

# **Multimedia Programming Interface and Data Specifications 1.0**

---

**Issued as a joint design by IBM Corporation and Microsoft Corporation**

**August 1991**

This document describes the programming interfaces and data specifications for multimedia that are common to both OS/2 and Windows environments. These specifications may be enhanced to incorporate new technologies or modified based on customer feedback and, as such, specifications incorporated into any final product may vary.

Microsoft is a registered trademark, and Windows is a trademark of Microsoft Corp.

IBM and OS/2 are registered trademarks of International Business Machines Corporation.

# Contents

---

## Contents

### **Chapter 1 ..... Overview of Multimedia File Formats**

Resource Interchange File Format.....	1-1
Multimedia File Formats.....	1-1
Media Control Interface.....	1-2
Registering Multimedia Formats.....	1-2

### **Chapter 2 ..... Resource Interchange File Format**

About the RIFF Tagged File Format.....	2-1
Notation Conventions.....	2-1
Chunks.....	2-2
RIFF Forms.....	2-3
Defining and Registering RIFF Forms.....	2-3
Registered Form and Chunk Types.....	2-4
Unregistered (Form-Specific) Chunk Types.....	2-4
Notation for Representing Sample RIFF Files.....	2-5
Basic Notation for Representing RIFF Files.....	2-5
Escape Sequences for Four-Character Codes and String Chunks.....	2-7
Extended Notation for Representing RIFF Form Definitions.....	2-8
Atomic Labels.....	2-10
A Sample RIFF Form Definition and RIFF Form.....	2-11
Storing Strings in RIFF Chunks.....	2-12
NULL-Terminated String (ZSTR) Format.....	2-12
String Table Format.....	2-13

NULL-Terminated, Byte Size Prefix String (BZSTR) Series.....	2-13
Multiline String Format .....	2-13
Choosing a Storage Method.....	2-13
LIST Chunk .....	2-14
INFO List Chunk .....	2-14
CSET (Character Set) Chunk.....	2-16
Country Codes .....	2-16
Language and Dialect Codes.....	2-17
JUNK (Filler) Chunk .....	2-18
Compound File Structure.....	2-18
Structural Overview.....	2-19
Compound File Table of Contents (CTOC) Chunk .....	2-19
Structural Overview .....	2-19
Header Information.....	2-21
Parameter Table Definition.....	2-21
Header Parameter Table.....	2-22
CTOC Table Entries.....	2-22
Usage Codes for Extra Header and Extra Entry Fields .....	2-24
Compression of Compound File Elements.....	2-26
Compound File Element Group (CGRP) Chunk.....	2-27
Placement of the CTOC and CGRP Chunks .....	2-27

**Chapter 3 ..... Multimedia File F**

Bundle File Format .....	3-1
Device Independent Bitmap File Format .....	3-1
Overview of DIB Structure .....	3-2
Bitmap File Header.....	3-2
Bitmap Information Header.....	3-3
Information Header Structures .....	3-4

Bitmap Color Table .....	3-6
Color Table Structure.....	3-6
Order of Colors.....	3-6
Field Descriptions .....	3-6
Locating the Color Table .....	3-7
Interpreting the Color Table .....	3-7
Bitmap Data .....	3-8
Windows 3.0 Bitmap Compression Formats.....	3-8
Compression of 8-Bit-Per-Pixel DIBs.....	3-8
Compression of 4-Bit-Per-Pixel DIBs.....	3-9
RIFF Device-Independent Bitmap File Format .....	3-10
Simple RDIB Format.....	3-10
Extended RDIB Format.....	3-10
Bitmap Header Chunk.....	3-11
Transitional Compression .....	3-16
CCC Compression.....	3-17
Palette Chunk .....	3-17
External Palette Chunk.....	3-17
Bitmap Data Chunk.....	3-17
MIDI and RIFF MIDI File Formats .....	3-18
Palette File Format .....	3-18
Simple PAL Format.....	3-18
Extended PAL Format.....	3-19
Rich Text Format (RTF) .....	3-22
Waveform Audio File Format (WAVE) .....	3-22
WAVE Format Chunk.....	3-22
WAVE Format Categories .....	3-23
Pulse Code Modulation (PCM) Format.....	3-24

Storage of WAVE Data .....	3-26
FACT Chunk.....	3-26
Cue-Points Chunk.....	3-27
Examples of File Position Values.....	3-28
Playlist Chunk.....	3-29
Associated Data Chunk .....	3-29
Label and Note Information.....	3-30
Text with Data Length Information.....	3-30
Embedded File Information .....	3-31

**Chapter 4 ..... Media Control In**

MCI Command Strings .....	4-1
Example of MCI Command Use.....	4-2
Categories of MCI Command Strings.....	4-2
Command Syntax Conventions.....	4-3
System Commands .....	4-3
Required Commands .....	4-3
Basic Commands .....	4-4
Extended Commands.....	4-4
Extended Commands Reserved for Future Use .....	4-4
Creating a Command String.....	4-5
About MCI Device Types .....	4-6
Using MCI Command Strings .....	4-6
Opening a Device .....	4-6
Opening Simple Devices.....	4-7
Opening Compound Devices.....	4-7
Using the Shareable Flag .....	4-8
Using the Alias Flag.....	4-8
Opening New Device Elements.....	4-8

Closing a Device.....	4-8
Shortcuts and Variations for MCI Commands .....	4-9
Using All as a Device Name.....	4-9
Combining the Device Type and Device Element Name.....	4-9
Automatic Open .....	4-9
Automatic Close.....	4-9
Using Wait and Notify Flags.....	4-10
Using the Notify Flag.....	4-10
Obtaining Information From MCI Devices .....	4-11
The Play Command .....	4-11
Stop, Pause, and Resume Commands .....	4-11
MCI System Commands .....	4-12
Required Commands for All Devices.....	4-13
Basic Commands for Specific Device Types.....	4-14
CD Audio (Redbook) Commands .....	4-17
MIDI Sequencer Commands .....	4-20
Videodisc Player Commands.....	4-25
Waveform Audio Commands.....	4-29

## Chapter 1

# Overview of Multimedia Specifications

---

This document describes the file format and control interface specifications for multimedia. These specifications allow developers to use common file format and device control interfaces.

## Resource Interchange File Format

The Resource Interchange File Format (RIFF), a tagged file structure, is a general specification upon which many file formats can be defined. The main advantage of RIFF is its extensibility; file formats based on RIFF can be future-proofed, as format changes can be ignored by existing applications.

