numbers are object numbers corresponding to the objects in the PDF document fragment contained in Example A.4.

Figure A.1 Pages tree for 62-page document example



Example A.4 Pages tree for a document containing 62 pages

```

337 0 obj
<<
/Kids [335 0 R 336 0 R]
/Count 62

```

March 11, 1999

/Type /Pages

>>

endobj

335 0 obj

<<

/Kids [4 0 R 43 0 R 77 0 R 108 0 R 139 0 R 170 0 R]

/Count 36

/Type /Pages

/Parent 337 0 R

>>

endobj

336 0 obj

<<

/Kids [201 0 R 232 0 R 263 0 R 294 0 R 325 0 R]

/Count 26

/Type /Pages

/Parent 337 0 R

>>

endobj

4 0 obj

<<

/Kids [3 0 R 16 0 R 21 0 R 26 0 R 31 0 R 37 0 R]

/Count 6

/Type /Pages

/Parent 335 0 R

>>

endobj

43 0 obj

<<

/Kids [42 0 R 48 0 R 53 0 R 58 0 R 63 0 R 70 0 R]

/Count 6

/Type /Pages

/Parent 335 0 R

>>

endobj

77 0 obj

<<

/Kids [76 0 R 82 0 R 87 0 R 92 0 R 97 0 R 102 0 R]

/Count 6

/Type /Pages

/Parent 335 0 R

>>

endobj

:

March 11, 1999

```
108 0 obj
<<
/Kids [107 0 R 113 0 R 118 0 R 123 0 R 128 0 R 133 0
R]
/Count 6
/Type /Pages
/Parent 335 0 R
>>
endobj
```

```
139 0 obj
<<
/Kids [138 0 R 144 0 R 149 0 R 154 0 R 159 0 R 164 0
R]
/Count 6
/Type /Pages
/Parent 335 0 R
>>
endobj
```

```
170 0 obj
<<
/Kids [169 0 R 175 0 R 180 0 R 185 0 R 190 0 R 195 0
R]
/Count 6
/Type /Pages
/Parent 335 0 R
>>
endobj
```

```
201 0 obj
<<
/Kids [200 0 R 206 0 R 211 0 R 216 0 R 221 0 R 226 0
R]
/Count 6
/Type /Pages
/Parent 336 0 R
>>
endobj
```

```
232 0 obj
<<
/Kids [231 0 R 237 0 R 242 0 R 247 0 R 252 0 R 257 0
R]
/Count 6
/Type /Pages
/Parent 336 0 R
>>
```

```
endobj

263 0 obj
<<
/Kids [262 0 R 268 0 R 273 0 R 278 0 R 283 0 R 288 0
R]
/Count 6
/Type /Pages
/Parent 336 0 R
>>
endobj

294 0 obj
<<
/Kids [293 0 R 299 0 R 304 0 R 309 0 R 314 0 R 319 0
R]
/Count 6
/Type /Pages
/Parent 336 0 R
>>
endobj

325 0 obj
<<
/Kids [324 0 R 330 0 R]
/Count 2
/Type /Pages
/Parent 336 0 R
>>
endobj
```

A.5 Outline

This section from a PDF file illustrates the structure of an outline tree with six entries. [Example A.5](#) shows the outline with all entries open, as illustrated in [Figure A.2](#).

Figure A.2 *Example of outline with six items, all open*

Onscreen appearance	Object number	Count
Document	21	6
Section 1	22	4
Section 2	25	0
Subsection 1	26	1
Subsection 1	27	0
Section 3	28	0
Summary	29	0

Example A.5 *Six entry outline, all items open*

```

21 0 obj
<<
/Count 6
/Type /Outlines
/First 22 0 R
/Last 29 0 R
>>
endobj

22 0 obj
<<
/Parent 21 0 R
/Dest [3 0 R /Top 0 792 0]
/Title (Document)
/Next 29 0 R
/First 25 0 R
/Last 28 0 R
/Count 4
>>
endobj

25 0 obj
<<
/Dest [3 0 R /FitR -38 255 650 792]
/Parent 22 0 R
/Title (Section 1)
/Next 26 0 R
>>
endobj

26 0 obj
<<

```


March 11, 1999

```
/Dest [3 0 R /FitR -38 255 650 792]
/Prev 25 0 R
/Next 28 0 R
/Parent 22 0 R
/Title (Section 2)
/First 27 0 R
/Last 27 0 R
/Count 1
>>
endobj

27 0 obj
<<
/Dest [3 0 R /FitR -38 255 650 792]
/Parent 26 0 R
/Title (Subsection 1)
>>
endobj

28 0 obj
<<
/Dest [3 0 R /FitR 3 255 622 792]
/Prev 26 0 R
/Parent 22 0 R
/Title (Section 3)
>>
endobj

29 0 obj
<<
/Prev 22 0 R
/Parent 21 0 R
/Dest [3 0 R /FitR 3 255 622 792]
/Title (Summary)
>>
endobj
```

[Example A.6](#) is the same as [Example A.5](#), except that one of the outline items has been closed in the display. The outline appears as shown in [Figure A.3](#).

Figure A.3 Example of outline with six items, five of which are open

Onscreen appearance	Object number	Count
Document	21	5
Section 1	22	3
Section 2	25	0
Section 3	26	-1
Section 3	28	0
Summary	29	0

Example A.6 Six entry outline, five entries open

```

21 0 obj
<<
/Count 5
/Type /Outlines
/First 22 0 R
/Last 29 0 R
>>
endobj

22 0 obj
<<
/Parent 21 0 R
/Dest [3 0 R /Top 0 792 0]
/Title (Document)
/Next 29 0 R
/First 25 0 R
/Last 28 0 R
/Count 3
>>
endobj

25 0 obj
<<
/Dest [3 0 R /FitR -38 255 650 792]
/Parent 22 0 R
/Title (Section 1)
/Next 26 0 R
>>
endobj

26 0 obj
<<
/Dest [3 0 R /FitR -38 255 650 792]

```

```
/Prev 25 0 R
/Next 28 0 R
/Parent 22 0 R
/Title (Section 2)
/First 27 0 R
/Last 27 0 R
/Count -1
>>
endobj

27 0 obj
<<
/Dest [3 0 R /FitR -38 255 650 792]
/Parent 26 0 R
/Title (Subsection 1)
>>
endobj

28 0 obj
<<
/Dest [3 0 R /FitR 3 255 622 792]
/Prev 26 0 R
/Parent 22 0 R
/Title (Section 3)
>>
endobj

29 0 obj
<<
/Prev 22 0 R
/Parent 21 0 R
/Dest [3 0 R /FitR 3 255 622 792]
/Title (Summary)
>>
endobj
```

A.6 Updated file

This example shows the structure of a PDF file as it is updated several times; multiple body sections, cross-reference sections, and trailers. In addition, it illustrates the fact that once an object ID has been assigned to an object, it keeps the ID until it is deleted, even if the object is altered. Finally, it illustrates the re-use of cross-reference entries for objects that have been deleted, along with the incrementing of the generation number after an object has been deleted.

The original file is that used in [Section A.1, “Minimal PDF file.”](#) This file is not shown again here. First, four text annotations are added and the file saved. Next, the text of one of the annotations is altered, and the file saved. Following this, two of the text annotations are deleted, and the file saved again. Finally, three text annotations are added, and the file saved again.

The segments added to the file at each stage are shown separately. Throughout this example, objects are referred to by their object IDs, made up of the object number and generation number, rather than simply by the object number, as was done in earlier examples. This is necessary because objects are re-used in this example, so that the object number is not a unique identifier.

Note The tables in this section show only the objects that are modified at some point during the updating process. Objects from the example file in [Section A.1, “Minimal PDF file”](#) that are never altered during the update are not shown.

A.6.1 Add four text annotations

Four text annotations were added to the initial file and the file saved. [Table A.4](#) lists the objects in this update.

Table A.4 *Object use after adding four text annotations*

<i>Object ID</i>	<i>Object type</i>
4 0	Page
7 0	Annots array
8 0	Text annotation
9 0	Text annotation
10 0	Text annotation
11 0	Text annotation

[Example A.7](#) shows the lines added to the file by this update. The Page object is updated because an **Annots** key has been added. Note that the file’s trailer now contains a **Prev** key, which points to the original cross-reference section in the file, while the **startxref** value at the end of the file points to the cross-reference section added by the update.

Example A.7 *Update section of PDF file when four text annotations are added*

```

4 0 obj
<<
  /Type /Page
  /Parent 3 0 R
  /Resources << /ProcSet 6 0 R >>
  /MediaBox [0 0 612 792]
  /Contents 5 0 R
  /Annots 7 0 R
>>

```

March 11, 1999

endobj

7 0 obj

[8 0 R 9 0 R 10 0 R 11 0 R]

endobj

8 0 obj

<<

/Type /Annot

/Subtype /Text

/Open true

/Rect [44 616 162 735]

/Contents (Text #1)

>>

endobj

9 0 obj

<<

/Type /Annot

/Subtype /Text

/Open false

/Rect [224 668 457 735]

/Contents (Text #2)

>>

endobj

10 0 obj

<<

/Type /Annot

/Subtype /Text

/Open true

/Rect [239 393 328 622]

/Contents (Text #3)

>>

endobj

11 0 obj

<<

/Type /Annot

/Subtype /Text

/Open false

/Rect [34 398 225 575]

/Contents (Text #4)

>>

endobj

xref

0 1

```

0000000000 65535 f
4 1
0000000612 00000 n
7 5
0000000747 00000 n
0000000792 00000 n
0000000897 00000 n
0000001004 00000 n
0000001111 00000 n
trailer
<<
/Size 12
/Root 1 0 R
/Prev 408
>>
startxref
1218
%%EOF

```

A.6.2 Modify text of one annotation

The lines shown in [Example A.8](#) were added to the file when it was saved after modifying one text annotation. Note that the file now contains two copies of the object with ID 10 0 (the text annotation that was modified), and that the cross-reference section added points to the more recent version of the object. The cross-reference section added contains one subsection. The subsection contains an entry only for the object that was modified. In addition, the **Prev** key in the file's trailer has been updated to point to the cross-reference section added by the previous update, while the **startxref** value at the end of the file points to the newly added cross-reference section.

Example A.8 *Update section of PDF file when one text annotation is modified*

```

10 0 obj
<<
/Type /Annot
/Subtype /Text
/Open true
/Rect [239 393 328 622]
/Contents (Modified Text #3)
>>
endobj

xref
10 1
0000001444 00000 n
trailer
<<
/Size 12

```

```

/Root 1 0 R
/Prev 1218
>>
startxref
1560
%%EOF

```

A.6.3 Delete two annotations

Table A.5 lists the objects updated when two text annotations were deleted and the file saved.

Table A.5 *Object use after deleting two text annotations*

<i>Object ID</i>	<i>Object type</i>
7 0	Annots array
8 0	Free
9 0	Free

The Annots array is the only object that is written in this update. It is updated because it now contains two fewer annotations.

Example A.9 shows the lines added when the file was saved. Note that objects with IDs 8 0 and 9 0 have been deleted, as can be seen from the fact that their entries in the cross-reference section end with an **f**. The cross-reference section added in this step contains four entries, corresponding to object number 0, the Annots array, and the two deleted text annotations. The cross-reference entry for object number 0 is updated because it is the head of the linked list of free objects, and must now point to the newly freed entry for object number 8. The entry for object number 8 points to the entry for object number 9 (the next free entry), while the entry for object number 9 is the last free entry in the cross-reference table, indicated by the fact that it points to object number 0. The entries for the two deleted text annotations are marked as free, and as having generation numbers of 1, which will be used for any objects that re-use these cross-reference entries. Keep in mind that, although the two objects have been deleted, they are still present in the file. It is the cross-reference table that records the fact that they have been deleted.

The **Prev** key in the trailer dictionary has again been updated, so that it points to the cross-reference section added in the previous step, and the **startxref** value points to the newly added cross-reference section.

Example A.9 *Update section of PDF file when two text annotations are deleted*

```

7 0 obj
[10 0 R 11 0 R]
endobj

xref
0 1

```

```

0000000008 65535 f
7 3
0000001658 00000 n
0000000009 00001 f
0000000000 00001 f
trailer
<<
/Size 12
/Root 1 0 R
/Prev 1560
>>
startxref
1691
%%EOF

```

A.6.4 Add three annotations

Finally, three text annotations were added to the file. [Table A.6](#) lists the objects involved in this update.

Table A.6 *Object use after adding three text annotations*

<i>Object ID</i>	<i>Object type</i>
7 0	Annots array
8 1	Text annotation
9 1	Text annotation
12 0	Text annotation

Object numbers 8 and 9, which were the object numbers used for the two annotations deleted in the previous step, have been re-used. The new objects have been given a generation number of 1, however. In addition, the third text annotation added was assigned the previously unused object ID of 12 0.

[Example A.10](#) shows the lines added to the file by this update. The cross-reference section added in this step contains five entries, corresponding object number 0, the Annots array, and the three annotations added. The entry for object number zero is updated because the previously free entries for object numbers 8 and 9 have been re-used. The entry for object number zero now shows that there are no free entries in the cross-reference table. The Annots array is updated to reflect the addition of the three new text annotations.

As in previous updates, the trailer's **Prev** key and **startxref** value have been updated.

The annotation with object ID 12 0 illustrates the splitting of a long text string across multiple lines, as well as the technique for including non-standard characters in a string. In this case, the character is an ellipsis (...), which is character code 203 (octal) in the **PDFDocEncoding** used for text annotations.

Example A.10 *Update section of PDF file after three text annotations are added*

```
7 0 obj
[10 0 R 11 0 R 8 1 R 9 1 R 12 0 R]
endobj

8 1 obj
<<
/Type /Annot
/Subtype /Text
/Open true
/Rect [58 657 172 742]
/Contents (New Text #1)
>>
endobj

9 1 obj
<<
/Type /Annot
/Subtype /Text
/Open false
/Rect [389 459 570 537]
/Contents (New Text Annotation #2)
>>
endobj

12 0 obj
<<
/Type /Annot
/Subtype /Text
/Open true
/Rect [44 253 473 337]
/Contents (A longer annotation which we'll call, \
for lack of a better name\203New T\
ext #3)
>>
endobj

xref
0 1
0000000000 65535 f
7 3
000001853 0000 n
000001905 0001 n
000002014 0001 n
12 1
000002136 0000 n
trailer
<<
```

:

March 11, 1999

```
/Size 13  
/Root 1 0 R  
/Prev 1691  
>>  
startxref  
2315  
%%EOF
```

APPENDIX **B**

Summary of Page Marking Operators

Following is a list of all page marking operators used in PDF files, arranged alphabetically. For each operator, a brief description is given, along with a reference to the page in this document where the operator is discussed in detail. Words shown in boldface in the summary column are PostScript language operators.