The RIFF file format is suitable for the following multimedia tasks:

- Playing back multimedia data
- Recording multimedia data
- Exchanging multimedia data between applications and across platforms

Chapter 2, “Resource Interchange File Format,” describes the RIFF format.

## Multimedia File Formats

A number of RIFF-based and non-RIFF file formats have been defined for the storage of multimedia data. Chapter 3, “Multimedia File Formats,” describes the following file formats:

- Bundle File Format
- Device-Independent Bitmap (DIB) and RIFF DIB file formats
- Musical Instrument Digital Interface (MIDI) and RIFF MIDI file formats
- Palette File Format
- Rich Text File Format
- Waveform Audio File Format



# Media Control Interface

The Media Control Interface (MCI) is a high-level control mechanism that provides a device-independent interface to multimedia devices and resource files.

The Media Control Interface (MCI) provides a command set for playing and recording multimedia devices and resource files. Developers creating multimedia applications are encouraged to use this high-level command interface rather than the low-level functions specific to each platform. The MCI command set acts as a platform-independent layer that sits between multimedia applications and the underlying system software.

The MCI command set is extensible in two ways:

- Developers can incorporate new multimedia devices and file formats in the MCI command set by creating new MCI drivers to interpret the commands.
- New commands and command options can be added to support special features or functions required by new multimedia devices or file formats.

Using MCI, an application can control multimedia devices using simple command strings like open, play, and close. The MCI command strings provide a generic interface to different multimedia devices, reducing the number of commands a developer needs to learn. A multimedia application might even accept MCI commands from an end user and pass them unchanged to the MCI driver, which parses the command and performs the appropriate action.

Chapter 3, “Media Control Interface,” describes MCI and its command set in detail.

## Registering Multimedia Formats

This document discusses several multimedia codes and formats that require registration. These multimedia elements include the following:

- Compression techniques
- RIFF form types, chunk IDs, and list types
- Compound-file usage codes
- Waveform audio format codes

To register these multimedia elements, request a *Multimedia Developer Registration Kit* from the following group:

Microsoft Corporation  
Multimedia Systems Group  
Product Marketing  
One Microsoft Way  
Redmond, WA 98052-6399

The *Multimedia Developer Registration Kit* also lists currently defined multimedia elements.

## Chapter 2

# Resource Interchange File Format

---

The Resource Interchange File Format (RIFF) is a tagged file structure developed for use on multimedia platforms. This chapter defines RIFF and describes the file structures based on RIFF. If your application requires a new file format, you should define it using the RIFF tagged file structure described in this chapter.

## About the RIFF Tagged File Format

RIFF (Resource Interchange File Format) is the tagged file structure developed for multimedia resource files. The structure of a RIFF file is similar to the structure of an Electronic Arts IFF file. RIFF is not actually a file format itself (since it does not represent a specific kind of information), but its name contains the words “interchange file format” in recognition of its roots in IFF. Refer to the EA IFF definition document, *EA IFF 85 Standard for Interchange Format Files*, for a list of reasons to use a tagged file format.

RIFF has a counterpart, RIFX, that is used to define RIFF file formats that use the Motorola integer byte-ordering format rather than the Intel format. A RIFX file is the same as a RIFF file, except that the first four bytes are ‘RIFX’ instead of ‘RIFF’, and integer byte ordering is represented in Motorola format.

## Notation Conventions

The following table lists some of the notation conventions used in this document. Further conventions and the notation for documenting RIFF forms are presented later in the document in the section “Notation for Representing Sample RIFF Files.”

---

<b>Notation</b>	<b>Description</b>
<code>&lt;element label&gt;</code>	RIFF file element with the label “element label”
<code>&lt;element label: TYPE&gt;</code>	RIFF file element with data type “TYPE”
<code>[&lt;element label&gt;]</code>	Optional RIFF file element
<code>&lt;element label&gt;...</code>	One or more copies of the specified element
<code>[&lt;element label&gt;]...</code>	Zero or more copies of the specified element

---

## Chunks

The basic building block of a RIFF file is called a *chunk*. Using C syntax, a chunk can be defined as follows:

```
typedef unsigned long DWORD;
typedef unsigned char BYTE;

typedef DWORD FOURCC;           // Four-character code

typedef FOURCC CKID;           // Four-character-code chunk identifier
typedef DWORD CKSIZE;          // 32-bit unsigned size value

typedef struct {               // Chunk structure
    CKID      ckID;             // Chunk type identifier
    CKSIZE    ckSize;           // Chunk size field (size of ckData)
    BYTE      ckData[ckSize];   // Chunk data
} CK;
```

A FOURCC is represented as a sequence of one to four ASCII alphanumeric characters, padded on the right with blank characters (ASCII character value 32) as required, with no embedded blanks.

For example, the four-character code 'FOO' is stored as a sequence of four bytes: 'F', 'O', 'O', ' ' in ascending addresses. For quick comparisons, a four-character code may also be treated as a 32-bit number.

The three parts of the chunk are described in the following table:

---

Part	Description
<b>ckID</b>	A four-character code that identifies the representation of the chunk data <b>data</b> . A program reading a RIFF file can skip over any chunk whose chunk ID it doesn't recognize; it simply skips the number of bytes specified by <b>ckSize</b> plus the pad byte, if present.
<b>ckSize</b>	A 32-bit unsigned value identifying the size of <b>ckData</b> . This size value does not include the size of the <b>ckID</b> or <b>ckSize</b> fields or the pad byte at the end of <b>ckData</b> .
<b>ckData</b>	Binary data of fixed or variable size. The start of <b>ckData</b> is word-aligned with respect to the start of the RIFF file. If the chunk size is an odd number of bytes, a pad byte with value zero is written after <b>ckData</b> . Word aligning improves access speed (for chunks resident in memory) and maintains compatibility with EA IFF. The <b>ckSize</b> value does not include the pad byte.

---

We can represent a chunk with the following notation (in this example, the **ckSize** and pad byte are implicit):

```
<ckID> ( <ckData> )
```

Two types of chunks, the 'LIST' and 'RIFF' chunks, may contain nested chunks, or subchunks. These special chunk types are discussed later in this document. All other chunk types store a single element of binary data in **<ckData>**.

## RIFF Forms

A RIFF form is a chunk with a ‘RIFF’ chunk ID. The term also refers to a file format that follows the RIFF framework. The following is the current list of registered RIFF forms. Each is described in Chapter 3, “Multimedia File Formats.”

---

<b>Form Type</b>	<b>Description</b>
PAL	RIFF Palette Format
RDIB	RIFF Device Independent Bitmap Format
RMID	RIFF MIDI Format
RMMP	RIFF Multimedia Movie File Format
WAVE	Waveform Audio Format

---

Using the notation for representing a chunk, a RIFF form looks like the following:

```
RIFF ( <formType> <ck>... )
```

The first four bytes of a RIFF form make up a chunk ID with values ‘R’, ‘I’, ‘F’, ‘F’. The **ckSize** field is required, but for simplicity it is omitted from the notation.

The first DWORD of chunk data in the ‘RIFF’ chunk (shown above as **<formType>**) is a four-character code value identifying the data representation, or *form type*, of the file. Following the form-type code is a series of subchunks. Which subchunks are present depends on the form type. The definition of a particular RIFF form typically includes the following:

- A unique four-character code identifying the form type
- A list of mandatory chunks
- A list of optional chunks
- Possibly, a required order for the chunks

## Defining and Registering RIFF Forms

The form-type code for a RIFF form must be unique. To guarantee this uniqueness, you must register any new form types before release. See “Registering Multimedia Formats” in Chapter 1, “Overview of Multimedia Specifications,” for information on registering RIFF forms.

Like RIFF forms, RIFX forms must also be registered. Registering a RIFF form does not automatically register the RIFX counterpart. No RIFX form types are currently defined.

## Registered Form and Chunk Types

By convention, the form-type code for registered form types contains only digits and uppercase letters. Form-type codes that are all uppercase denote a registered, unique form type. Use lowercase letters for temporary or prototype chunk types.

Certain chunk types are also globally unique and must also be registered before use. These registered chunk types are not specific to a certain form type; they can be used in any form. If a registered chunk type can be used to store your data, you should use the registered chunk type rather than define your own chunk type containing the same type of information.

For example, a chunk with chunk ID 'INAM' always contains the name or title of a file. Also, within all RIFF files, filenames or titles are contained within chunks with ID 'INAM' and have a standard data format.

## Unregistered (Form-Specific) Chunk Types

Chunk types that are used only in a certain form type use a lowercase chunk ID. A lowercase chunk ID has specific meaning only within the context of a specific form type. After a form designer is allocated a registered form type, the designer can choose lowercase chunk types to use within that form. See "Registering Multimedia Formats" in Chapter 1, "Overview of Multimedia Specifications," for information on registering form types.

For example, a chunk with ID 'scln' inside one form type might contain the "number of scan lines." Inside some other form type, a chunk with ID 'scln' might mean "secondary lambda number."

# Notation for Representing Sample RIFF Files

RIFF is a binary format, but it is easier to comprehend an ASCII representation of a RIFF file. This section defines a standard notation used to present samples of various types of RIFF files. If you define a RIFF form, we urge you to use this notation in any file format samples you provide in your documentation.

## Basic Notation for Representing RIFF Files

The following table summarizes the elements of the RIFF notation required for representing sample RIFF files:

---

Notation	Description														
<b>&lt;ckID&gt; (&lt;ckData&gt;)</b>	<p>The chunk with ID <b>&lt;ckID&gt;</b> and data <b>&lt;ckData&gt;</b>. As previously described, <b>&lt;ckID&gt;</b> is a four-character code which may be enclosed by single quotes for emphasis.</p> <p>For example, the following notation describes a 'RIFF' chunk with a form type of 'QRST'. The data portion of this chunk contains a 'FOO' subchunk.</p> <pre>RIFF('QRST' FOO(17 23))</pre> <p>The following example describes an 'ICOP' chunk containing the string "Copyright Encyclopedia International.":</p> <pre>'ICOP' ("Copyright Encyclopedia International."Z)</pre>														
<b>&lt;number&gt;[&lt;modifier&gt;]</b>	<p>A number in Intel format, where <b>&lt;number&gt;</b> is an optional sign (+ or -) followed by one or more digits and modified by the optional <b>&lt;modifier&gt;</b>. Valid <b>&lt;modifier&gt;</b> values follow:</p> <hr/> <table><thead><tr><th>Modifier</th><th>Meaning</th></tr></thead><tbody><tr><td>None</td><td>16-bit number in decimal format</td></tr><tr><td>H</td><td>16-bit number in hexadecimal format</td></tr><tr><td>C</td><td>8-bit number in decimal format</td></tr><tr><td>CH</td><td>8-bit number in hexadecimal format</td></tr><tr><td>L</td><td>32-bit number in decimal format</td></tr><tr><td>LH</td><td>32-bit number in hexadecimal format</td></tr></tbody></table> <hr/>	Modifier	Meaning	None	16-bit number in decimal format	H	16-bit number in hexadecimal format	C	8-bit number in decimal format	CH	8-bit number in hexadecimal format	L	32-bit number in decimal format	LH	32-bit number in hexadecimal format
Modifier	Meaning														
None	16-bit number in decimal format														
H	16-bit number in hexadecimal format														
C	8-bit number in decimal format														
CH	8-bit number in hexadecimal format														
L	32-bit number in decimal format														
LH	32-bit number in hexadecimal format														

Several examples follow:

```
0
65535
-1
0L
4a3c89HL
-1C
```

21HC

Note that -1 and 65535 represent the same value. The application reading this file must know whether to interpret the number as signed or unsigned.

'<chars>'

A four-character code (32-bit quantity) consisting of a sequence of zero to four ASCII characters <chars> in the given order. If <chars> is less than four characters long, it is implicitly padded on the right with blanks. Two single quotes is equivalent to four blanks. Examples follow.

```
'RIFF'  
'xyz'  
' '
```

<chars> can include escape sequences, which are combinations of characters introduced by a backslash (\) and used to represent other characters. Escape sequences are listed in the following section.

"<string>"[<modifier>]

The sequence of ASCII characters contained in <string> and modified by the optional modifier <modifier>. The quoted text can include any of the escape sequences listed in the following section. Valid <modifier> values follow:

---

Modifier	Meaning
none	No NULL terminator or size prefix.
Z	String is NULL-terminated
B	String has an 8-bit (byte) size prefix
W	String has a 16-bit (word) size prefix
BZ	String has a byte-size prefix and is NULL-terminated
WZ	String has a word-size prefix and is NULL-terminated

---

NULL-terminated means that the string is followed by a character with ASCII value 0. A size prefix is an unsigned integer, stored as a byte or a word in Intel format preceding the string characters, that specifies the length of the string. In the case of strings with BZ or WZ modifiers, the size prefix specifies the size of the string without the terminating NULL.