Table B.1 *PDF page marking operators*

Operator	Summary	Page
b	closepath , fill , and stroke path	320
B	fill and stroke path	338
b*	closepath , eofill , and stroke path	338
B*	eofill and stroke path	338
BDC	begin marked content, with a dictionary	353
BI	begin image	350
BMC	begin marked content	353
BT	begin text object	345
BX	begin section allowing undefined operators	351
c	curveto	335
cm	concat . Concatenates the matrix to the current transformation matrix.	323
cs	setcolorspace for fill	331
CS	setcolorspace for stroke	332
d	setdash	325
d0	setcharwidth for Type 3 font	351
d1	setcachedevice for Type 3 font	351
Do	execute the named XObject	348

PDF 1.2**PDF 1.2****PDF 1.1****PDF 1.1****PDF 1.1**

Table B.1 PDF page marking operators

Operator	Summary	Page
DP	mark a place in the content stream, with a dictionary	353
EI	end image	350
EMC	end marked content	353
ET	end text object	345
EX	end section that allows undefined operators	351
f	fill path	338
F	fill path	338
f*	eofill path	338
g	setgray (fill)	331
G	setgray (stroke)	331
gs	set parameters in the extended graphics state	327
h	closepath	336
i	setflat	323
ID	begin image data	350
j	setlinejoin	326
J	setlinecap	324
k	setcmykcolor (fill)	331
K	setcmykcolor (stroke)	331
l	lineto	335
m	moveto	335
M	setmiterlimit	327
MP	mark a place in the content stream	353
n	end path without fill or stroke	338
q	save graphics state	323
Q	restore graphics state	323
re	rectangle	336
rg	setrgbcolor (fill)	331
RG	setrgbcolor (stroke)	331

PDF 1.2

PDF 1.2

PDF 1.1

PDF 1.2

PDF 1.2

Table B.1 *PDF page marking operators*

Operator	Summary	Page
ri	set the rendering intent	333
s	closepath and stroke path	338
S	stroke path	338
sc	setcolor (fill)	332
SC	setcolor (stroke)	332
scn	setcolor (fill, in pattern and separation color spaces)	333
SCN	setcolor (stroke, in pattern and separation color spaces)	333
sh	shfill (shaded fill)	338
Tc	set character spacing	340
Td	move text current point	345
TD	move text current point and set leading	345
Tf	set font name and size	342
Tj	show text	347
TJ	show text, allowing individual character positioning	347
TL	set leading	342
Tm	set text matrix	345
Tr	set text rendering mode	343
Ts	set super/subscripting text rise	344
Tw	set word spacing	341
Tz	set horizontal scaling	341
T*	move to start of next line	346
v	curveto	335
w	setlinewidth	326
W	clip	339
W*	eoclip	339
y	curveto	335
'	move to next line and show text	347
"	move to next line and show text	347

PDF 1.1

PDF 1.1

PDF 1.2

PDF 1.2

PDF 1.3

:

March 11, 1999

APPENDIX C

Predefined Font Encodings

PDF provides several predefined font encodings:

- **MacRomanEncoding**, **MacExpertEncoding**, and **WinAnsiEncoding** may be used in Font and Encoding objects.
- **PDFDocEncoding** is the encoding used in outline entries, text annotations, and strings in the Info dictionary.
- **StandardEncoding** is the built-in encoding for many fonts.

This appendix contains three tables describing these encodings. The first table shows all encodings except **MacExpertEncoding** and is arranged alphabetically by character name. The second table is similar, except that it is arranged numerically by character code. The third table shows the encoding for **MacExpertEncoding**, which is shown in a separate table because it has a substantially different character set than the other encodings.

C.1 Predefined encodings sorted by character name

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
A	A	65	101	65	101	65	101	65	101
Æ	AE	225	341	174	256	198	306	198	306
Á	Aacute	—	—	231	347	193	301	193	301
Â	Acircumflex	—	—	229	345	194	302	194	302
Ä	Adieresis	—	—	128	200	196	304	196	304
À	Agrave	—	—	203	313	192	300	192	300
Å	Aring	—	—	129	201	197	305	197	305
Ã	Atilde	—	—	204	314	195	303	195	303
B	B	66	102	66	102	66	102	66	102
C	C	67	103	67	103	67	103	67	103
Ç	Ccedilla	—	—	130	202	199	307	199	307
D	D	68	104	68	104	68	104	68	104
E	E	69	105	69	105	69	105	69	105
É	Eacute	—	—	131	203	201	311	201	311
Ê	Ecircumflex	—	—	230	346	202	312	202	312
Ë	Edieresis	—	—	232	350	203	313	203	313
È	Egrave	—	—	233	351	200	310	200	310
€	Euro	—	—	—	—	128	200	160	240
F	F	70	106	70	106	70	106	70	106
G	G	71	107	71	107	71	107	71	107
H	H	72	110	72	110	72	110	72	110
I	I	73	111	73	111	73	111	73	111
Í	Iacute	—	—	234	352	205	315	205	315
Î	Icircumflex	—	—	235	353	206	316	206	316
Ï	Idieresis	—	—	236	354	207	317	207	317
Ì	Igrave	—	—	237	355	204	314	204	314
J	J	74	112	74	112	74	112	74	112
K	K	75	113	75	113	75	113	75	113
L	L	76	114	76	114	76	114	76	114
	Lslash	232	350	—	—	—	—	149	225
M	M	77	115	77	115	77	115	77	115
N	N	78	116	78	116	78	116	78	116
Ñ	Ntilde	—	—	132	204	209	321	209	321
O	O	79	117	79	117	79	117	79	117
Œ	OE	234	352	206	316	140	214	150	226
Ó	Oacute	—	—	238	356	211	323	211	323
Ô	Ocircumflex	—	—	239	357	212	324	212	324
Ö	Odieresis	—	—	133	205	214	326	214	326
Ò	Ograve	—	—	241	361	210	322	210	322
Ø	Oslash	233	351	175	257	216	330	216	330
Õ	Otilde	—	—	205	315	213	325	213	325
P	P	80	120	80	120	80	120	80	120

PDF 1.3

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
Q	Q	81	121	81	121	81	121	81	121
R	R	82	122	82	122	82	122	82	122
S	S	83	123	83	123	83	123	83	123
Š	Scaron	—	—	—	—	138	212	151	227
T	T	84	124	84	124	84	124	84	124
Þ	Thorn	—	—	—	—	222	336	222	336
U	U	85	125	85	125	85	125	85	125
Ú	Uacute	—	—	242	362	218	332	218	332
Û	Ucircumflex	—	—	243	363	219	333	219	333
Ü	Udieresis	—	—	134	206	220	334	220	334
Û	Ugrave	—	—	244	364	217	331	217	331
V	V	86	126	86	126	86	126	86	126
W	W	87	127	87	127	87	127	87	127
X	X	88	130	88	130	88	130	88	130
Y	Y	89	131	89	131	89	131	89	131
Ý	Yacute	—	—	—	—	221	335	221	335
ÿ	Ydieresis	—	—	217	331	159	237	152	230
Z	Z	90	132	90	132	90	132	90	132
	Zcaron	—	—	—	—	142	216	153	231
a	a	97	141	97	141	97	141	97	141
á	aacute	—	—	135	207	225	341	225	341
â	acircumflex	—	—	137	211	226	342	226	342
´	acute	194	302	171	253	180	264	180	264
ä	adieresis	—	—	138	212	228	344	228	344
æ	ae	241	361	190	276	230	346	230	346
à	agrave	—	—	136	210	224	340	224	340
&	ampersand	38	46	38	46	38	46	38	46
å	aring	—	—	140	214	229	345	229	345
^	asciicircum	94	136	94	136	94	136	94	136
~	asciitilde	126	176	126	176	126	176	126	176
*	asterisk	42	52	42	52	42	52	42	52
@	at	64	100	64	100	64	100	64	100
ã	atilde	—	—	139	213	227	343	227	343
b	b	98	142	98	142	98	142	98	142
\	backslash	92	134	92	134	92	134	92	134
	bar	124	174	124	174	124	174	124	174
{	braceleft	123	173	123	173	123	173	123	173
}	braceright	125	175	125	175	125	175	125	175
[bracketleft	91	133	91	133	91	133	91	133
]	bracketright	93	135	93	135	93	135	93	135
˘	breve	198	306	249	371	—	—	24	30
‡	brokenbar	—	—	—	—	166	246	166	246
•	bullet	183	267	165	245	149	225	128	200
c	c	99	143	99	143	99	143	99	143
ÿ	caron	207	317	255	377	—	—	25	31

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
ç	cedilla	—	—	141	215	231	347	231	347
¸	cedilla	203	313	252	374	184	270	184	270
¢	cent	162	242	162	242	162	242	162	242
^	circumflex	195	303	246	366	136	210	26	32
:	colon	58	72	58	72	58	72	58	72
,	comma	44	54	44	54	44	54	44	54
©	copyright	—	—	169	251	169	251	169	251
¤	currency	168	250	219	333	164	244	164	244
d	d	100	144	100	144	100	144	100	144
†	dagger	178	262	160	240	134	206	129	201
‡	daggerdbl	179	263	224	340	135	207	130	202
°	degree	—	—	161	241	176	260	176	260
¨	dieresis	200	310	172	254	168	250	168	250
÷	divide	—	—	214	326	247	367	247	367
\$	dollar	36	44	36	44	36	44	36	44
•	dotaccent	199	307	250	372	—	—	27	33
◦	dotlessi	245	365	245	365	—	—	154	232
e	e	101	145	101	145	101	145	101	145
é	eacute	—	—	142	216	233	351	233	351
ê	ecircumflex	—	—	144	220	234	352	234	352
ë	edieresis	—	—	145	221	235	353	235	353
è	egrave	—	—	143	217	232	350	232	350
8	eight	56	70	56	70	56	70	56	70
...	ellipsis	188	274	201	311	133	205	131	203
—	emdash	208	320	209	321	151	227	132	204
–	endash	177	261	208	320	150	226	133	205
=	equal	61	75	61	75	61	75	61	75
ð	eth	—	—	—	—	240	360	240	360
!	exclam	33	41	33	41	33	41	33	41
¡	exclamdown	161	241	193	301	161	241	161	241
f	f	102	146	102	146	102	146	102	146
?fi	fi	174	256	222	336	—	—	147	223
5	five	53	65	53	65	53	65	53	65
?fl	fl	175	257	223	337	—	—	148	224
f	florin	166	246	196	304	131	203	134	206
4	four	52	64	52	64	52	64	52	64
/	fraction	164	244	218	332	—	—	135	207
g	g	103	147	103	147	103	147	103	147
ß	germandbls	251	373	167	247	223	337	223	337
`	grave	193	301	96	140	96	140	96	140
>	greater	62	76	62	76	62	76	62	76
«	guillemotleft	171	253	199	307	171	253	171	253
»	guillemotright	187	273	200	310	187	273	187	273
<	guilsinglleft	172	254	220	334	139	213	136	210
>	guilsinglright	173	255	221	335	155	233	137	211

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
h	h	104	150	104	150	104	150	104	150
"	hungarumlaut	205	315	253	375	—	—	28	34
-	hyphen	45	55	45	55	45	55	45	55
i	i	105	151	105	151	105	151	105	151
í	iacute	—	—	146	222	237	355	237	355
î	icircumflex	—	—	148	224	238	356	238	356
ï	idieresis	—	—	149	225	239	357	239	357
ì	igrave	—	—	147	223	236	354	236	354
j	j	106	152	106	152	106	152	106	152
k	k	107	153	107	153	107	153	107	153
l	l	108	154	108	154	108	154	108	154
<	less	60	74	60	74	60	74	60	74
¬	logicalnot	—	—	194	302	172	254	172	254
/	lslash	248	370	—	—	—	—	155	233
m	m	109	155	109	155	109	155	109	155
-	macron	197	305	248	370	175	257	175	257
-	minus	—	—	—	—	—	—	138	212
μ	mu	—	—	181	265	181	265	181	265
³	multiply	—	—	—	—	215	327	215	327
n	n	110	156	110	156	110	156	110	156
9	nine	57	71	57	71	57	71	57	71
ñ	ntilde	—	—	150	226	241	361	241	361
#	numbersign	35	43	35	43	35	43	35	43
o	o	111	157	111	157	111	157	111	157
ó	oacute	—	—	151	227	243	363	243	363
ô	ocircumflex	—	—	153	231	244	364	244	364
ö	odieresis	—	—	154	232	246	366	246	366
œ	oe	250	372	207	317	156	234	156	234
•	ogonek	206	316	254	376	—	—	29	35
ò	ograve	—	—	152	230	242	362	242	362
1	one	49	61	49	61	49	61	49	61
½	onehalf	—	—	—	—	189	275	189	275
¼	onequarter	—	—	—	—	188	274	188	274
¹	onesuperior	—	—	—	—	185	271	185	271
ª	ordfeminine	227	343	187	273	170	252	170	252
º	ordmasculine	235	353	188	274	186	272	186	272
ø	oslash	249	371	191	277	248	370	248	370
õ	otilde	—	—	155	233	245	365	245	365
p	p	112	160	112	160	112	160	112	160
¶	paragraph	182	266	166	246	182	266	182	266
(parenleft	40	50	40	50	40	50	40	50
)	parenright	41	51	41	51	41	51	41	51
%	percent	37	45	37	45	37	45	37	45
.	period	46	56	46	56	46	56	46	56
·	periodcentered	180	264	225	341	183	267	183	267

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
%	perthousand	189	275	228	344	137	211	139	213
+	plus	43	53	43	53	43	53	43	53
±	plusminus	—	—	177	261	177	261	177	261
q	q	113	161	113	161	113	161	113	161
?	question	63	77	63	77	63	77	63	77
¿	questiondown	191	277	192	300	191	277	191	277
"	quotedbl	34	42	34	42	34	42	34	42
„	quotedblbase	185	271	227	343	132	204	140	214
“	quotedblleft	170	252	210	322	147	223	141	215
”	quotedblright	186	272	211	323	148	224	142	216
‘	quoteleft	96	140	212	324	145	221	143	217
’	quoteright	39	47	213	325	146	222	144	220
,	quotesinglbase	184	270	226	342	130	202	145	221
’	quotesingle	169	251	39	47	39	47	39	47
r	r	114	162	114	162	114	162	114	162
®	registered	—	—	168	250	174	256	174	256
°	ring	202	312	251	373	176	260	30	36
s	s	115	163	115	163	115	163	115	163
š	scaron	—	—	—	—	154	232	157	235
§	section	167	247	164	244	167	247	167	247
;	semicolon	59	73	59	73	59	73	59	73
7	seven	55	67	55	67	55	67	55	67
6	six	54	66	54	66	54	66	54	66
/	slash	47	57	47	57	47	57	47	57
	space	32	40	32,202	40,312	32	40	32	40
£	sterling	163	243	163	243	163	243	163	243
t	t	116	164	116	164	116	164	116	164
þ	thorn	—	—	—	—	254	376	254	376
3	three	51	63	51	63	51	63	51	63
¾	threequarters	—	—	—	—	190	276	190	276
¾	threesuperior	—	—	—	—	179	263	179	263
~	tilde	196	304	247	367	152	230	31	37
™	trademark	—	—	170	252	153	231	146	222
2	two	50	62	50	62	50	62	50	62
²	twosuperior	—	—	—	—	178	262	178	262
u	u	117	165	117	165	117	165	117	165
ú	uacute	—	—	156	234	250	372	250	372
û	ucircumflex	—	—	158	236	251	373	251	373
ü	udieresis	—	—	159	237	252	374	252	374
ù	ugrave	—	—	157	235	249	371	249	371
_	underscore	95	137	95	137	95	137	95	137
v	v	118	166	118	166	118	166	118	166
w	w	119	167	119	167	119	167	119	167
x	x	120	170	120	170	120	170	120	170
y	y	121	171	121	171	121	171	121	171

Char	Name	StandardEncoding		MacRomanEncoding		WinAnsiEncoding		PDFDocEncoding	
		Decimal	Octal	Decimal	Octal	Decimal	Octal	Decimal	Octal
ý	yacute	—	—	—	—	253	375	253	375
ÿ	ydiereasis	—	—	216	330	255	377	255	377
¥	yen	165	245	180	264	165	245	165	245
z	z	122	172	122	172	122	172	122	172
	zcaron	—	—	—	—	158	236	158	236
0	zero	48	60	48	60	48	60	48	60

PDF 1.3

Note In the **WinAnsiEncoding**, the hyphen character can also be accessed using a character code of 173, the space using 160, and bullets are used for the otherwise unused character codes 127, 129, 141, 143, 144, and 157.