The various string formats referred to above are discussed in "Storing Strings in RIFF Chunks," following later in this section., +

Examples follow:

```
"No prefix, no NULL terminator"  
"No prefix, NULL terminator"Z  
"Byte prefix, NULL terminator"BZ
```

---

## Escape Sequences for Four-Character Codes and String Chunks

The following escape sequences can be used in four-character codes and string chunks:

---

<b>Escape Sequence</b>	<b>ASCII Value</b>	<b>Description</b>
<code>\n</code>	10	Newline character
<code>\t</code>	9	Horizontal tab character
<code>\b</code>	8	Backspace character
<code>\r</code>	13	Carriage return character
<code>\f</code>	12	Form feed character
<code>\\</code>	92	Backslash
<code>\'</code>	39	Single quote
<code>\"</code>	34	Double quote
<code>\ddd</code>	Octal <i>ddd</i>	Arbitrary character

---



## Extended Notation for Representing RIFF Form Definitions

To unambiguously define the structure of new RIFF forms, document the RIFF form using the basic notation along with the following extended notation:

---

Notation	Description
<b>&lt;name&gt;</b>	<p>A label that refers to some element of the file, where <b>&lt;name&gt;</b> is the name of the label. Examples follow:</p> <pre>&lt;NAME-ck&gt; &lt;GOBL-form&gt; &lt;bitmap-bits&gt; &lt;foo&gt;</pre> <p>Conventionally, a label that refers to a chunk is named <b>&lt;ckID-ck&gt;</b>, where 'ckID' is the chunk ID. Similarly, a label that refers to a RIFF form is named <b>&lt;formType-form&gt;</b>, where "formType" is the name of the form's type.</p>
<b>&lt;name&gt; → elements</b>	<p>The actual data represented by <b>&lt;name&gt;</b> is defined as <b>elements</b>.</p> <p>This states that <b>&lt;name&gt;</b> is an abbreviation for <b>elements</b>, where <b>elements</b> is a sequence of other labels and literal data. An example follows:</p> <pre>&lt;GOBL-form&gt; → RIFF ( 'GOBL' &lt;form-data&gt; )</pre> <p>This example defines label <b>&lt;GOBL-form&gt;</b> as representing a RIFF form with chunk ID 'GOBL' and data equal to <b>&lt;form-data&gt;</b>, where <b>&lt;form-data&gt;</b> is a label that would be defined in another rule. Note that a label may represent any data, not just a RIFF chunk or form.</p> <p><i>Note:</i> A number of atomic labels are defined in the section "Atomic Labels" later in this document. These labels refer to primitive data types.</p>
<b>&lt;name:type&gt;</b>	<p>This is the same as <b>&lt;name&gt;</b>, but it also defines <b>&lt;name&gt;</b> to be equivalent to <b>&lt;type&gt;</b>. This notation obviates the following rule:</p> <pre>&lt;name&gt; → &lt;type&gt;</pre> <p>This allows you to give a symbolic name to an element of a file format and to specify the element data type. An example follows:</p> <pre>&lt;xyz-coordinate&gt; → &lt;x:INT&gt; &lt;y:INT&gt; &lt;z:INT&gt;</pre> <p>This defines <b>&lt;xyz-coordinate&gt;</b> to consist of three parts concatenated together: <b>&lt;x&gt;</b>, <b>&lt;y&gt;</b>, and <b>&lt;z&gt;</b>. The definition also specifies that <b>&lt;x&gt;</b>, <b>&lt;y&gt;</b>, and <b>&lt;z&gt;</b> are integers. This notation is equivalent to the following:</p> <pre>&lt;xyz-coordinate&gt; → &lt;x&gt; &lt;y&gt; &lt;z&gt; &lt;x&gt; → &lt;INT&gt; &lt;y&gt; → &lt;INT&gt; &lt;z&gt; → &lt;INT&gt;</pre>
<b>[elements]</b>	<p>An optional sequence of labels and literal data. Surrounded by square brackets, it may be considered an element itself. An example follows:</p> <pre>&lt;FOO-form&gt; → RIFF( 'FOO' [ &lt;header-ck&gt; ] &lt;data-ck&gt; )</pre>

This example defines form “FOO” with an optional header chunk followed by a mandatory data chunk.

**e1 | e2 | ... | eN**

Exactly one of the listed elements must be present. An example follows:

`<hdr-ck> → hdr(<hdr-x> | <hdr-y> | <hdr-z>)`

This example defines the ‘hdr’ chunk's data as containing one of <hdr-x>, <hdr-y>, or <hdr-z>.

**element...**

One or more occurrences of **element** may be present. An ellipsis has this meaning only if it follows an element; in cases such as “e1 | e2 | ... | eN,” the ellipsis has its ordinary English meaning. If there is any possibility of confusion, an ellipsis should only be used to indicate one or more occurrences. An example follows:

`<data-ck> → data(<count:INT> <item:INT>...)`

This example defines the data of the ‘data’ chunk to contain an integer <count>, followed by one or more occurrences of the integer <item>.

**[element]...**

Zero or more occurrences of **element** may be present. An example follows.

`<data-ck> → data(<count:INT> [<item:INT>]...)`

This example defines the data of the ‘data’ chunk to contain an integer <count> followed by zero or more occurrences of an integer <item>.

**{elements}**

The group of elements within the braces should be considered a single element. An example follows:

`<blog> → <this> | {<that> | <other>}...`

This example defines <blog> to be either <this> or one or more occurrences of <that> or <other>, intermixed in any way. Contrast this with the following example:

`<blog> → <this> | <that> | <other>...`

This example defines <blog> to be either <this> or <that> or one or more occurrences of <other>.

### **struct { ...} name**

A structure defined using C syntax. This can be used instead of a sequence of labels if a C header (include) file is available that defines the structure. The label used to refer to the structure should be the same as the structure's typedef name. An example follows:

```
<3D_POINT> → struct {  
    INT x;           // x-coordinate  
    INT y;           // y-coordinate  
    INT z;           // z-coordinate  
} 3D_POINT
```

Wherever possible, the types used in the structure should be the types listed in the following section, "Atomic Labels," because these types are more portable than C types such as int. The structure fields are assumed to be present in the file in the order given, with no padding or forced alignment.

Unless the RIFF chunk ID is 'RIFX', integer byte ordering is assumed to be in Intel format.

### **// comment**

An explanatory comment to a rule. An example follows:

```
<weekend> → 'Sat' | 'Sun'           // Four-character code  
                                           // for day
```

---

## **Atomic Labels**

The following are atomic labels, which are labels that refer to primitive data types. Where available, the equivalent Microsoft C data type is also listed.