Note Note that in the **WinAnsiEncoding**, character codes 142 and 158 are now used for Zcaron and zcaron, respectively, and that character code 128 is used for Euro. Position 219 in the **MacRomanEncoding** is still used for currency even though Apple has redefined it to be Euro. Documents that wish to access Euro using **MacRomanEncoding** must explicitly re-encode the font.

PDF 1.3

C.2 Predefined encodings sorted by character code

Note Character codes 0 through 23 are not used in any of the predefined encodings.

Code		StandardEncoding	MacRomanEncoding	WinAnsiEncoding	PDFDocEncoding
Decimal	Octal				
24	30	—	—	—	breve
25	31	—	—	—	caron
26	32	—	—	—	circumflex
27	33	—	—	—	dotaccent
28	34	—	—	—	hungarumlaut
29	35	—	—	—	ogonek
30	36	—	—	—	ring
31	37	—	—	—	tilde
32	40	space	space	space	space
33	41	exclam	exclam	exclam	exclam
34	42	quotedbl	quotedbl	quotedbl	quotedbl
35	43	numbersign	numbersign	numbersign	numbersign
36	44	dollar	dollar	dollar	dollar
37	45	percent	percent	percent	percent
38	46	ampersand	ampersand	ampersand	ampersand
39	47	quoteright	quotesingle	quotesingle	quotesingle
40	50	parenleft	parenleft	parenleft	parenleft
41	51	parenright	parenright	parenright	parenright
42	52	asterisk	asterisk	asterisk	asterisk
43	53	plus	plus	plus	plus
44	54	comma	comma	comma	comma
45	55	hyphen	hyphen	hyphen	hyphen
46	56	period	period	period	period
47	57	slash	slash	slash	slash
48	60	zero	zero	zero	zero

<i>Code</i>		<i>StandardEncoding</i>	<i>MacRomanEncoding</i>	<i>WinAnsiEncoding</i>	<i>PDFDocEncoding</i>
<i>Decimal</i>	<i>Octal</i>				
49	61	one	one	one	one
50	62	two	two	two	two
51	63	three	three	three	three
52	64	four	four	four	four
53	65	five	five	five	five
54	66	six	six	six	six
55	67	seven	seven	seven	seven
56	70	eight	eight	eight	eight
57	71	nine	nine	nine	nine
58	72	colon	colon	colon	colon
59	73	semicolon	semicolon	semicolon	semicolon
60	74	less	less	less	less
61	75	equal	equal	equal	equal
62	76	greater	greater	greater	greater
63	77	question	question	question	question
64	100	at	at	at	at
65	101	A	A	A	A
66	102	B	B	B	B
67	103	C	C	C	C
68	104	D	D	D	D
69	105	E	E	E	E
70	106	F	F	F	F
71	107	G	G	G	G
72	110	H	H	H	H
73	111	I	I	I	I
74	112	J	J	J	J
75	113	K	K	K	K
76	114	L	L	L	L
77	115	M	M	M	M
78	116	N	N	N	N
79	117	O	O	O	O
80	120	P	P	P	P
81	121	Q	Q	Q	Q
82	122	R	R	R	R
83	123	S	S	S	S
84	124	T	T	T	T
85	125	U	U	U	U
86	126	V	V	V	V
87	127	W	W	W	W
88	130	X	X	X	X
89	131	Y	Y	Y	Y
90	132	Z	Z	Z	Z
91	133	bracketleft	bracketleft	bracketleft	bracketleft
92	134	backslash	backslash	backslash	backslash
93	135	bracketright	bracketright	bracketright	bracketright
94	136	asciicircum	asciicircum	asciicircum	asciicircum
95	137	underscore	underscore	underscore	underscore
96	140	quoteleft	grave	grave	grave
97	141	a	a	a	a

<i>Code</i>		<i>StandardEncoding</i>	<i>MacRomanEncoding</i>	<i>WinAnsiEncoding</i>	<i>PDFDocEncoding</i>
<i>Decimal</i>	<i>Octal</i>				
98	142	b	b	b	b
99	143	c	c	c	c
100	144	d	d	d	d
101	145	e	e	e	e
102	146	f	f	f	f
103	147	g	g	g	g
104	150	h	h	h	h
105	151	i	i	i	i
106	152	j	j	j	j
107	153	k	k	k	k
108	154	l	l	l	l
109	155	m	m	m	m
110	156	n	n	n	n
111	157	o	o	o	o
112	160	p	p	p	p
113	161	q	q	q	q
114	162	r	r	r	r
115	163	s	s	s	s
116	164	t	t	t	t
117	165	u	u	u	u
118	166	v	v	v	v
119	167	w	w	w	w
120	170	x	x	x	x
121	171	y	y	y	y
122	172	z	z	z	z
123	173	braceleft	braceleft	braceleft	braceleft
124	174	bar	bar	bar	bar
125	175	braceright	braceright	braceright	braceright
126	176	asciitilde	asciitilde	asciitilde	asciitilde
127	177	—	—	bullet	—
128	200	—	Adieresis	Euro	bullet
129	201	—	Aring	bullet	dagger
130	202	—	Ccedilla	quotesinglbase	daggerdbl
131	203	—	Eacute	florin	ellipsis
132	204	—	Ntilde	quotedblbase	emdash
133	205	—	Odieresis	ellipsis	endash
134	206	—	Udieresis	dagger	florin
135	207	—	aacute	daggerdbl	fraction
136	210	—	agrave	circumflex	guilsinglleft
137	211	—	acircumflex	perthousand	guilsinglright
138	212	—	adieresis	Scaron	minus
139	213	—	atilde	guilsinglleft	perthousand
140	214	—	aring	OE	quotedblbase
141	215	—	ccedilla	bullet	quotedblleft
142	216	—	eacute	Zcaron	quotedblright
143	217	—	egrave	bullet	quoteleft
144	220	—	ecircumflex	bullet	quoteright
145	221	—	edieresis	quoteleft	quotesinglbase
146	222	—	iacute	quoteright	trademark

PDF 1.3

PDF 1.3

<i>Code</i>		<i>StandardEncoding</i>	<i>MacRomanEncoding</i>	<i>WinAnsiEncoding</i>	<i>PDFDocEncoding</i>
<i>Decimal</i>	<i>Octal</i>				
147	223	—	igrave	quotedblleft	fi
148	224	—	icircumflex	quotedblright	fl
149	225	—	idieresis	bullet	Lslash
150	226	—	ntilde	endash	OE
151	227	—	oacute	emdash	Scaron
152	230	—	ograve	tilde	Ydieresis
153	231	—	ocircumflex	trademark	Zcaron
154	232	—	odieresis	scaron	dotlessi
155	233	—	otilde	guilsingright	lslash
156	234	—	uacute	oe	oe
157	235	—	ugrave	bullet	scaron
158	236	—	ucircumflex	zcaron	zcaron
159	237	—	udieresis	Ydieresis	—
160	240	—	dagger	space	Euro
161	241	exclamdown	degree	exclamdown	exclamdown
162	242	cent	cent	cent	cent
163	243	sterling	sterling	sterling	sterling
164	244	fraction	section	currency	currency
165	245	yen	bullet	yen	yen
166	246	florin	paragraph	brokenbar	brokenbar
167	247	section	germandbls	section	section
168	250	currency	registered	dieresis	dieresis
169	251	quotesingle	copyright	copyright	copyright
170	252	quotedblleft	trademark	ordfeminine	ordfeminine
171	253	guillemotleft	acute	guillemotleft	guillemotleft
172	254	guilsingleft	dieresis	logicalnot	logicalnot
173	255	guilsingright	—	hyphen	—
174	256	fi	AE	registered	registered
175	257	fl	Oslash	macron	macron
176	260	—	—	degree	degree
177	261	endash	plusminus	plusminus	plusminus
178	262	dagger	—	twosuperior	twosuperior
179	263	daggerdbl	—	threesuperior	threesuperior
180	264	periodcentered	yen	acute	acute
181	265	—	mu	mu	mu
182	266	paragraph	—	paragraph	paragraph
183	267	bullet	—	periodcentered	periodcentered
184	270	quotesinglbase	—	cedilla	cedilla
185	271	quotedblbase	—	onesuperior	onesuperior
186	272	quotedblright	—	ordmasculine	ordmasculine
187	273	guillemotright	ordfeminine	guillemotright	guillemotright
188	274	ellipsis	ordmasculine	onequarter	onequarter
189	275	perthousand	—	onehalf	onehalf
190	276	—	ae	threequarters	threequarters
191	277	questiondown	oslash	questiondown	questiondown
192	300	—	questiondown	Agrave	Agrave
193	301	grave	exclamdown	Aacute	Aacute
194	302	acute	logicalnot	Acircumflex	Acircumflex
195	303	circumflex	—	Atilde	Atilde

PDF 1.3

PDF 1.3

<i>Code</i>		<i>StandardEncoding</i>	<i>MacRomanEncoding</i>	<i>WinAnsiEncoding</i>	<i>PDFDocEncoding</i>
<i>Decimal</i>	<i>Octal</i>				
196	304	tilde	florin	Adieresis	Adieresis
197	305	macron	—	Aring	Aring
198	306	breve	—	AE	AE
199	307	dotaccent	guillemotleft	Ccedilla	Ccedilla
200	310	dieresis	guillemotright	Egrave	Egrave
201	311	—	ellipsis	Eacute	Eacute
202	312	ring	space	Ecircumflex	Ecircumflex
203	313	cedilla	Agrave	Edieresis	Edieresis
204	314	—	Atilde	Igrave	Igrave
205	315	hungarumlaut	Otilde	Iacute	Iacute
206	316	ogonek	OE	Icircumflex	Icircumflex
207	317	caron	oe	Idieresis	Idieresis
208	320	emdash	endash	Eth	Eth
209	321	—	emdash	Ntilde	Ntilde
210	322	—	quotedblleft	Ograve	Ograve
211	323	—	quotedblright	Oacute	Oacute
212	324	—	quoteleft	Ocircumflex	Ocircumflex
213	325	—	quoteright	Otilde	Otilde
214	326	—	divide	Odieresis	Odieresis
215	327	—	—	multiply	multiply
216	330	—	ydieresis	Oslash	Oslash
217	331	—	Ydieresis	Ugrave	Ugrave
218	332	—	fraction	Uacute	Uacute
219	333	—	currency	Ucircumflex	Ucircumflex
220	334	—	guilsinglleft	Udieresis	Udieresis
221	335	—	guilsinglright	Yacute	Yacute
222	336	—	fi	Thorn	Thorn
223	337	—	fl	germandbls	germandbls
224	340	—	daggerdbl	agrave	agrave
225	341	AE	periodcentered	aacute	aacute
226	342	—	quotesinglbase	acircumflex	acircumflex
227	343	ordfeminine	quotedblbase	atilde	atilde
228	344	—	perthousand	adieresis	adieresis
229	345	—	Acircumflex	aring	aring
230	346	—	Ecircumflex	ae	ae
231	347	—	Aacute	ccedilla	ccedilla
232	350	Lslash	Edieresis	egrave	egrave
233	351	Oslash	Egrave	eacute	eacute
234	352	OE	Iacute	ecircumflex	ecircumflex
235	353	ordmasculine	Icircumflex	edieresis	edieresis
236	354	—	Idieresis	igrave	igrave
237	355	—	Igrave	iacute	iacute
238	356	—	Oacute	icircumflex	icircumflex
239	357	—	Ocircumflex	idieresis	idieresis
240	360	—	—	eth	eth
241	361	ae	Ograve	ntilde	ntilde
242	362	—	Uacute	ograve	ograve
243	363	—	Ucircumflex	oacute	oacute
244	364	—	Ugrave	ocircumflex	ocircumflex

:

March 11, 1999

<i>Code</i>		<i>StandardEncoding</i>	<i>MacRomanEncoding</i>	<i>WinAnsiEncoding</i>	<i>PDFDocEncoding</i>
<i>Decimal</i>	<i>Octal</i>				
245	365	dotlessi	dotlessi	otilde	otilde
246	366	—	circumflex	odieresis	odieresis
247	367	—	tilde	divide	divide
248	370	lslash	macron	oslash	oslash
249	371	oslash	breve	ugrave	ugrave
250	372	oe	dotaccent	uacute	uacute
251	373	germandbls	ring	ucircumflex	ucircumflex
252	374	—	cedilla	udieresis	udieresis
253	375	—	hungarumlaut	yacute	yacute
254	376	—	ogonek	thorn	thorn
255	377	—	caron	ydieresis	ydieresis

C.3 MacExpert encoding

Char	Name	Code		Char	Name	Code	
		Decimal	Octal			Decimal	Octal
æ	AEmall	190	276	¬	Lslashsmall	194	302
á	Aacutesmall	135	207	l	Lsmall	108	154
â	Acircumflexsmall	137	211	Û	Macronsmall	244	364
´	Acutesmall	39	47	m	Msmall	109	155
ä	Adieresissmall	138	212	n	Nsmall	110	156
à	Agravesmall	136	210	ñ	Ntildesmall	150	226
å	Aringsmall	140	214	œ	OEmall	207	317
a	Asmall	97	141	ó	Oacutesmall	151	227
ã	Atildesmall	139	213	ô	Ocircumflexsmall	153	231
Û	Brevesmall	243	363	ö	Odieresissmall	154	232
b	Bsmall	98	142	Ú	Ogoneksmall	242	362
Æ	Caronsmall	174	256	ò	Ogravesmall	152	230
ç	Ccedillasmall	141	215	ø	Oslashsmall	191	277
...	Cedillasmall	201	311	o	Osmall	111	157
^	Circumflexsmall	94	136	õ	Otildesmall	155	233
c	Csmall	99	143	p	Psmall	112	160
¨	Dieresissmall	172	254	q	Qsmall	113	161
•	Dotaccentsmall	250	372	°	Ringsmall	251	373
d	Dsmall	100	144	r	Rsmall	114	162
é	Eacutesmall	142	216	ß	Scaronsmall	167	247
ê	Ecircumflexsmall	144	220	s	Ssmall	115	163
ë	Edieresissmall	145	221	¼	Thornsmall	185	271
è	Egravesmall	143	217	~	Tildesmall	126	176
e	Esmall	101	145	t	Tsmall	116	164
D	Ethsmall	68	104	ú	Uacutesmall	156	234
f	Fsmall	102	146	û	Ucircumflexsmall	158	236
`	Gravesmall	96	140	ü	Udieresissmall	159	237
g	Gsmall	103	147	ù	Ugravesmall	157	235
h	Hsmall	104	150	u	Usmall	117	165
"	Hungarumlautsmall	34	42	v	Vsmall	118	166
í	Iacutesmall	146	222	w	Wsmall	119	167
î	Icircumflexsmall	148	224	x	Xsmall	120	170
ï	Idieresissmall	149	225	¥	Yacutesmall	180	264
ì	Igravesmall	147	223	ÿ	Ydieresissmall	216	330
i	Ismall	105	151	y	Ysmall	121	171
j	Jsmall	106	152	¾	Zcaronsmall	189	275
k	Ksmall	107	153	z	Zsmall	122	172
&	ampersandsmall	38	46	Ò	lsuperior	241	361
Å	asuperior	129	201	~	msuperior	247	367
•	bsuperior	245	365	ª	nineinferior	187	273
©	centinferior	169	251	9	nineoldstyle	57	71
#	centoldstyle	35	43	·	ninesuperior	225	341
Ç	centsuperior	130	202	^	nsuperior	246	366
:	colon	58	72	+	onedotenleader	43	53