---

<b>Label</b>	<b>Meaning</b>	<b>MS C Type</b>
<CHAR>	8-bit signed integer	signed char
<BYTE>	8-bit unsigned quantity	unsigned char
<INT>	16-bit signed integer in Intel format	signed int
<WORD>	16-bit unsigned quantity in Intel format	unsigned int
<LONG>	32-bit signed integer in Intel format	signed long
<DWORD>	32-bit unsigned quantity in Intel format	unsigned long
<FLOAT>	32-bit IEEE floating point number	float
<DOUBLE>	64-bit IEEE floating point number	double
<STR>	String (a sequence of characters)	
<ZSTR>	NULL-terminated string	
<BSTR>	String with byte (8-bit) size prefix	

<WSTR>	String with word (16-bit) size prefix
<BZSTR>	NULL-terminated string with byte size prefix
<WZSTR>	NULL-terminated string with word size prefix

---

NULL-terminated means that the string is followed by a character with ASCII value 0.

A size prefix is an unsigned integer, stored as a byte or a word in Intel format, that specifies the length of the string. In the case of strings with BZ or WZ modifiers, the size prefix specifies the size of the string without the terminating NULL.

**Note:** The WINDOWS.H header file defines the C types BYTE, WORD, LONG, and DWORD. These types correspond to labels <BYTE>, <WORD>, <LONG>, and <DWORD>, respectively.

## A Sample RIFF Form Definition and RIFF Form

The following example defines <GOBL-form>, the hypothetical RIFF form of type 'GOBL'. To fully document a new RIFF form definition, a developer would also provide detailed descriptions of each file element, including the semantics of each chunk and sample files documented using the standard notation.

```

<GOBL-form> → RIFF( 'GOBL' // RIFF form header
                [<org-ck>] // Origin chunk (default (0,0,0))
                <obj-list> // Series of graphical objects

<org-ck> → org( <origin:3D_POINT> ) // Object-list origin

// An object is a:
<obj-list> → LIST( 'obj' { <sqr-ck> | // square,
                        <circ-ck> | // circle,
                        <poly-ck> }... ) // or polygon

<sqr-ck> → sqr( <pt1:3D_POINT> // one vertex
                <pt2:3D_POINT> // another vertex
                <pt3:3D_POINT> ) // a third vertex

<circ-ck> → circ( <center:3D_POINT> // Center of circle
                  <circumPt:3D_POINT> ) // Point on circumference

<poly-ck> → poly( <pt:3D_POINT>... ) // List of points in a polygon

<3D_POINT> → struct // Defined in "gobl.h"
{ INT x; // x-coordinate
  INT y; // y-coordinate
  INT z; // z-coordinate
} 3D_POINT

```

### Sample RIFF Form

The following sample RIFF form adheres to the form definition for form type GOBL. The file contains three subchunks:

- An 'INFO' list

- An 'org' chunk
- An 'obj' chunk

The 'INFO' list and 'org' chunk each have two subchunks. The 'INFO' list is a registered global chunk that can be used within any RIFF file. The 'INFO' list is described in the 'INFO List Chunk,' later in this chapter.

Since the definition of the GOBL form does not refer to the INFO chunk, software that expects only 'org' and 'obj' chunks in a GOBL form would ignore the unknown 'INFO' chunk.

```
RIFF( 'GOBL'
    LIST('INFO'          // INFO list containing filename and copyright
        INAM("A House"Z)
        ICOP("(C) Copyright Encyclopedia International 1991"Z)
    )

    org(2, 0, 0)        // Origin of object list

    LIST('obj'          // Object list containing two polygons
        poly(0,0,0 2,0,0 2,2,0, 1,3,0, 0,2,0)
        poly(0,0,5 2,0,5 2,2,5, 1,3,5, 0,2,5)
    )
) // End of form
```

## Storing Strings in RIFF Chunks

This section lists methods for storing text strings in RIFF chunks. While these guidelines may not make sense for all applications, you should follow these conventions if you must make an arbitrary decision regarding string storage.

### NULL-Terminated String (ZSTR) Format

A NULL-terminated string (ZSTR) consists of a series of characters followed by a terminating NULL character. The ZSTR is better than a simple character sequence (STR) because many programs are easier to write if strings are NULL-terminated. ZSTR is preferred to a string with a size prefix (BSTR orWSTR) because the size of the string is already available as the **<ckSize>** value, minus one for the terminating NULL character.

## String Table Format

In a string table, all strings used in a structure are stored at the end of the structure in packed format. The structure includes fields that specify the offsets from the beginning of the string table to the individual strings. An example follows:

```
typedef struct
{
    INT     iWidgetNumber;    // the widget number
    WORD    offszWidgetName;  // an offset to a string in <rgchStrTab>
    WORD    offszWidgetDesc;  // an offset to a string in <rgchStrTab>
    INT     iQuantity;        // how many widgets
    CHAR    rgchStrTab[1];    // string table (allocate as large as needed)
} WIDGET;
```

If multiple chunks within the file need to reference variable-length strings, you can store the strings in a single chunk that acts as a string table. The chunks that refer to the strings contain offsets relative to the beginning of the data part of the string table chunk.

## NULL-Terminated, Byte Size Prefix String (BZSTR) Series

In a BZSTR series, a series of strings is stored in packed format. Each string is a BZSTR, with a byte size prefix and a NULL terminator. This format retains the ease-of-use characteristics of the ZSTR while providing the string size, allowing the application to quickly skip unneeded strings.

## Multiline String Format

When storing multiline strings, separate lines with a carriage return/line feed pair (ASCII 13/ASCII 10 pair). Although applications vary in their requirements for new line symbols (carriage return only, line feed only, or both), it is generally easier to strip out extra characters than to insert extra ones. Inserting characters might require reallocating memory blocks or pre-scanning the chunk before allocating memory for it.

## Choosing a Storage Method

The following lists guidelines for deciding which storage method is appropriate for your application.

---

<b>Usage</b>	<b>Recommended Format</b>
Chunk data contains nothing except a string	ZSTR (NULL-terminated string) format.
Chunk data contains a number of fields, some of which are variable-length strings	String-table format
Multiple chunks within the file need to reference variable-length strings	String-table format
Chunk data stores a sequence of strings, some of which the application may want to skip	BZSTR (NULL-terminated string with byte size prefix) series

---

## LIST Chunk

A LIST chunk contains a list, or ordered sequence, of subchunks. A LIST chunk is defined as follows:

```
LIST( <list-type> [<chunk>]... )
```

The **<list-type>** is a four-character code that identifies the contents of the list.

If an application recognizes the list type, it should know how to interpret the sequence of subchunks. However, since a LIST chunk may contain only subchunks (after the list type), an application that does not know about a specific list type can still walk through the sequence of subchunks.

Like chunk IDs, list types must be registered, and an all-lowercase list type has meaning relative to the form that contains it. See “Registering Multimedia Formats” in Chapter 1, “Overview of Multimedia Specifications,” for information on registering list types.

## INFO List Chunk

The ‘INFO’ list is a registered global form type that can store information that helps identify the contents of the chunk. This information is useful but does not affect the way a program interprets the file; examples are copyright information and comments. An ‘INFO’ list is a ‘LIST’ chunk with list type ‘INFO’. The following shows a sample ‘INFO’ list chunk:

```
LIST('INFO'  INAM("Two Trees"Z)
        ICMT("A picture for the opening screen"Z) )
```

An ‘INFO’ list should contain only the following chunks. New chunks may be defined, but an application should ignore any chunk it doesn't understand. The chunks listed below may only appear in an ‘INFO’ list. Each chunk contains a ZSTR, or null-terminated text string.

---

Chunk ID	Description
IARL	<i>Archival Location.</i> Indicates where the subject of the file is archived.
IART	<i>Artist.</i> Lists the artist of the original subject of the file. For example, “Michaelangelo.”
ICMS	<i>Commissioned.</i> Lists the name of the person or organization that commissioned the subject of the file. For example, “Pope Julian II.”
ICMT	<i>Comments.</i> Provides general comments about the file or the subject of the file. If the comment is several sentences long, end each sentence with a period. Do <i>not</i> include newline characters.
ICOP	<i>Copyright.</i> Records the copyright information for the file. For example, “Copyright Encyclopedia International 1991.” If there are multiple copyrights, separate them by a semicolon followed by a space.
ICRD	<i>Creation date.</i> Specifies the date the subject of the file was created. List dates in year-month-day format, padding one-digit months and days with a zero on the left. For example, “1553-05-03” for May 3, 1553.

ICRP	<i>Cropped.</i> Describes whether an image has been cropped and, if so, how it was cropped. For example, “lower right corner.”
IDIM	<i>Dimensions.</i> Specifies the size of the original subject of the file. For example, “8.5 in h, 11 in w.”
IDPI	<i>Dots Per Inch.</i> Stores dots per inch setting of the digitizer used to produce the file, such as “300.”
IENG	<i>Engineer.</i> Stores the name of the engineer who worked on the file. If there are multiple engineers, separate the names by a semicolon and a blank. For example, “Smith, John; Adams, Joe.”
IGNR	<i>Genre.</i> Describes the original work, such as, “landscape,” “portrait,” “still life,” etc.
IKEY	<i>Keywords.</i> Provides a list of keywords that refer to the file or subject of the file. Separate multiple keywords with a semicolon and a blank. For example, “Seattle; aerial view; scenery.”
ILGT	<i>Lightness.</i> Describes the changes in lightness settings on the digitizer required to produce the file. Note that the format of this information depends on hardware used.
IMED	<i>Medium.</i> Describes the original subject of the file, such as, “computer image,” “drawing,” “lithograph,” and so forth.
INAM	<i>Name.</i> Stores the title of the subject of the file, such as, “Seattle From Above.”
IPLT	<i>Palette Setting.</i> Specifies the number of colors requested when digitizing an image, such as “256.”
IPRD	<i>Product.</i> Specifies the name of the title the file was originally intended for, such as “Encyclopedia of Pacific Northwest Geography.”
ISBJ	<i>Subject.</i> Describes the contents of the file, such as “Aerial view of Seattle.”
ISFT	<i>Software.</i> Identifies the name of the software package used to create the file, such as “Microsoft WaveEdit.”
ISHP	<i>Sharpness.</i> Identifies the changes in sharpness for the digitizer required to produce the file (the format depends on the hardware used).
ISRC	<i>Source.</i> Identifies the name of the person or organization who supplied the original subject of the file. For example, “Trey Research.”
ISRF	<i>Source Form.</i> Identifies the original form of the material that was digitized, such as “slide,” “paper,” “map,” and so forth. This is not necessarily the same as IMED.
ITCH	<i>Technician.</i> Identifies the technician who digitized the subject file. For example, “Smith, John.”

---



# CSET (Character Set) Chunk

To define character-set and language information for a RIFF file, use the CSET chunk. The CSET chunk defines the code page and country, language, and dialect codes for the file. These values can be overridden for specific file elements; see “Usage Codes for Extra Header and Extra Entry Fields,” later in this chapter, for information on specifying character set information in a compound file.

The CSET chunk is defined as follows:

```
<CSET chunk> → CSET( <wCodePage:WORD>  
                      <wCountryCode:WORD>  
                      <wLanguageCode:WORD>  
                      <wDialect:WORD> )
```

The fields are as follows:

---

Field	Description
<b>wCodePage</b>	Specifies the code page used for file elements. If the CSET chunk is not present, or if this field has value zero, assume standard ISO 8859/1 code page (identical to code page 1004 without code points defined in hex columns 0, 1, 8, and 9).
<b>wCountryCode</b>	Specifies the country code used for file elements. See “Country Codes,” following this section, for a list of currently defined country codes.  If the CSET chunk is not present, or if this field has value zero, assume USA (country code 001).
<b>wLanguage, wDialect</b>	Specify the language and dialect used for file elements. See “Language and Dialect Codes,” later in this chapter, for a list of language and dialect codes.  If the CSET chunk is not present, or if these fields have value zero, assume US English (language code 9, dialect code 1).