Char	Name	Code		Char	Name	Code	
		Decimal	Octal			Decimal	Octal
{	colonmonetary	123	173	J	oneeighth	74	112
,	comma	44	54		onefitted	124	174
ø	commainferior	178	262	H	onehalf	72	110
-	commasuperior	248	370	ı	oneinferior	193	301
1	dollarinferior	182	266	l	oneoldstyle	49	61
\$	dollaroldstyle	36	44	G	onequarter	71	107
%	dollarsuperior	37	45	/	onesuperior	218	332
Î	dsuperior	235	353	N	onethird	78	116
•	eightinferior	165	245	∅	osuperior	175	257
8	eightoldstyle	56	70	[parenleftinferior	91	133
°	eightsuperior	161	241	(parenleftsuperior	40	50
‰	esuperior	228	344]	parenrightinferior	93	135
÷	exclamdownsmall	214	326)	parenrightsuperior	41	51
!	exclamsmall	33	41	.	period	46	56
V	ff	86	126	Š	periodinferior	179	263
Y	ffi	89	131	•	periodsuperior	249	371
Z	ffl	90	132	ı	questiondownsmall	192	300
W	fi	87	127	?	questionsmall	63	77
-	figuredash	208	320	Â	rsuperior	229	345
L	fiveeighths	76	114	}	rupiah	125	175
x	fiveinferior	176	260	;	semicolon	59	73
5	fiveoldstyle	53	65	M	seveneighths	77	115
	fivesuperior	222	336	¶	seveninferior	166	246
X	fl	88	130	7	sevenoldstyle	55	67
¢	fourinferior	162	242	‡	sevensuperior	224	340
4	fouroidstyle	52	64	§	sixinferior	164	244
>	foursuperior	221	335	6	sixoldstyle	54	66
/	fraction	47	57	?	sixsuperior	223	337
-	hyphen	45	55		space	32	40
_	hypheninferior	95	137	Í	ssuperior	234	352
—	hyphensuperior	209	321	K	threeeighths	75	113
È	isuperior	233	351	£	threeinferior	163	243
3	threeoldstyle	51	63				
I	threequarters	73	111				
=	threequartersemdash	61	75				
<	threesuperior	220	334				
Ê	tsuperior	230	346				
*	twodotenleader	42	52				
™	twoinferior	170	252				
2	twooldstyle	50	62				
²	twosuperior	219	333				
O	twothirds	79	117				
°	zeroinferior	188	274				
0	zerooldstyle	48	60				
,	zerosuperior	226	342				

March 11, 1999

:

March 11, 1999

APPENDIX D**Implementation Limits**

In general, PDF does not restrict the size or quantity of things described in the file format, such as numbers, arrays, images, and so on. However, a PDF viewer application running on a particular processor and in a particular operating environment does have such limits. If a viewer application attempts to perform an action that exceeds one of the limits, it will display an error.

PostScript interpreters also have implementation limits, listed in Appendix B of the *PostScript Language Reference Manual, Third Edition* [1]. It is possible to construct a PDF file that does not violate viewer application limits but will not print on a PostScript printer. Keep in mind that these limits vary according to the PostScript language level, interpreter version, and the amount of memory available to the interpreter.

All limits are sufficiently large that most PDF files should never approach them. However, using the techniques described in Chapters 10 through 14 of this book will further reduce the chance of reaching these limits.

This appendix describes typical limits for Acrobat. These limits fall into two main classes:

- *Architectural limits.* The hardware on which a viewer application executes imposes certain constraints. For example, an integer is usually represented in 32 bits, limiting the range of allowed integers. In addition, the design of the software imposes other constraints, such as a limit of 65,535 elements in an array or string.
- *Memory limits.* The amount of memory available to a viewer application limits the number of memory-consuming objects that can be held simultaneously.

PDF itself has one architectural limit. Because ten digits are allocated to byte offsets, the size of a file is limited to 10^{10} bytes (approximately 10GB).

[Table D.1](#) describes the architectural limits for most PDF viewer applications running on 32-bit machines. These limits are likely to remain constant across a wide variety of implementations. However, memory limits will often be exceeded before architectural limits, such as the limit on the number of PDF objects, are reached.

Table D.1 *Architectural limits*

<i>Quantity</i>	<i>Limit</i>	<i>Explanation</i>
integer	2,147,483,647	Largest positive value, $2^{31} - 1$.
	-2,147,483,648	Largest negative value, -2^{31} .
real	$\pm 32,767$	Approximate range of values.
	$\pm 1/65,536$	Approximate smallest non-zero value.
	5	Approximate number of decimal digits of precision in fractional part.
array	8,191	Maximum number of elements in an array.
dictionary	4,095	Maximum number of key-value pairs in a dictionary.
string	65,535	Maximum number of characters in a string.
name	127	Maximum number of characters in a name.
indirect object	250,000	Maximum number of indirect objects in a PDF file.

Memory limits cannot be characterized so precisely, because the amount of available memory and the way in which it is allocated vary from one implementation to another.

Memory is automatically reallocated from one use to another when necessary. When more memory is needed for a particular purpose, it can be taken away from memory allocated to another purpose if that memory is currently unused or its use is non-essential (a cache, for example). Also, data is often saved to a temporary file when memory is limited. Because of this behavior, it is not possible to state limits for such items as the number of pages, number of text annotations or hypertext links on a page, number of graphics objects on a page, or number of fonts on a page or in a document.

Acrobat has some additional architectural limits:

- Thumbnails may be no larger than 106×106 samples, and should be created at one-eighth scale for 8.5×11 inch and A4 size pages. Thumbnails should use either the **DeviceGray** or direct or indexed **DeviceRGB** color space.
- The minimum allowed page size is 0.25 X 0.25 inch (18×18 units in the default user space coordinate system), and the maximum allowed page size is 200×200 inches (14,400×14,400 units in the default user space coordinate system).
- The zoom factor of a view is constrained to be between approximately 8% and 3,200%. These limits aren't fixed, they vary with the size of the page being displayed, as well as with the size of the pages previously viewed within the file.

- When Acrobat reads a PDF file with a damaged or missing cross-reference table, it attempts to rebuild the table by scanning all the objects in the file. However, the generation numbers of deleted entries are lost if the cross-reference table is missing or severely damaged. Reconstruction fails if any object identifiers do not occur at the start of a line or if the **endobj** keyword does not appear at the start of a line. Also, reconstruction fails if a stream contains a line beginning with the word **endstream**, aside from the required **endstream** that delimits the end of the stream.

Implementation limits affecting Web Capture

PDF 1.3

The data structures constructed by Web Capture (see [page 148](#)) depend on the maximum size of an array, k , which is 8,191 in the Acrobat 4 implementation.

- A Web Capture Content Set array can associate at most k Spider Content Sets with a given name.
- A Web Capture Content Set can reference at most k objects.
- There may be at most k Source Info objects associated with a Web Capture Content Set.
- An Aliased URLs Object can contain at most k chains, and each chain can contain at most k URLs.
- A total of k Web Capture Command objects can be stored in the Commands array of the Web Capture Info dictionary.
- There may be at most $k/2$ Conversion Settings dictionaries stored in a Web Capture Command Settings dictionary.

:

March 11, 1999

APPENDIX E

Obtaining XUIDs and Technical Notes

Creators of widely distributed forms who wish to use the XUID mechanism must obtain an organization ID from Adobe Systems Incorporated at the addresses listed below. XUID organization IDs are held by Adobe in a public registry. The registry is mechanism for preventing conflict among products of different developers.

Technical notes, technical support, and periodic mailings are available to members of the Adobe Developers Association. In particular, the PostScript language software development kit (SDK) contains all the technical notes mentioned in this book. The Adobe Developers Association can be contacted at the addresses listed below:

Europe:

Adobe Developers Association
PO Box 12356
Edinburgh EH11 4GJ
United Kingdom
ADA Helpdesk: +44.131.458 6801
E-mail: euroADA@adobe.com

U.S. and the rest of the world:

Adobe Developers Association
Adobe Systems Incorporated
345 Park Avenue
San Jose, CA 95110
Telephone: (408) 536-9000

In addition, some technical notes and other information may be available from Adobe's World Wide Web server

<http://www.adobe.com>

and from an anonymous ftp site

<ftp.adobe.com>

When accessing the anonymous ftp site, use **anonymous** as the user name, and provide your E-mail address as the password (for example, **smith@acme.com**).

:

March 11, 1999

APPENDIX F

PDF Name Registry

With the introduction of Adobe Acrobat 2.0, it has become easy for third parties to add private data to PDF documents and to add plug-ins that change viewer behavior based on this data. However, Acrobat users have certain expectations when opening a PDF document, no matter what plug-ins are available. PDF enforces certain restrictions on private data in order to meet these expectations.

PDF 1.1

A PDF producer or Acrobat viewer plug-in may define new action, destination, annotation, and security handler types. If a user opens a PDF document and the plug-in that implements the new type of object is unavailable, the viewers will behave as described in [Appendix G](#).

A PDF producer or Acrobat plug-in may also add keys to any PDF object that is implemented as a dictionary except the trailer dictionary. In addition, a PDF producer or Acrobat plug-in may create tags that indicate the role of Marked Content operators, as described in [Section 8.10.3 on page 351](#). The names of such tags have additional requirements beyond those specified below.

PDF 1.2

To avoid conflicts with third-party names and with future versions of PDF, Adobe maintains a registry, similar to the registry it maintains for Document Structuring Conventions. Third-party developers must only add private data that conforms to the registry rules. The registry includes three classes:

- *First-class* — Names and data of value to a wide range of developers. All the names defined in PDF 1.0 and 1.1 are first-class names. Plug-ins that are publicly available should often use first-class names for their private data. First class names and data formats must be registered with Adobe, and will be made available for all developers to use. To submit a private data name and format for consideration as first-class, contact Adobe's Developer Support group, as described later in this section.
- *Second-class* — Names that are applicable to a specific developer. (Adobe does not register second-class data formats.) Adobe distributes second-class names by registering developer-specific prefixes, which must be used as the first characters in the names of all private data added by the developer. Adobe will not register the same prefix to two different developers, ensuring that different developers' second-class names will not conflict. It is up to each developer to ensure that they do not use the same name in conflicting ways themselves. To request a prefix for second-class names, contact Adobe's Developer Support group, as described later in this section.

- *Third-class* — Names that can be used only in files that will never be seen by other third parties, because they may conflict with third-class names defined by others. Third-class names all begin with a specific prefix reserved by Adobe for private plug-ins; this prefix is **XX**. This prefix must be used as the first characters in the names of all private data added by the developer. It is not necessary to contact Adobe to register third-class names.

Note New keys for the Info dictionary in the Catalog and in Threads need not be registered.

To register either first- or second-class names, contact Adobe's Developer Support group at (408) 536-9000, or send e-mail to

devsup-person@adobe.com

APPENDIX G

Compatibility and Implementation Notes

The goal of the Adobe Acrobat family of products is to enable people to exchange and view electronic documents easily and reliably. Ideally, that means that any Acrobat viewer should be able to display the contents of any PDF file, even if the PDF file was created long before or long after the viewer. Of course, new versions of viewers are introduced to provide additional capabilities not present before. Furthermore, beginning with Acrobat 2.0, viewers may accept plug-in extensions, making some Acrobat 2.0 viewers more capable than others depending on what extensions are present.

PDF 1.1

Both the viewers and PDF itself have been designed to enable users to view everything in the document that the viewer understands and to ignore or inform the user about objects not understood. The decision whether to ignore or inform the user is made on a feature-by-feature basis.

The original PDF specification did not specify how a viewer should behave when it reads a file that does not conform to the specification. This appendix provides this information. The PDF version number associated with a file determines how it should be treated when a viewer encounters a problem.

In addition, this appendix includes notes on the Adobe Acrobat implementation for details that are not strictly defined by the PDF specifications.

G.1 Version numbers

The PDF version number consists of a major and minor version. The version number is part of the PDF header, the first line of the file. This header takes the form:

`%PDF-M.m`

where *M* is the major number and *m* is the minor number.

If PDF changes in a way that current viewers will be unlikely to read a document without a serious error, the major version number will be incremented. A serious error is an error that prevents pages from being viewed.

If PDF changes in a way that a viewer will display an error message but continue its work, the minor version number will change. Adding new page description operators would require a change in the minor version number.

If PDF changes in a way that current viewers are unlikely to detect, the version number need not change. This includes the addition of private data that can be gracefully ignored by consumers that do not understand that data. An example is adding a key to a dictionary object such as the Catalog.

An Acrobat viewer will try to read any file with a valid PDF header, even if the version number is newer than the viewer itself. It will read without errors any file that does not require a plug-in, even if the version number is older than the viewer. Some documents may require a plug-in to display an annotation, follow a link, or execute an action. Viewer behavior in this situation is described below. However, a plug-in is never required to display the contents of a page.

If a viewer opens a document with a newer major version number than it expects, it warns the user that it is unlikely to be able to read the document successfully and that the user will not be able to change or save the document. At the first error related to document processing, the viewer notifies the user that an error has occurred but that no further errors will be reported. (Some errors are always reported, including file I/O errors, extension loading errors, out-of-memory errors, and notification that a command failed.) Processing continues if possible. Acrobat Exchange does not permit a document with a newer major version number to be inserted into another document.

If a viewer opens a document with a newer minor version number than it expects, it silently remembers the version number. Only if it encounters an error does it alert the user. At this point it notifies the user that the document is newer than expected, that an error has occurred, and that no further errors will be reported. The document may not be incrementally saved but can be saved to a new file. The saved file will continue to have the new version number. A user may insert a document with a newer minor version into another document. The resulting document can be saved. Its version number will be the greater of the version number of the original document and the documents inserted into the original.

When opening a file, the Acrobat viewers are very liberal in their check for a valid PDF header. All viewers allow the header to appear anywhere in the first 1,000 bytes of the file. The 1.0 viewers require only that “%PDF-” appear in the header, but ignore the rest of the header. Subsequent viewers search for a header of the form described above. However, they also accept a header of the form:

```
%!PS-Adobe-N.n PDF-M.m
```

where *N.n* is an Adobe Document Structuring Conventions version number and *M.m* is a PDF version number. (See *PostScript Language Document Structuring Conventions Specification Version 3.0* [14] for more information).

G.2 Dictionary keys

Adding key-value pairs not described in the PDF specification to dictionary objects usually does not affect the behavior of 1.0 viewers and never affects the behavior of subsequent viewers. These keys are ignored. If a dictionary object such as an annotation is copied into another document during a page insertion (or in Acrobat

2.0 and 3.0 viewers during a page extraction), all key-value pairs are copied. If a value is an indirect reference to another object, that object may be copied as well, depending on the key.

In some cases a 1.0 viewer displays an error if it finds an unknown key in a dictionary. These cases are keys in image dictionaries (both XObjects and in-line images) and keys in DecodeParms dictionaries for filters.

See [Appendix F](#) for information on how to choose key names that are compatible with future versions of PDF.

G.3 Implementation notes

The following sections give details of the implementation of Adobe Acrobat. They are presented in the same order in which they appear in the main text.

[Section 4.4, “Strings and text”](#)

All Acrobat viewers can read strings that include non-printable ASCII.

[Section 7.2, “Date”](#)

Acrobat 1.0 viewers report date strings as ordinary strings. Later viewers report date strings as dates when used as the value of the **CreationDate** or **ModDate** in the Info dictionary or as the value of the **Date** key in annotations. The viewers ignore the GMT information.

[Section 4.5, “Names”](#)

In cases where a PostScript name must be preserved, or where a string is permitted in PostScript but not in PDF, the Acrobat Distiller application uses the # convention as necessary. When an Acrobat viewer generates PostScript, it “inverts” the convention by writing a string, where that is permitted, or a name otherwise. For example, if the string (**Adobe Green**) were used as a key in a dictionary, the Distiller program would use the name **/Adobe#20Green**, and the viewer would generate (**Adobe Green**).

[Section 4.8, “Streams” \(Filters\)](#)

PDF uses stream objects to encapsulate the data in images, indexed color spaces, thumbnails, and many other objects. These streams usually use filters to compress their data. The legal PDF 1.0 filters are the same as those available in PostScript Level 2. The 1.0 viewer behavior when encountering an unknown filter depends on its context, as described in [Table G.1](#).

Table G.1 *Acrobat 1.0 Viewer behavior with unknown filters*

<i>Context</i>	<i>Behavior</i>
Image resource	The image does not appear but no error is reported.

In-line image	(An in-line image is specified directly in a page description, while an image resource is specified outside of a page and referenced from the page.) An error is reported, and page processing stops.
Indexed color space	An error is reported, but page processing continues.
Thumbnail	An error is reported, no more thumbnails are displayed, but the thumbnails can be deleted and created again.
Embedded font	An error is reported, and the viewer behaves as if the font is not embedded.
Page description	An error is reported, and page processing stops.
Form description	An error is reported, and page processing stops.
Type 3 character description	An error is reported, and page processing stops.

Subsequent viewers do not allow plug-ins to provide additional filters. If an unrecognized filter is encountered, these viewers specify the context in which the filter was found. If an error occurs while displaying a page, only the first error is reported. Subsequent behavior depends on the context, as described in [Table G.2](#).

Table G.2 *Acrobat 2.0 Viewer behavior with unknown filters*

<i>Context</i>	<i>Behavior</i>
Image resource	The image does not appear but page processing continues.
In-line image	Page processing stops.
Indexed color space	The image does not appear but page processing continues.
Thumbnail	An error is reported, no more thumbnails are displayed, but the thumbnails can be deleted and created again.
Embedded font	The viewer behaves as if the font had not been embedded.
Page description	Page processing stops.
Form description	The form does not appear but page processing continues.
Type 3 character description	The character does not appear but page processing continues. The current point is adjusted based on the character's width.

Operations that process pages, such as Find and Create Thumbnails, stop as soon as an error occurs.

Versions of the Acrobat viewer prior to 3.0 do not understand the **FlateDecode** filter. They display an error message indicating that they are unable to process a page.

When a stream specifies an external file, PDF 1.0 and PDF 1.1 parsers ignore the file and always use the characters between **stream** and **endstream**.

[Section 4.8.9, “DCTDecode filter”](#)

Acrobat 4.0 viewers do not support the combination of the DCTDecode filter with any other filter if the encoded data uses progressive JPEG format.

If an Acrobat viewer earlier than version 4.0 encounters DCTDecode data encoded in progressive JPEG encoding, an error will occur that will be handled according to [Table G.1](#).

[Section 5.12, “PDF files”](#)

The Acrobat 1.0 viewers successfully read files that contain binary data. The restriction on line length is not enforced by any Acrobat viewer.

The Acrobat 1.0 products on the Apple Macintosh computer create files with type '**TEXT**'. Later Acrobat products create files with type '**PDF**'. A user can open these documents from a 1.0 viewer but not from the Finder.

[Section 5.18, “Encryption”](#)

On opening a protected document, a version 1.0 Acrobat viewer displays a blank page or reports that an error was found while processing a page. Subsequent viewers report that a plug-in is required to open the document if the security handler for the document is not available.

[Section 6.2.1, “Viewer preferences”](#)

The **Direction** key is ignored by Acrobat 3.0 and earlier versions of the Acrobat viewer.

[Section 6.2.2, “PageLabel dictionaries”](#)

Older viewers will ignore the **PageLabels** tree and label pages using decimal numbers starting at 1, as is implemented in Acrobat 3.0 and earlier viewers. Manipulating such a PDF file in an older viewer may cause the **PageLabels** array to become out-of-date; at worst this could confuse readers using a more up-to-date viewer.

[Section 6.4, “Page objects”](#)

Acrobat viewers rebuild the Beads array for all pages of a document containing beads if the first page with a bead does not have a Beads array.

[Section 6.6, “Annotations”](#)

An annotation is a dictionary element of a page's **Annots** array. Its **Subtype** specifies the type of annotation. Only **Text** and **Link** are defined by PDF 1.0. If a 1.0 viewer reads a page with an annotation whose **Subtype** is not **Text** or **Link**, it displays an error. It displays one error per page no matter how many annotations are present.

Subsequent viewers display unknown annotations in a closed form similar to text annotations, with an icon containing a question mark. If the user opens the annotation, an alert appears with a message giving the annotation type and explaining that an unavailable plug-in is required to open it. An unknown annotation can be selected, moved, and deleted. Every annotation type must specify its position and size using the **Rect** key.

Acrobat viewers ignore the first two numbers in the **Border** array of an annotation.

Acrobat viewers support a maximum of 10 entries in the dash array (the fourth element of the **Border** array).

Acrobat viewers update the **ModDate** string only for text annotations.

If an Acrobat viewer encounters an annotation type it does not understand (i.e., for which there is no annotation handler), the viewer displays it as an unknown annotation unless the annotation's **F** (Flags) key specifies that the "invisible" flag is set. The **C**, **T**, **M**, and **F** keys are ignored by Acrobat 1.0 viewers. The **H**, **BS**, **AP**, and **AS** keys are ignored by Acrobat 2.0 and 2.1 viewers.

The use of the *Hidden* and *Print* flags (bits 2 and 3) in an annotation has no effect on Acrobat viewers prior to 3.0. Annotations that should be hidden are shown; annotations that should be printed are not printed.

In version 3.0 of Acrobat, neither text annotations nor link annotations respect the *Print* flag.

[Section 6.6.1, "Annotation borders"](#)

Acrobat viewers prior to 3.0 ignore the **BS** key.

If an Acrobat 3.0 viewer encounters a border subtype it doesn't recognize, the border defaults to **S** (solid).

[Section 6.6.2, "Annotation highlighting"](#)

Acrobat viewers prior to 3.0 ignore highlighting modes. The Acrobat 3.0 viewer uses highlighting modes only for **Link** and **Widget** annotations.

[Section 6.6.3, "Annotation appearances"](#)

The presence of an appearance dictionary has no effect in versions of Acrobat prior to 3.0. The additional functionality provided by this construct is not available in older viewers.

[Section 6.6.5, "Link annotations"](#)

Acrobat 1.0 viewers do not report an error when a user activates a link or outline entry that has an unknown destination type or is missing a destination. Links and outline entries with an **A** key will appear to have no destination. Subsequent viewers report an error when the destination or action type is unknown.

[Section 6.6.5, “Link annotations”](#)

A link or a bookmark in PDF 1.0 is a dictionary that contains a **Dest** key whose value specifies a view of the document that is displayed when the link or bookmark is activated. A destination is an array. Its first element is a name that serves as destination type; it determines the interpretation of subsequent array elements. If a 1.0 viewer encounters an unknown destination type, no action is performed and no error is reported when the user activates the link or bookmark. Subsequent viewers display a message when they find an unknown destination type.

An Acrobat 1.0 viewer does nothing if it does not find a **Dest** key in a link or bookmark.

[Section 6.6.6, “Movie Player annotations”](#)

Acrobat 1.0 viewers report the following error when they encounter an annotation of type **Movie**: “An error occurred while reading a note or link. Unknown annotation type.” The annotation does not appear on the document. Subsequent viewers report the following error when they encounter an annotation of type **Movie**: “The Plug-in required by this ‘Movie’ annotation is unavailable.” The annotation is displayed as a gray rectangle with a question mark.

[Section 6.6.16, “PDF Trapping”](#)

Older viewers may fail to maintain the **TrapNet** annotation’s required position at the end of the **Annots** array.

Older viewers may fail to validate trapping networks before printing.

[Section 6.8, “Actions”](#)

Action has superseded destinations in PDF 1.1. An Acrobat 1.0 viewer ignores actions.

The existence of the **Next** key in an action dictionary has no effect in Acrobat viewers prior to version 3.0.

[Section 6.8.1, “Action Trigger Points”](#)

The existence of an Additional Actions dictionary in an annotation, page, outline, or document has no effect in Acrobat viewers prior to 3.0.

In Acrobat 3.0, the **Open** and **Close** actions are disabled if the display is not in a page-oriented mode (e.g., if it is in multi-column mode). This prevents actions for multiple pages from being executed simultaneously, which would be confusing to the user.

In PDF 1.2, the actions associated with the **FP**, **PP**, **NP**, and **LP** keys were not implemented. As of PDF 1.3, these keys are obsolete.

Section 6.8.3, “GoToR action”

The **NewWindow** attribute is ignored by Acrobat viewers prior to 3.0.

Section 6.8.4, “Launch action”

Some implementations of Acrobat viewers may check for alternative keys whose values provide platform-specific parameters for the **Launch** action. For example, the Acrobat viewer for Windows uses the dictionary corresponding to the **Win** key to determine its launch parameters.

The Acrobat viewer for Windows use the Windows function **ShellExecute** to launch an application. The **Win** dictionary entries correspond to the parameters of **ShellExecute**.

The **NewWindow** attribute is ignored by Acrobat viewers prior to 3.0.

Section 6.8.6, “URI action”

Acrobat 1.0 viewers report no error when a link annotation that uses the URI action is invoked. The link inverts its color and performs no action. Subsequent viewers report the following error when a link annotation that uses the URI action is invoked: “The plug-in required by this URI action is not available.”

When resolving the fragment identifier, the WebLink plug-in checks all named destinations defined for the document. If one is found whose name matches the fragment identifier, that destination is invoked.

Section 6.8.8, “Movie Player actions”

Acrobat viewers prior to version 3.0 report an error when they encounter an action of type **Movie**.

Section 6.8.9, “SetState action”

Acrobat viewers prior to 3.0 report the following error when encountering an action of type **SetState**: “The plug-in needed for this SetState action is not available.”

In Acrobat Viewer 3.0 the effect of a **SetState** action is temporary in nature in the same manner that opening or closing text annotations is temporary. If the user saves the document, the changes become permanent. Otherwise, the user is not prompted to save the document and the change in state is not permanent.

Section 6.8.10, “Hide action”

Acrobat viewers prior to 3.0 report the following error when encountering an action of type **Hide**: “The plug-in needed for this Hide action is not available.”

In Acrobat 3.0, the effect of the Hide action is temporary in nature in the same manner that opening or closing text annotations is temporary. If the user saves the document, the changes become permanent. Otherwise, the user is not prompted to save the document and the hiding/showing of the annotation is not permanent.

[Section 6.8.11, “Named actions”](#)

Acrobat viewers prior to 3.0 report the following error when encountering an action of type **Named**: “The plug-in needed for this Named action is not available.”

The Acrobat 3.0 viewer extends the list of Named Actions in [Table 6.49 on page 118](#) by allowing most viewer menu item names to be specified. For further details, see the listing of menu item names in the Acrobat Plug-In Developer’s SDK.

[Section 6.8.12, “NOP action”](#)

Acrobat 1.0 viewers ignore all actions, including NOP. If an Acrobat 2.0 viewer attempts to perform a NOP action, it displays a warning that says that the plug-in required for this action is not present. It is unlikely, however, that such a warning would occur, as this type of action is defined only in places where behavior may be inherited, and there are no such places in Acrobat 2.0.

[Section 6.11, “Articles”](#)

The thread array and dictionary objects are invisible to 1.0 viewers. Consequently, operations that insert or delete pages do not carry along any threads.

[Section 6.12, “File ID”](#)

Although the **ID** key is not required, all Adobe applications that produce PDF include this key. Acrobat Exchange adds this key when saving a file if it is not present.

[Section 6.13, “Encryption dictionary”](#)

In Acrobat 2.0 and 2.1 viewers, the standard security handler uses the empty string if there is no owner password in step 1 of [Algorithm 6.8](#).

[Section 6.14.3, “Field dictionaries”](#)

In Acrobat 3.0, partial field names may not contain a period.

[Section 6.14.10, “Choice”](#)

In Acrobat 3.0, the **Opt** array in a Choice field must be homogenous: the elements must be either all strings or all arrays.

[Section 6.14.13.1, “SubmitForm Action”](#)

In Acrobat 3.0, if the response to a **SubmitForm** action uses Forms Data Format, then the URL must end in **#FDF** so that it is recognized as such by the Acrobat software and handled properly. Conversely, if the response is anything else, then the URL should not end in **#FDF**.

Section 6.14.13.3, “ImportData Action”

Acrobat 3.0 puts a relative file specification (of the FDF with respect to the Form) as the value of **F**. When performing the action, if the FDF is not found, then Acrobat 3.0 tries to locate the file in a few “well-known” platform-dependent locations. For example, on the Windows platform, it looks in the directory from which Acrobat loaded, the current directory, the System directory, the Windows directory, and the directories that are listed in the PATH environment variable; on the Macintosh, it looks in the Preferences folder, and in the Acrobat folder.

When executing the **ImportData** action, Acrobat 3.0 imports the FDF into the current Form, ignoring the **F** and **ID** keys inside the FDF.

Section 7.6, “ProcSets”

Each page includes a ProcSet resource that describe the PostScript procedure sets required to print the page. A 1.0 viewer ignores requests for unknown procedure sets. An Acrobat 2.0 viewer warns the user that a procedure set is unavailable and cancels printing.

Section 7.7, “Fonts”

All Acrobat viewers ignore the **Name** key in a Font resource.

Section 7.7.6, “Type 3 fonts”

For compatibility with Acrobat 1.0, 2.0, and 2.1, the names of resources in a Type 3 font’s Resources dictionary must match the names of resources in the Page’s Resources dictionary. If backwards compatibility is not required, then any valid names may be used.

Section 7.11, “Font descriptors”

Acrobat viewers prior to version 3.0 ignore the **FontFile3** value. If a font uses the Adobe Standard Roman Character CharSet, then Acrobat creates a substitute font. Otherwise, Acrobat displays an error message (once per document), and substitutes any characters in the font with the bullet character.

Section 7.11.1, “Font files”

Embedded TrueType fonts are ignored by Acrobat 1.0 viewers.

Section 7.12, “Color spaces”

An image has a **ColorSpace** key. A 1.0 viewer displays an error each time it finds an image with a color space that is not one of the PDF 1.0 color spaces. Color spaces may not be added by plug-ins. If an Acrobat 2.0 viewer encounters an unknown color space, such as the special color spaces defined in PDF 1.2 (**Pattern**, **Separation**, and some uses of **Indexed**), it will be in a document with a PDF version number greater than 1.1. The viewer displays an error specifying the type of color space but it reports no further errors.