---

## Country Codes

Use one of the following country codes in the **wCountryCode** field:

---

Country Code	Country
000	None (ignore this field)
001	USA
002	Canada
003	Latin America
030	Greece
031	Netherlands

032	Belgium
033	France
034	Spain
039	Italy
041	Switzerland
043	Austria
044	United Kingdom
045	Denmark
046	Sweden
047	Norway
049	West Germany
052	Mexico
055	Brazil
061	Australia
064	New Zealand
081	Japan
082	Korea
086	People's Republic of China
088	Taiwan
090	Turkey
351	Portugal
352	Luxembourg
354	Iceland
358	Finland

---

## Language and Dialect Codes

Specify one of the following pairs of language-code and dialect-code values in the **wLanguage** and **wDialect** fields:

Language Code	Dialect Code	Language
0	0	None (ignore these fields)
1	1	Arabic
2	1	Bulgarian
3	1	Catalan
4	1	Traditional Chinese
4	2	Simplified Chinese
5	1	Czech
6	1	Danish
7	1	German
7	2	Swiss German
8	1	Greek
9	1	US English
9	2	UK English
10	1	Spanish
10	2	Spanish Mexican
11	1	Finnish
12	1	French
12	2	Belgian French
12	3	Canadian French
12	4	Swiss French
13	1	Hebrew
14	1	Hungarian

15	1	Icelandic
16	1	Italian
16	2	Swiss Italian
17	1	Japanese
18	1	Korean
19	1	Dutch
19	2	Belgian Dutch
20	1	Norwegian - Bokmal
20	2	Norwegian - Nynorsk
21	1	Polish
22	1	Brazilian Portuguese
22	2	Portuguese
23	1	Rhaeto-Romanic
24	1	Romanian
25	1	Russian
26	1	Serbo-Croatian (Latin)
26	2	Serbo-Croatian (Cyrillic)
27	1	Slovak
28	1	Albanian
29	1	Swedish
30	1	Thai
31	1	Turkish
32	1	Urdu
33	1	Bahasa

---

## JUNK (Filler) Chunk

A JUNK chunk represents padding, filler or outdated information. It contains no relevant data; it is a space filler of arbitrary size. The JUNK chunk is defined as follows:

`<JUNK chunk> → JUNK( <filler> )`

where **<filler>** contains random data.

## Compound File Structure

The compound file structure is a RIFF-based structure upon which multimedia file formats can be defined. The compound file structure is a parameterized structure that provides for the following:

- Storage of multimedia data elements
- Direct access to multimedia data elements (as opposed to sequential searching)

The goals of the compound file structure are to maximize flexibility and extensibility while minimizing implementation costs. Using the compound file structure, developers of multimedia data formats can define both simple and complex file formats.

The structure is flexible enough to be used for many purposes, but it can be simplified for use with simple file formats. Designers of new multimedia file formats can restrict the use of standard header fields, requiring some and removing others.

For example, a developer might define a compound file format that stores a series of bitmaps in a single file, thus reducing compact disc seek times. Another developer might define a compound file format that contains a special type of audio resource, using the compound file header information to identify the attributes of the audio data stored within.

## Structural Overview

Files based upon the compound file structure contain the following two RIFF chunks at their topmost level:

- Compound File Table of Contents (CTOC) chunk
- Compound File Element Group (CGRP) chunk

The CTOC chunk indexes the CGRP chunk, which contains the actual multimedia data elements. Defined using the standard chunk notation, a compound file is represented as follows:

```
<compound file> → RIFF('type' <CTOC> <CGRP>)
```

where 'type' is a FOURCC indicating the file type.

This section describes the CTOC and CGRP chunks in detail.

## Compound File Table of Contents (CTOC) Chunk

The CTOC chunk functions mainly as an index, allowing direct access to elements within a compound file. The CTOC chunk also contains information about the attributes of the entire file and of each media element within the file.

To provide the maximum flexibility for defining compound file formats, the CTOC chunk can be customized at several levels. The CTOC chunk contains fields whose length and usage is defined by other CTOC fields. This parameterization adds complexity, but it provides flexibility to file format designers and allows applications to correctly read data without necessarily knowing the specific file format definition.

## Structural Overview

The CTOC chunk defines the contents of the CGRP chunk. The CTOC chunk has the following components:

- Header information defining the size of the CTOC chunk, the number of entries in the CGRP chunk, the size of the CGRP chunk, and general information about the entire header file
- A parameter table definition defining the size and contents of the header parameter table and CTOC table entries
- A header parameter table defining attributes that apply to the entire compound file.
- CTOC table entries defining the location, size, name, and attributes of the compound file elements contained in the CGRP chunk.

These components appear sequentially in the CTOC chunk. The individual fields in the CTOC chunk are the following:

```
<CTOC-chunk> → CTOC (
    <dwHeaderSize:DWORD>           // Header information
    <dwEntriesTotal:DWORD>
    <dwEntriesDeleted:DWORD>
    <dwEntriesUnused:DWORD>
    <dwBytesTotal:DWORD>
    <dwBytesDeleted:DWORD>
    <dwHeaderFlags:DWORD>

    <wEntrySize:WORD>             // Parameter table definition
    <wNameSize:WORD>
    <wExHdrFields:WORD>
    <wExEntFields:WORD>
    <awExHdrFldUsage:WORD[wExHdrFields]>
    <awExEntFldUsage:WORD[wExEntFields]>

    // Header parameter table
    <adwExHdrField:DWORD[wExHdrFields]>
    [<bHeaderPad:BYTE>]
    [<CTOC-table-entry>]         // CTOC table entries
)
```

Each CTOC table entry is defined as follows:

```
<CTOC-table-entry> →
    <dwOffset:DWORD>
    <dwSize:DWORD>
    <dwMedType:DWORD>
    <dwMedUsage:DWORD>
    <dwCompressTech:DWORD>
    <dwUncompressBytes:DWORD>
    <adwExEntField:DWORD[wExEntFields]>
    <bEntryFlags:BYTE>
    <achName:CHAR[wNameSize]>
    [<bEntryPad:BYTE>]...
```

The following sections describe each field in detail.

## Header Information

The header information section defines general information about the CTOC header and about the entire compound file. It contains the following fields:

---

Field Name	Description
<b>dwHeaderSize</b>	Combined size of header information, parameter table definition, and header parameter table. Use this value to locate the start of the CTOC table entries within the CTOC chunk.
<b>dwEntriesTotal</b>	Total number of CTOC table entries, including unused entries and entries corresponding to deleted elements.
<b>dwEntriesDeleted</b>	Number of CTOC table entries that correspond to deleted elements.
<b>dwEntriesUnused</b>	Number of CTOC table entries that are unused.
<b>dwBytesTotal</b>	Combined size of all CGRP elements, including deleted elements.
<b>dwBytesDeleted</b>	Combined size of all deleted CGRP elements.
<b>dwHeaderFlags</b>	Flags that give information about the entire compound file. The following flags may be used:  CTOC_HF_SEQUENTIAL Valid CTOC table entries are arranged in sequential order. If this flag is not set, the CTOC table entries may be in an arbitrary order.  CTOC_HF_MEDSUBTYPE The <b>dwMedUsage</b> field of each CTOC table entry contains a FOURCC that indicates how the element is used. If this flag is not set, the <b>dwMedUsage</b> field contains information as defined by the form type.