PDF 1.1 defines three additional color spaces: **CalGray**, **CalRGB**, and **Lab**. To be more compatible with 1.0 viewers, PDF 1.1 allows an image color space to be specified indirectly through the Resources dictionary. When an Acrobat 2.0 viewer processes an image and the image's **ColorSpace** key specifies **DeviceRGB**, the viewer looks in the page's resources for a color space called **DefaultRGB**. If this key is present, the color space associated with it is used instead of **DeviceRGB**. Similarly, if an image's **ColorSpace** key specifies **DeviceGray**, the viewer looks for **DefaultGray**. The 1.0 viewer ignores **DefaultRGB** and **DefaultGray**.

See [page 230](#) for an explanation of the use of color spaces in page descriptions. The presence of **DefaultRGB** or **DefaultGray** change the interpretation of some color operators.

Acrobat viewers allow a user to approximate device-independent colors on screen with device-dependent colors with no transformation. **CalGray** colors are viewed as **DeviceGray**, and **CalRGB** colors are viewed as **DeviceRGB**.

[Section 7.12.8, "Separation color spaces"](#)

The Acrobat 3.0 viewer applies the *tintTransform* function, as specified in this section, for displaying graphics that use separation color spaces.

[Section 7.12.9, "DeviceN color spaces"](#)

Acrobat 4.0 supports up to 10 color components for **DeviceN** color spaces.

[Section 7.13, "XObjects"](#)

An XObject is a stream or dictionary that is referred to by name from a page description by the **Do** operator. The effect of the operator is determined by the type of the XObject. A 1.0 viewer displays an error for each XObject of a different type, no matter how many are on a page.

Plug-ins may not add XObject types, since they are considered part of the page and a viewer without plug-ins should always be able to display a page. If an Acrobat 2.0 viewer encounters an unknown XObject type, it will be in a document with a PDF version number greater than 1.1. The viewer displays an error specifying the type of XObject but it reports no further errors.

[Section 7.13.1, "Images"](#)

The **Name** key in an Image resource is ignored by all Acrobat viewers.

Image XObjects in PDF 1.2 and earlier versions are all implicitly unmasked images. A PDF consumer that does not recognize the **Mask** key will treat the image as unmasked, and no error will be raised.

[Section 7.13.4, "Color rendering intent"](#)

The Acrobat 1.0 viewers display an error if an image specifies an **Intent**.

Because of the large gamut of most displays, Acrobat viewers ignore the **Intent** key when displaying a PDF file and always use **RelativeColorimetric**. When printing to a PostScript printer, the Acrobat viewers do not specify an intent unless one was explicitly specified.

Section 7.13.5, “Masking images by position”

Masked images are a feature of PostScript Level 3. Acrobat 4.0 does not attempt to emulate the effect of masked images when printing to Level 1 or Level 2 printers; it prints the base image without the mask.

The Acrobat 4.0 viewer will display masked images, but only when the amount of data in the imagemask is below a certain limit. Above that, the viewer will display the base image without the mask.

Section 7.13.9, “OPI dictionary”

The Acrobat 3.0 Distiller program converts OPI comments into OPI dictionaries, and when the Acrobat 3.0 Viewer prints a PDF file to a PostScript file or printer, it converts the OPI dictionary to OPI comments. However, the OPI information has no effect on the displayed XObject (form or image).

In Acrobat 3.0, the value of the **F** key in an OPI dictionary must be a string.

The Acrobat 3.0 Distiller application and the Acrobat 3.0 Viewer do not support OPI 2.0.

Section 7.14.1, “Sampled functions (Function Type 0)”

Acrobat 3.0 supports only linear interpolation (**Order** 1).

Section 7.14.4, “PostScript Calculator Function (FunctionType 4)”

Functions of types 2, 3, and 4 are not compatible with PDF 1.2 or earlier versions. Acrobat will report an error, “Invalid Function Resource,” if it encounters a function of this type. The syntax of numbers in Type 4 functions is the same as in the rest of PDF; this is slightly different from the syntax for numbers in PostScript.

Section 7.16.3, “Predefined spot functions”

If the Acrobat 3.0 Distiller encounters a call to **setscreen** or **sethalftone**, and if that call includes a spot function, the Distiller examines the code for the spot function. If the code matches the PostScript shown in [Table 7.51 on page 277](#), then the Distiller puts the corresponding name in the halftone dictionary, and Acrobat uses that PostScript code when printing the PDF file to a PostScript printer. Otherwise, the Distiller samples the spot function and generate a function for the halftone dictionary; when printing to a PostScript printer, Acrobat generates a spot function that interpolates values from that function.

Section 7.17, “Patterns”

The Acrobat 3.0 does not display patterns on the screen, but it does print them to PostScript.

Chapter 8, “Page Descriptions”

A 1.0 viewer reports an error the first time it finds an unknown operator or an operator with too few operands, but it continues processing the marking context (e.g., page or form). If it finds ten errors on a page, it reports back to the user and asks whether to continue processing. No further errors are reported. Each time an error occurs, the operand stack is cleared. Later Acrobat viewers behave the same, although there is no additional warning if ten errors are encountered.

PDF 1.1 provides new page description operators for specifying device-independent color and pass-through PostScript fragments. Since these operators are incompatible with 1.0 viewers, PDF 1.1 provides alternative compatible methods as well.

PDF 1.2 defines new operators for setting parameters in the graphics state (gs) and for setting the color in **Pattern** and **Separation** color spaces. There is no compatible mechanism for these operators in viewers prior to Acrobat 3.0.

Section 8.5.2, “Color operators”

For compatibility with PDF 1.0 viewers, it is recommended that device-dependent colors be specified using the 1.0 operators and that device-independent colors be specified using the color space substitution method defined in [Section 7.12.12 on page 245](#).

If an Acrobat 1.0 viewer reads a page containing any of the **setcolorspace**, **setcolor**, or **intent** operators (**cs**, **CS**, **sc**, **SC**, **scn**, **SCN**, or **ri**), it reports an error. Errors can be ignored by the user and objects are displayed, but colors will most likely be black (the default).

The **scn** and **SCN** operators are not compatible with versions of Acrobat prior to 3.0.

Section 8.7.5, “Text string operators”

In versions of Acrobat prior to 3.0, when using the **TJ** operator, the x-coordinate of the current point after drawing a character and moving by any specified offset must not be less than the x-coordinate of the current point before the character was drawn.

Section 8.8.1, “XObject operators”

If an Acrobat 1.0 viewer reads a page containing the **PS** operator, it reports an error. The operator is otherwise ignored.

Section 8.10.2, “Compatibility operators”

If an Acrobat 1.0 viewer reads a page containing the **BX** or **EX** operators, it reports an error. The operators are otherwise ignored.

Section 9.3.2, “Header and linearization information”

:

March 11, 1999

In the Acrobat 3.0 viewer, linearization is an option that is available when a PDF file is saved. It rewrites the entire file (a “full save”) and always uses version 1.2 in the header: **%PDF-1.2**.

APPENDIX H

Forms Data Format

This appendix describes FDF, the file format used for Acrobat Forms data. FDF is used when submitting Form data to a server, receiving the response, and incorporating it into the Form. It can also be used to generate (i.e. “export”) stand-alone files containing Form data that can be stored, transmitted electronically (e.g., via e-mail), and imported back into the corresponding Form.

PDF 1.2

FDF can also be used to control more of the document structure. That is, constructs within FDF allow it to control which Acrobat Forms are used in the creation of a new PDF document. This functionality can be used to create complex documents dynamically.

PDF 1.3

FDF is also used to define a container for annotations that are separate from the PDF document to which the annotations apply. (See [Section H.5 on page 492](#) for an example.)

PDF 1.3

- FDF is based on PDF, and uses the same syntax and set of basic object types as PDF.
- FDF also has the same file structure as PDF, except for the fact that the cross-reference is optional.
- The document structure is much simpler than PDF, since the body of an FDF document consists of only one required object.
- Objects in FDF can only be of generation 0; no two objects can have the same object number, and FDF files cannot have updates appended to them.
- The value of the **Length** attribute in the dictionary of any stream object appearing inside an FDF document must be a direct object.
- FDF uses the MIME type **application/vnd.fdf**. On Windows and UNIX it uses the ***.fdf** extension, and on the Mac it has the **'FDF'** file type.

H.1 File Structure

An FDF file consists of a one-line header, a body, and a trailer. It can optionally contain a cross-reference table. In other words, FDF is structured in the same way as PDF, but need only contain those elements required for Acrobat Forms data export and import, which are described below.

Figure 14.5 *FDF File Structure*

FDF File Structure

Header
Body
Cross Reference Table (optional)
Trailer

H.1.1 Header

The first line of an FDF file specifies the version number of the PDF specification that FDF is a part of. The current version of PDF is 1.3; therefore the first line of an FDF file is

%FDF-1.3

H.1.2 Body

The body consists of one Catalog object and any additional indirect objects that it may reference. The Catalog object is a dictionary with only one (required) key in it, **FDF**. Its value is a dictionary, whose entries are described in Section H.2, “[The FDF Catalog Object](#).”

It is legal for the body to contain additional objects, and for the Catalog object to contain additional key-value pairs. Comments can appear anywhere in the body section of an FDF file.

Just as in PDF, objects in FDF can be direct or indirect.

H.1.3 Trailer

The trailer consists of a trailer dictionary, followed by the last line of the FDF file, containing the end-of-file marker, **%%EOF**. The trailer dictionary consists of the keyword **trailer**, followed by at least one key-value pair enclosed in double angle brackets. The only required key is **Root**, and its value is an indirect reference to the Catalog object in the FDF body.

It is legal for the trailer dictionary to contain the additional key-value pairs described in the PDF specification.

H.2 The FDF Catalog Object

The value of the **FDF** key in the Catalog object is a dictionary whose entries are described in the following table.

Table H.1 *FDF attributes*

Key	Type	Semantics
Fields	array	<i>(Optional)</i> This array contains the root fields being exported or imported. A root field is one with no parent (i.e., it is not in the Kids array of another field). The attributes of the fields are described in Section H.2.1, “FDF fields.”
Status	string	<i>(Optional)</i> A status to be displayed indicating the result of an action, typically a SubmitForm action (see Section 6.14.13.1 on page 143). This string is encoded with PDFDocEncoding .
	<i>Implementation note</i>	<i>The Acrobat 3.0 implementation of Forms displays the Status, if any, in an Alert Note, when importing an FDF.</i>
F	File specification	<i>(Optional)</i> File specification for the PDF document that this FDF was exported from, or is meant to be imported into.
ID	array	<i>(Optional)</i> The value of the ID field in the trailer dictionary of the PDF document that this FDF was exported from, or is meant to be imported into.
Pages	array	<i>(Optional)</i> This array causes new pages to be added to a PDF document. See Section H.2.3 on page 489 for a description of the elements in this array. Use of the Pages key precludes the use of the Fields and Status keys. PDF 1.3
Encoding	name	<i>(Optional)</i> The encoding to be used for any FDF field value (V key) or option (Opt key) that is a string and does not begin with the Unicode prefix <FE FF>. The default is PDFDocEncoding . PDF 1.3
	<i>Implementation note</i>	<i>The only Encoding value supported by Acrobat 4.0 is Shift-JIS. All other values will use the default, PDFDocEncoding.</i>
Annot	array	<i>(Optional)</i> An array of FDF annotation dictionaries. See Section H.2.4 on page 491 . PDF 1.3

H.2.1 FDF fields

[Table H.2](#) describes the attributes of each field in the FDF. The majority of the attributes described in this table represent the same information and have the same semantics as the attributes of the same name described in [Section 6.14.3 on page 131](#).

Table H.2 *Field attributes*

<i>Key</i>	<i>Type</i>	<i>Semantics</i>
T	string	<i>(Required)</i> The partial field name.
Kids	array	<i>(Optional)</i> Contains the child field dictionaries.
V	<i>various</i>	<i>(Optional)</i> Field value.
Opt	array	<i>(Optional)</i> Options.
Ff	integer	<i>(Optional)</i> Field flags. When imported into an Acrobat Form, it replaces the current value of the Ff key in the corresponding field inside the Form. If SetFf and/or ClrFf are also present, they are ignored.
SetFf	integer	<i>(Optional)</i> Field flags. When imported into an Acrobat Form, it is OR'ed with the current value of the Ff key in the corresponding field inside the Form.
ClrFf	integer	<i>(Optional)</i> Field flags. When imported into an Acrobat Form, for each bit that is set to one in this value, sets the corresponding bit in the Form field's Ff flags to zero. If SetFf is also present, ClrFf is applied after SetFf .
F	integer	<i>(Optional)</i> Widget annotation flags. When imported into an Acrobat Form, it replaces the current value of the F key in the corresponding field inside the Form. If SetF and/or ClrF are also present, they are ignored.
SetF	integer	<i>(Optional)</i> Widget annotation flags. When imported into an Acrobat Form, it is OR'ed with the current value of the F key in the corresponding field inside the Form.
ClrF	integer	<i>(Optional)</i> Widget annotation flags. When imported into an Acrobat Form, for each bit that is set to one in this value, sets the corresponding bit in the Form's F flags to zero. If SetF is also present, ClrF is applied after SetF .
AP	dictionary	<i>(Optional)</i> Dictionary containing the appearances for a Push Button field (see Section 6.14.7 on page 135). Similar to Table 6.13 on page 88 , except that the values of the N , R , and D keys must all be streams.
AS	name	<i>(Optional)</i> Appearance state.
A	dictionary	<i>(Optional)</i> Action to be performed on activation of this Widget annotation.
AA	dictionary	<i>(Optional)</i> Additional actions.
APRef	dictionary	<i>(Optional)</i> Dictionary that includes references to external PDF files containing the pages to use for the appearances for a Push Button field. Similar to Table 6.13 on page 88 , except that the values of the N , R , and D keys must all be named page reference dictionaries, described in Table H.6 . If both an AP and an APRef are provided, the AP is used.
IF	dictionary	<i>(Optional)</i> See Section H.2.2, "Icon-fit dictionary."

PDF 1.3

PDF 1.3

H.2.2 Icon-fit dictionary

The Icon-fit dictionary controls how the button icon is to be manipulated within the boundaries of the button. The icon-fit dictionary, if provided, must contain three keys, described in [Table H.3](#).

Table H.3 *Icon-fit attributes*

PDF 1.3

Key	Type	Semantics
SW (Scale When)	name	<p><i>(Required)</i> Indicates when the icon should be scaled inside the button. The value is one of:</p> <p>B Scale when icon is too big, or only scale down. S Scale when icon is too small, or only scale up. N Never scale. A Always scale. This is the default value, used if no icon-fit dictionary is provided.</p>
S (Scaling)	name	<p><i>(Required)</i> Indicates how the icon should be scaled inside the button. Possible values are:</p> <p>A (Anamorphic) Always exactly fills the BBox for the button. Anamorphic scaling does not maintain the aspect ratio of the icon.</p> <p>P (Proportional) Aspect ratio is maintained. The icon is scaled until the contents fill the BBox for the field annotation. This is the default value, used if no icon-fit dictionary is provided.</p>
A (Align)	array	<p><i>(Required)</i> An array of two numbers between 0 and 1 indicating the fraction of leftover space to assign to the left and bottom location of the icon. A value of [0.5 0.5] centers the icon in the Bbox, and a value of [0 0] makes the icon appear in the lower left corner of the Bbox.</p> <p>This array is not used if the icon is scaled anamorphically. If no icon-fit dictionary is provided, a value of [0.5 0.5] is used.</p>

H.2.3 FDF Pages object

Table H.4 *Pages-object attributes*

PDF 1.3

Key	Type	Semantics
Templates	array	<p><i>(Required)</i> An array of named page dictionaries that describe the named pages that serve as templates on a page. The Templates dictionary attributes are described in Table H.5.</p>

Info dictionary (Optional) Contains other information about the page described by this dictionary. Currently there are no attributes defined.

Table H.5 *Templates dictionary attributes*

PDF 1.3

Key	Type	Semantics
TRef	dictionary	(Required) A named page reference dictionary that describes the location of the template. The named page reference dictionary attributes are described in Table H.6 .
Fields	array	(Optional) This array contains the root fields being imported. A root field is one with no parent (i.e., it is not in the Kids array of another field). The attributes of the fields are described in Table H.2 .
Rename	boolean	(Optional) If <i>false</i> , prevents the fields from being renamed during FDF importing. The default value is <i>true</i> . Rename affects the fields described in the same dictionary and their children. The effect of Rename is described more fully below.

The **Rename** key is used when a new page is being added to a PDF document under the control of the FDF, and a new field on the new page has the same name as an existing field. This can occur if the same template page is imported more than once, or if two different templates are imported, but have fields with the same names.

If **Rename** is *true*, fields from the document being imported are renamed to guarantee their uniqueness. If **Rename** is *false*, the fields are not renamed, and each time the FDF provides keys for that field, all fields with that name will be updated.

Implementation note Acrobat Exchange renames fields by prepending a page number, a template name, and an ordinal number to the field name. The template ordinal number corresponds to the order in which the template is applied to a page, with 0 being the first template specified for the page. For example, suppose the first template used on the fifth page has the name **Template** and has **Rename** set to *true*. In this case, the fields defined in that template will be renamed by prepending the character string **P5.Template_0.** to the field names.

Implementation note Adobe Extreme™ printing systems require that the **Rename** key have a value of *true*.

The named page reference dictionary describes the location of external templates or page elements, as shown in [Table H.6](#).

Table H.6 *Named page reference dictionary attributes***PDF 1.3**

Key	Type	Semantics
F	File specification	(Optional) Refers to an external PDF file. If the F key is omitted, it is assumed that the Name refers to a page in the document being imported into.
Name	string	(Required) The name of the referenced named page (template or page element) in the document specified by the F key.

H.2.4 FDF Annotation dictionaries

Each annotation dictionary in an FDF file must have a **Page** key indicating the page of the source document to which the annotation is attached.

Table H.7 *FDF annotation attributes (in addition to those in Table 6.10 on page 82)***PDF 1.3**

Key	Type	Semantics
Page	integer	(Required for annotations in FDF files) The ordinal page number on which this annotations should appear. Page 0 is the first page.

PDF 1.3

H.3 Use of FDF

For most of the keys, unless otherwise indicated in Table H.2, importing consists of taking the value of each key as received in the FDF, and using it to replace the value of the corresponding key in the field inside the Form with the same fully qualified name.

Implementation note Of all the possible keys shown in Table H.2, Acrobat 3.0 will export only the **V** key of a field when generating FDF, and Acrobat 4.0 will export only the **V** and **AP** keys. It will, however, import FDF files containing fields using any of the described keys.

Implementation note Acrobat, when importing an FDF that came back as a result of a SubmitForm action, if the Form currently being displayed is not the one specified in the **F** key of the **FDF** dictionary (which, as stated in Table H.1, is an optional key), then that Form is first fetched, and then the FDF gets imported.

Implementation note When exporting FDF, Acrobat computes a relative path from the location the FDF is being stored, to the location the Form is in, and uses that as the value of the **F** key in the FDF dictionary.

Implementation note Under Acrobat, if an FDF being imported contains fields whose fully qualified names are not present in the Form, those fields will be discarded. This feature can be useful, among other cases, if an FDF containing commonly used fields (such as name, address, etc.) is used to populate various types of Acrobat Forms, each of which does not necessarily include all the fields available in the FDF.

Note As shown in [Table H.2](#), the only required key in the field dictionary is **T**. One possible use for exporting FDF with fields that only contain **T** but no **V**, is as an indication to a server of which fields are desired in the FDF coming back as a response. For example, a server accessing a database might choose between sending all fields in a record, vs. just some selected ones, based on the use of this feature in the request FDF. The implementation of Acrobat Forms will ignore, during import, fields in the FDF that do not exist in the Form.

Implementation note The Acrobat implementation of Forms allows the choice for a SubmitForm action to send the data using HTML format. This is for the benefit of existing server scripts written to process such forms. Note, however, that any such existing scripts that generate new HTML forms as a response will need to be modified to generate FDF instead.

H.4 Sample FDF

```
%FDF-1.2
1 0 obj <<
  /FDF <<
    /Fields
    [
      <<
        /T (My Children)
        /V (Tali)
        /Opt [(Maya) (Adam) (Tali)]
      >>
    ]
  /F (Dependents.pdf)
  >>
  >>
endobj
trailer
<</Root 1 0 R>>
%%EOF
```

H.5 FDF for annotations

The following example shows an FDF file used to express a set of annotations for a separate PDF document.

```
%FDF-1.3
%\342\343\317\323
1 0 obj
  /FDF <<
    /Annot
    [
      <<
        /Type /Annot
```

March 11, 1999

```

    /Subtype /Text
    /Page 0
    /Rect [88 370 359 442]
    /Contents (Converted to Word for Windows)
    /M (D:19970620115631)
    /T (Joe Carousel)
  >>
  <<
    /Type /Annot
    /Subtype /Circle
    /Page 0
    /Rect [232 282 341 444]
    /Contents (Was Macintosh version 5.1a)
    /M (D:19970616124809)
    /T (Joe Carousel)
  >>
]
>>
endobj
%%EOF
```

:

March 11, 1999

APPENDIX I

ISO 639 Language Codes

code	language
aa	Afar
ab	Abkhazian
af	Afrikaans
am	Amharic
ar	Arabic
as	Assamese
ay	Aymara
az	Azerbaijani
ba	Bashkir
be	Byelorussian
bg	Bulgarian
bh	Bihari
bi	Bislama
bn	Bengali; Bangla
bo	Tibetan
br	Breton
ca	Catalan
co	Corsican
cs	Czech

code	language
cy	Welsh
da	Danish
de	German
dz	Bhutani
el	Greek
en	English
eo	Esperanto
es	Spanish
et	Estonian
eu	Basque
fa	Persian
fi	Finnish
fj	Fiji
fo	Faeroese
fr	French
fy	Frisian
ga	Irish
gd	Scots Gaelic
gl	Galician

code	language
gn	Guarani
gu	Gujarati
ha	Hausa
hi	Hindi
hr	Croatian
hu	Hungarian
hy	Armenian
ia	Interlingua
ie	Interlingue
ik	Inupiak
in	Indonesian
is	Icelandic
it	Italian
iw	Hebrew
ja	Japanese
ji	Yiddish
jw	Javanese

code	language
ka	Georgian
kk	Kazakh
kl	Greenlandic
km	Cambodian
kn	Kannada
ko	Korean
ks	Kashmiri
ku	Kurdish
ky	Kirghiz
la	Latin
ln	Lingala
lo	Laothian
lt	Lithuanian
lv	Latvian, Let-tish
mg	Malagasy
mi	Maori
mk	Macedonian
ml	Malayalam
mn	Mongolian
mo	Moldavian
mr	Marathi
ms	Malay
mt	Maltese
my	Burmese
na	Nauru
ne	Nepali

code	language
nl	Dutch
no	Norwegian
oc	Occitan
om	(Afan) Oromo
or	Oriya
pa	Punjabi
pl	Polish
ps	Pashto, Pushto
pt	Portuguese
qu	Quechua
rm	Rhaeto-Romance
rn	Kirundi
ro	Romanian
ru	Russian
rw	Kinyarwanda
sa	Sanskrit
sd	Sindhi
sg	Sangro
sh	Serbo-Croat-ian
si	Singhalese
sk	Slovak
sl	Slovenian
sm	Samoan
sn	Shona

code	language
so	Somali
sq	Albanian
sr	Serbian
ss	Siswati
st	Sesotho
su	Sundanese
sv	Swedish
sw	Swahili
ta	Tamil
te	Tegulu
tg	Tajik
th	Thai
ti	Tigrinya
tk	Turkmen
tl	Tagalog
tn	Setswanato Tonga
tr	Turkish
ts	Tsonga
tt	Tatar
tw	Twi
uk	Ukrainian
ur	Urdu
uz	Uzbek
vi	Vietnamese
vo	Volapuk

March 11, 1999

code	language
wo	Wolof
xh	Xhosa
yo	Yoruba
zh	Chinese
zu	Zulu

:

March 11, 1999

APPENDIX J

ISO 3166 Country Codes

country	code
AFGHANISTAN	AF
ALBANIA	AL
ALGERIA	DZ
AMERICAN SAMOA	AS
ANDORRA	AD
ANGOLA	AO
ANGUILLA	AI
ANTARCTICA	AQ
ANTIGUA AND BARBUDA	AG
ARGENTINA	AR
ARMENIA	AM
ARUBA	AW
AUSTRALIA	AU
AUSTRIA	AT
AZERBAIJAN	AZ
BAHAMAS	BS
BAHRAIN	BH
BANGLADESH	BD
BARBADOS	BB
BELARUS	BY
BELGIUM	BE

country	code
BELIZE	BZ
BENIN	BJ
BERMUDA	BM
BHUTAN	BT
BOLIVIA	BO
BOSNIA AND HERZEGOWINA	BA
BOTSWANA	BW
BOUVET ISLAND	BV
BRAZIL	BR
BRITISH INDIAN OCEAN TERRITORY	IO
BRUNEI DARUSSALAM	BN
BULGARIA	BG
BURKINA FASO	BF
BURUNDI	BI
CAMBODIA	KH
CAMEROON	CM
CANADA	CA
CAPE VERDE	CV
CAYMAN ISLANDS	KY
CENTRAL AFRICAN REPUBLIC	CF
CHAD	TD
CHILE	CL

country	code
CHINA	CN
CHRISTMAS ISLAND	CX
COCOS (KEELING) ISLANDS	CC
COLOMBIA	CO
COMOROS	KM
CONGO	CG
COOK ISLANDS	CK
COSTA RICA	CR
COTE D'IVOIRE	CI
CROATIA (local name: Hrvatska)	HR
CUBA	CU
CYPRUS	CY
CZECH REPUBLIC	CZ
DENMARK	DK
DJIBOUTI	DJ
DOMINICA	DM
DOMINICAN REPUBLIC	DO
EAST TIMOR	TP
ECUADOR	EC
EGYPT	EG
EL SALVADOR	SV
EQUATORIAL GUINEA	GQ
ERITREA	ER
ESTONIA	EE
ETHIOPIA	ET
FALKLAND ISLANDS (MALVINAS)	FK
FAROE ISLANDS	FO
FIJI	FJ
FINLAND	FI
FRANCE	FR

country	code
FRANCE, METROPOLITAN	FX
FRENCH GUIANA	GF
FRENCH POLYNESIA	PF
FRENCH SOUTHERN TERRITORIES	TF
GABON	GA
GAMBIA	GM
GEORGIA	GE
GERMANY	DE
GHANA	GH
GIBRALTAR	GI
GREECE	GR
GREENLAND	GL
GRENADA	GD
GUADELOUPE	GP
GUAM	GU
GUATEMALA	GT
GUINEA	GN
GUINEA-BISSAU	GW
GUYANA	GY
HAITI	HT
HEARD AND MC DONALD ISLANDS	HM
HONDURAS	HN
HONG KONG	HK
HUNGARY	HU
ICELAND	IS
INDIA	IN
INDONESIA	ID
IRAN (ISLAMIC REPUBLIC OF)	IR
IRAQ	IQ
IRELAND	IE

March 11, 1999

country	code
ISRAEL	IL
ITALY	IT
JAMAICA	JM
JAPAN	JP
JORDAN	JO
KAZAKHSTAN	KZ
KENYA	KE
KIRIBATI	KI
KOREA, DEMOCRATIC PEOPLE'S REPUBLIC OF	KP
KOREA, REPUBLIC OF	KR
KUWAIT	KW
KYRGYZSTAN	KG
LAO PEOPLE'S DEMOCRATIC REPUBLIC	LA
LATVIA	LV
LEBANON	LB
LESOTHO	LS
LIBERIA	LR
LIBYAN ARAB JAMAHIRIYA	LY
LIECHTENSTEIN	LI
LITHUANIA	LT
LUXEMBOURG	LU
MACAU	MO
MACEDONIA, THE FORMER YUGO- SLAV REPUBLIC OF	MK
MADAGASCAR	MG
MALAWI	MW
MALAYSIA	MY
MALDIVES	MV
MALI	ML

country	code
MALTA	MT
MARSHALL ISLANDS	MH
MARTINIQUE	MQ
MAURITANIA	MR
MAURITIUS	MU
MAYOTTE	YT
MEXICO	MX
MICRONESIA, FEDERATED STATES OF	FM
MOLDOVA, REPUBLIC OF	MD
MONACO	MC
MONGOLIA	MN
MONTSERRAT	MS
MOROCCO	MA
MOZAMBIQUE	MZ
MYANMAR	MM
NAMIBIA	NA
NAURU	NR
NEPAL	NP
NETHERLANDS	NL
NETHERLANDS ANTILLES	AN
NEW CALEDONIA	NC
NEW ZEALAND	NZ
NICARAGUA	NI
NIGER	NE
NIGERIA	NG
NIUE	NU
NORFOLK ISLAND	NF
NORTHERN MARIANA ISLANDS	MP
NORWAY	NO

country	code
OMAN	OM
PAKISTAN	PK
PALAU	PW
PANAMA	PA
PAPUA NEW GUINEA	PG
PARAGUAY	PY
PERU	PE
PHILIPPINES	PH
PITCAIRN	PN
POLAND	PL
PORTUGAL	PT
PUERTO RICO	PR
QATAR	QA
REUNION	RE
ROMANIA	RO
RUSSIAN FEDERATION	RU
RWANDA	RW
SAINT KITTS AND NEVIS	KN
SAINT LUCIA	LC
SAINT VINCENT AND THE GRENADINES	VC
SAMOA	WS
SAN MARINO	SM
SAO TOME AND PRINCIPE	ST
SAUDI ARABIA	SA
SENEGAL	SN
SEYCHELLES	SC
SIERRA LEONE	SL
SINGAPORE	SG
SLOVAKIA (Slovak Republic)	SK

country	code
SLOVENIA	SI
SOLOMON ISLANDS	SB
SOMALIA	SO
SOUTH AFRICA	ZA
SPAIN	ES
SRI LANKA	LK
ST. HELENA	SH
ST. PIERRE AND MIQUELON	PM
SUDAN	SD
SURINAME	SR
SVALBARD AND JAN MAYEN ISLANDS	SJ
SWAZILAND	SZ
SWEDEN	SE
SWITZERLAND	CH
SYRIAN ARAB REPUBLIC	SY
TAIWAN, PROVINCE OF CHINA	TW
TAJKISTAN	TJ
TANZANIA, UNITED REPUBLIC OF	TZ
THAILAND	TH
TOGO	TG
TOKELAU	TK
TONGA	TO
TRINIDAD AND TOBAGO	TT
TUNISIA	TN
TURKEY	TR
TURKMENISTAN	TM
TURKS AND CAICOS ISLANDS	TC
TUVALU	TV
UGANDA	UG

March 11, 1999

country	code
UKRAINE	UA
UNITED ARAB EMIRATES	AE
UNITED KINGDOM	GB
UNITED STATES	US
UNITED STATES MINOR OUTLYING ISLANDS	UM
URUGUAY	UY
UZBEKISTAN	UZ
VANUATU	VU
VATICAN CITY STATE (HOLY SEE)	VA
VENEZUELA	VE
VIET NAM	VN
VIRGIN ISLANDS (BRITISH)	VG
VIRGIN ISLANDS (U.S.)	VI
WALLIS AND FUTUNA ISLANDS	WF
WESTERN SAHARA	EH
YEMEN	YE
YUGOSLAVIA	YU
ZAIRE	ZR
ZAMBIA	ZM
ZIMBABWE	ZW

:

March 11, 1999

Bibliography

Adobe Technical Notes are available online at the following URL:

<http://partners.adobe.com/supportservice/devrelations/technotes.html>

Documents in the Acrobat Software Development Kit (SDK) are available at the following URL:

<http://partners.adobe.com/supportservice/devrelations/nonmember/acrosdk/main.html>

In addition, developers can find these documents on the Acrobat SDK CD, in the **Docs/OtherDocs** folder.

[1] Adobe Systems Incorporated, *PostScript Language Reference Manual, Third Edition*, Addison-Wesley, 1990, ISBN 0-201-37922-8. Reference manual describing the imaging model used in the PostScript language and the language itself.

<http://partners.adobe.com/supportservice/devrelations/PDFS/TN/PLRM.pdf>

[2] Adobe Systems Incorporated, *Supporting Data Compression in PostScript Level 2 and the Filter Operator*, Adobe Developer Support Technical Note 5115.

[3] Adobe Systems Incorporated, *Supporting the DCT Filters in PostScript Level 2*, Adobe Developer Support Technical Note 5116. Contains errata for the JPEG discussion in the *PostScript Language Reference Manual, Second Edition*. Also describes the compatibility of the JPEG implementation with various versions of the JPEG standard.

[4] Adobe Systems Incorporated, *Adobe CMap and CIDfont File Specification, version 1*, Adobe Developer Support Technical Note 5014.

[5] Adobe Systems Incorporated, *Adobe Type 1 Font Format*, Addison-Wesley, 1990, ISBN 0-201-57044-0. Explains the internal organization of a PostScript language Type 1 font program.

[6] Adobe Systems Incorporated, *Type 1 Font Format Supplement*, Adobe Developer Support Technical Note 5015.

[7] Adobe Systems Incorporated, *The Compact Font Format Specification*, Adobe Developer Support Technical Note 5176. Describes the additions made to the Type 1 font format to support Type1C fonts.

- [8] Adobe Systems Incorporated. *The Type 2 Charstring Format*. Adobe Developer Support Technical Note 5176.
- [9] Adobe Systems Incorporated, *Portable Job Ticket (PJTF) Specification*. To be published.
- [10] Adobe Systems Incorporated, *Adobe Acrobat Forms JavaScript Object Specification*. To be published.
- [11] Adobe Systems Incorporated, *PDF Public-Key Digital Signature and Encryption Specification*. This is available on the Acrobat SDK Web site.
- [12] Adobe Systems Incorporated, *Standard Element Types for Logical Structure in PDF*. This is available on the Acrobat SDK Web site.
- [13] Adobe Systems Incorporated, *Adobe Font Metrics File Format Specification*. Adobe Developer Support Technical Note 5004. Adobe Font Metrics (AFM) files are available at
<ftp://ftp.adobe.com/pub/adobe/type/mac/all/afmfiles/> and
<ftp://ftp.adobe.com/pub/adobe/type/win/all/afmfiles/>.
- [14] Adobe Systems Incorporated, *PostScript Language Document Structuring Conventions Specification Version 3.0*, Adobe Developer Support technical Note 5001.
- [15] Aho, Alfred V., John E. Hopcroft, and Jeffrey D. Ullman, *Data Structures and Algorithms*, Addison-Wesley, 1983, ISBN 0-201-00023-7. Includes a discussion of balanced trees.
- [16] Arvo, James (ed.), *Graphics Gems II*, Academic Press, 1991, ISBN 0-12-064480-0. The section “Geometrically Continuous Cubic Bézier Curves” by Hans-Peter Seidel describes the mathematics used to smoothly join two cubic Bézier curves.
- [17] Berners-Lee, T., and D. Connolly. Internet RFC 1866, *Hypertext Markup Language 2.0 Proposed Standard*. November 1995.

<URL:ftp://ds.internic.net/rfc/rfc1866.txt;type=a>

For updates, see

<http://www.w3.org/pub/WWW/MarkUp/html-spec>

- [18] Berners-Lee, T., Masinter, McCahill, and the Network Working Group. Internet RFC 1738, *Uniform Resource Locators*.

<URL:ftp://ds.internic.net/rfc/rfc1738.txt;type=a>

- [19] CCITT, *Blue Book*, Volume VII.3, 1988. ISBN 92-61-03611-2. Recommendations T.4 and T.6 are the CCITT standards for Group 3 and Group 4 facsimile encoding. This document may be purchased from Global Engineering Documents, P.O. Box 19539, Irvine, California 92713.

March 11, 1999

[20] CCITT, *Recommendation X.208: Specification of Abstract Syntax Notation One (ASN.1)*, 1988.

[21] Deutsch, L. Peter, and Jean-Loup Gailly. Internet RFC 1950, *ZLIB Compressed Data Format Specification version 3.3*.

[<URL:ftp://ds.internic.net/rfc/rfc1950.txt;type=a>](ftp://ds.internic.net/rfc/rfc1950.txt)

[22] Deutsch, L. Peter. Internet RFC 1951, *DEFLATE Compressed Data Format Specification version 1.3*.

[<URL:ftp://ds.internic.net/rfc/rfc1951.txt;type=a>](ftp://ds.internic.net/rfc/rfc1951.txt)

[23] Fielding, Network Working Group. Internet RFC 1808, *Relative Uniform Resource Locators*.

[<URL:ftp://ds.internic.net/rfc/rfc1808.txt;type=a>](ftp://ds.internic.net/rfc/rfc1808.txt)

[24] Fielding, Gettys, Mogul, Frystyk, Berners-Lee. Internet RFC 2068, *Hypertext Transfer Protocol – HTTP/1.1*, January 1997.

[<URL:ftp://ds.internic.net/rfc/rfc2068.txt;type=a>](ftp://ds.internic.net/rfc/rfc2068.txt)

[25] Foley, James D., Andries van Dam, Steven K. Feiner, and John F. Hughes, *Computer Graphics: Principles and Practice, Second Edition*, Addison-Wesley, 1990, ISBN 0-201-12110-7. Section 11.2, “Parametric Cubic Curves,” contains a description of the mathematics of cubic Bézier curves and a comparison of various types of parametric cubic curves

[26] Freed, N. and N. Borenstein, *RFC 2045: Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*, November 1996.

[<URL:ftp://ds.internic.net/rfc/rfc2045.txt;type=a>](ftp://ds.internic.net/rfc/rfc2045.txt)

[27] RFC 2046, *Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types*.

[28] Glassner, Andrew S. (ed.), *Graphics Gems*, Academic Press, 1990, ISBN 0-12-286165-5. The section “An Algorithm For Automatically Fitting Digitized Curves” by Philip J. Schneider describes an algorithm for determining the set of Bézier curves approximating an arbitrary set of user-provided points. Appendix 2 contains an implementation of the algorithm, written in the C programming language. Other sections relevant to the mathematics of Bézier curves include “Solving the Nearest-Point-On-Curve Problem” by Philip J. Schneider, “Some Properties of Bézier Curves” by Ronald Goldman, and “A Bézier Curve-Based Root-Finder” by Philip J. Schneider. The source code appearing in the appendix is available via anonymous ftp, as described in the preface to *Graphics Gems III* [31].

[29] International Color Consortium. See <http://www.color.org>. Version 3.4 of the specification for ICC profiles is available at

<ftp://sgigate.sgi.com/pub/icc/ICC34.pdf>

[30] Joint Photographic Experts Group (JPEG) "Revision 8 of the JPEG Technical Specification," ISO/IEC JTC1/SC2/WG8, CCITT SGVIII, August 14, 1990. Defines a set of still-picture grayscale and color image data compression algorithms.

[31] Kirk, David (ed.), *Graphics Gems III*, Academic Press, 1992, ISBN 0-12-409670-0 (with IBM Disk) or ISBN 0-12-409671-9 (with Macintosh disk). The section "Interpolation Using Bézier Curves" by Gershon Elber contains an algorithm for calculating a Bézier curve that passes through a user-specified set of points. The algorithm utilizes not only cubic Bézier curves, which are supported in PDF, but also higher-order Bézier curves. The appendix contains an implementation of the algorithm, written in the C programming language. All of the source code appearing in the appendix is available via anonymous ftp, as described in the preface.

[32] Microsoft Corp., *TrueType 1.0 Font Files*, Revision 1.00, May 1992.

[33] Netscape Communications Incorporated, *Netscape JavaScript Reference Manual*.

<http://developer.netscape.com/docs/manuals/communicator/jsref/index.htm>

[34] Pennebaker, W. B. and Joan L. Mitchell, *JPEG Still Image Data Compression Standard*, Van Nostrand Reinhold, 1993, ISBN 0-442-01272-1.

[35] Ron Rivest, *RFC 1321: The MD5 Message-Digest Algorithm*, April 1992.

[URL:ftp://ds.internic.net/rfc/rfc1321.txt](ftp://ds.internic.net/rfc/rfc1321.txt);type=a

[36] TrueType Conference Technical Committee, *Japanese TrueType Font Property Selection Guidelines* (TrueType Kyougikai Gijutsu Iinkai, *Nihongo TrueTypeFont Shotai Zokusei-chi Sentei Gaidorain*)

[37] Warnock, John and D. Wyatt, "A Device Independent Graphics Imaging Model for Use with Raster Devices," *Computer Graphics* (ACM SIGGRAPH), Volume 16, Number 3, July 1982. Technical background for the imaging model used in the PostScript language.

Index

Symbols

" 445
40
% 56
%%EOF 59
%PDF-1.3 56
' 445
/ 39
<< 40
>> 40

A

Acrobat Distiller 20
Aliased URLs Object 155
array 40
ASCII85Decode 45
ASCIHexDecode 44

B

B 443
b 443
B* 443
b* 443
base 14 fonts 202
BaseFont 201, 203, 204, 205
BDC 443
begin image 443
begin image data 444
begin marked content 443
BI 443
blue screen mask 255
BMC 443
bookmark 104
BT 443

C

c 443
Catalog 67
CCITT 23
CCITTFaxDecode 50
chroma-key mask 255
clip 445
closepath 443, 444, 445

cm 443
comments 56
compatibility note 15
concat 443
content stream 195, 444
cross-reference table 24, 57
current Resources dictionary 197
curveto 443, 445

D

d 443
d0 443
d1 443
Date (data structure definition) 183
DCTDecode 52
deleted object 58
Destination (data structure definition) 184
dictionary 40
direct object 53
Distiller 20
Do 443
DP 444

E

EBNF 15
Edge flag 309
EI 444
EMC 444
Encryption 62
end image 444
end marked content 444
end path 444
end text object 444
endobj 53
endstream 41
eoclip 445
eofill 443, 444
escape character 38
ET 444
EX 444
Extended Backus–Naur Form 15
extended graphics state 444

F

F 444
f 58, 444

f* 444
false 37
File specification (data structure definition) 186
file specification dictionary 188
fill 443, 444, 445
fill path 444
Flate 23
font
 multiple master Type 1 205
 subsets 204
 TrueType 203
 Type 0 207
 Type 1 199
 Type 3 206
free object 58
full file specification 186

G

G 444
g 444
GDI 20
generation number 53
glyphs 24
gradient 294
gs 444

H

h 444
Header 56

I

i 444
ID 444
implementation note 15
Incremental update 61
indirect object 37, 53
indirect reference 54
inheritance of attributes 72
ITC Zapf Dingbats font 24

J

J 444
j 444
JPEG 23

:

March 11, 1999

K

K 444

k 444

L

l 444

Lempel-Ziv-Welch compression 23

lineto 444

LZW 23

LZWDecode 45

M

M 444

m 444

moveto 444

MP 444

N

n 57, 444

name 39

named destinations 185

named resources 195

namedXObject 443

null 53

numbers 37

O

obj 53

object number 53

operator 443

P

page description 19

page label 70

page marking operators 443

Pages tree 71

path 444, 445

pattern 445

PDF consumer 9

PDF document 19

PDF file 19

PDF generator 9

PDF Writer 20

PostScript 9, 19, 443

procsets 26

Q

Q 444

q 444

QuickDraw 20

R

R 54

RC4 encryption algorithm 63

re 444

rectangle 444

Rectangle (data structure definition) 183

rendering intent 445

Resources dictionary 195

restore graphics state 444

RG 444

rg 444

ri 445

Run Length compression 23

RunLengthDecode 50

S

S 445

s 445

save graphics state 444

SC 445

sc 445

SCN 445

scn 445

section 57

separation color spaces 445

set character spacing 445

set file size 445

set font name 445

set horizontal scaling 445

set leading 445

set text matrix 445

set text rendering mode 445

set text rise 445

set word spacing 445

setcachedevice 443

setcharwidth 443

setcmykcolor 444

setcolor 445

setcolorspace 443

setdash 443
setflat 444
setgray 444
setlinecap 444
setlinejoin 444
setlinewidth 445
setmiterlimit 444
setrgbcolor 444
sh 445
shaded fill 445
shfill 445
show text 445
simple file specification 186
snapshot (Multiple Master font) 205
startxref 59
stream 41
string 37
stroke 443, 444, 445
style (TrueType fonts) 204
subsection 57
Symbol font 24

T

T* 445
Tc 445
TD 445
Td 445
text 39
text current point 445
Tf 445
TJ 445
Tj 445
TL 445
Tm 445
Tr 445
trailer 59, 59
Trapping 99
true 37
TrueType 24
Ts 445
Tw 445
Type 41
Type (restrictions on the use of the key name) 14
Type 3 font 443
Tz 445

U

undefined object 54
undefined operators 444
Unicode Byte Order marker 39
URL 194

V

v 445
viewer preferences dictionary 69

W

W 445
w 445
W* 445

X

xref 57

Y

y 445

:

March 11, 1999



Colophon

This book was produced electronically using Adobe FrameMaker[®] on Macintosh, Dell Dimension, and Sun[™] SPARCstation[®] computers. Art was produced using Adobe Photoshop[®], Adobe Illustrator[®], and Adobe FrameMaker[®]. The final version is a PDF 1.2 file produced by Adobe Acrobat Distiller from PostScript created by the AdobePS[™] print driver.

Authors—Tim Bienz, Richard Cohn, and Jim Meehan

Key Contributors—Nabeel Al-Shamma, Ed Taft

Reviewers—Nabeel Al-Shamma, David Gelpman, Sherri Nichols, Paul Rovner, Alan Wootton, Jim Pravetz, Dick Sites, Terry O'Donnell and numerous others at Adobe Systems; and L. Peter Deutsch of Aladdin Enterprises.

Editing—Gary Staas, Diana Wynne, Karl Schmidtman

Project Management—Tim Oey, Bob Wulff

:

March 11, 1999