---

## Parameter Table Definition

The parameter table definition defines the size and contents of the header parameter table and CTOC table. It contains the following fields:

---

Field Name	Description
<b>wEntrySize</b>	Size of each CTOC table entry, including any pad bytes.
<b>wNameSize</b>	Size of the <b>achName</b> field of each CTOC table entry. Each <b>achName</b> field must be padded with null characters to this length. The <b>achName</b> field is a null-terminated string, so it always contains at least one trailing null character.
<b>wExHdrFields</b>	Number of extra header fields, or entries in the <b>awExHdrFldUsage</b> and <b>adwExHdrField</b> arrays.
<b>wExEntFields</b>	Number of extra entry fields, or entries in the <b>awExEntFldUsage</b> and <b>adwExHdrField</b> arrays.

<b>awExHdrFldUsage</b>	Array of extra header field usage fields. Each entry in this array corresponds to the same numbered entry in the <b>adwExHdrField</b> array and defines how that entry is interpreted. Valid usage codes for each field in this array are listed in “Usage Codes for Extra Header and Extra Entry Fields,” later in this chapter. The number of WORDs in this array is defined by the <b>wExHdrFields</b> value.
<b>awExEntFldUsage</b>	Array of extra entry field usage fields. Each entry in this array corresponds to the same numbered entry in the <b>adwExEntField</b> array, present in each CTOC table entry, and defines how that entry is interpreted. Valid usage codes for each field in this array are listed in “Usage Codes for Extra Header and Extra Entry Fields,” later in this chapter. The number of WORDs in this array is defined by the <b>wExEntFields</b> value.

---

## Header Parameter Table

The header parameter table is an optional component generally used to define attributes of the entire compound file.

---

<b>Field Name</b>	<b>Type</b>
<b>adwExHdrField</b>	Extra header fields. The usage of each cell in the array is defined by the corresponding cell in the <b>awExHdrFldUsage</b> array.  The number of DWORDs in this array is defined by the <b>wExHdrFields</b> value.
<b>bHeaderPad</b>	Zero or more NULL pad bytes. There must be enough padding in this field to make the CTOC header an even number of bytes in length.

---

## CTOC Table Entries

The CTOC table entries define the location, size, name, and other information about the individual compound file elements contained in the CGRP chunk. The number of CTOC table entries is determined by the **dwEntriesTotal** field in the header information of the CTOC chunk.

Each CTOC table entry is a structure containing the following fields:

<b>Field Name</b>	<b>Description</b>
<b>dwOffset</b>	<p>Byte offset of the compound file element measured from the beginning of the data portion of the CGRP chunk.</p> <p>For example, if <b>dwOffset</b> is 1000 and the chunk ID of the CGRP chunk is at offset 500, the element is at offset 1508 (1000+500+4 (chunk ID)+4 (chunk size field)).</p>
<b>dwSize</b>	Size of the element in bytes.
<b>dwMedType</b>	FOURCC value identifying the media element type of the compound file element. This field may be zero if the compound file element is not to be interpreted as a standalone file. If the compound file element is a RIFF form, then the media element type is the same as the RIFF form type.
<b>dwMedUsage</b>	<p>Extra usage information for the compound file element.</p> <p>If the CTOC_HF_MEDSUBTYPE flag is set in the <b>dwHeaderFlags</b> field, this field contains a FOURCC that indicates how the element is used. To avoid name conflicts, this FOURCC must be registered. See “Registering Multimedia Formats” in Chapter 1, “Overview of Multimedia Specifications,” for information on usage codes.</p> <p>If the CTOC_HF_MEDSUBTYPE flag is not set in the <b>dwHeaderFlags</b> field, this field contains 32 bits of information interpreted as defined by the form type.</p>
<b>dwCompressTech</b>	Compression technique used to compress the media element. If this value is zero, the element is not compressed. See “Compression of Compound File Elements,” later in this chapter, for more information.
<b>dwUncompressBytes</b>	Number of bytes the compound file element occupies in memory after decompression. This value assumes the decompression technique identified in the <b>dwCompressTech</b> field. If the <b>dwCompressTech</b> field is 0, then the compound file element is not compressed, and this field should equal the <b>dwSize</b> field.
<b>adwExEntField</b>	<p>Array of extra entry fields defining attributes of this compound file element. The usage of each cell in the array is defined by the corresponding cell in the <b>awExEntFldUsage</b> array.</p> <p>The number of DWORDs in this array is defined by the <b>wExEntFields</b> value.</p>
<b>bEntryFlags</b>	<p>Flags giving information about the compound file element or this CTOC table entry. Possible values follow; these may be combined:</p> <p>CTOC_EF_DELETED Compound file is marked as deleted and should not be accessed. Do not combine this flag with the CTOC_EF_UNUSED flag.</p> <p>CTOC_EF_UNUSED CTOC table entry is unused and does not refer to any compound file element. This entry can be used to refer to a new compound file element. Do not combine this flag with the CTOC_EF_DELETED flag.</p>
<b>achName</b>	Array of characters containing the name of the compound file element. The number of bytes in this array is defined by the



**wNameSize** value.

The string must be padded on the right with NULL characters and must be terminated by at least one NULL character. This field must be an odd number of bytes in length and must be at least one byte long.

**bEntryPad**

Zero or more NULL pad bytes as needed to make the table entry an even number of bytes in length.

---

## Usage Codes for Extra Header and Extra Entry Fields

The following are valid usage codes for elements in the **awExHdrFldUsage** and **awExEntFldUsage** arrays, both of which are fields of the CTOC header. These arrays define the meaning of data stored in the **adwExHdrField** and **adwExEntField** “extra fields.” All usage codes apply to both header fields and entry fields, unless explicitly stated otherwise.

Values marked in the extra header field arrays generally apply to all elements in the CFRG chunk, while values marked in the extra entry field arrays generally apply only to the element referenced by the corresponding CTOC table entry.

---

<b>Flag</b>	<b>Description</b>
CTOC_EFU_UNUSED (0x00)	The field is unused. This usage code may be used to logically delete a header field.
CTOC_EFU_LASTMODTIME (0x01)	When used to describe an extra header field, the field contains the time that any portion of the CTOC or CGRP was last modified.  When used to describe an extra entry field, the field contains the time that the corresponding CTOC table entry, or the compound file element it refers to, was last modified.  The field is interpreted as a DWORD containing the number of seconds that have elapsed since 00:00:00 Greenwich Mean Time (GMT), January 1, 1970.
CTOC_EFU_CODEPAGE	The field contains the code page and country code for the <b>achName</b> field. These values override any values specified in a CSET chunk.  When used to describe an extra header field, the field contains code-page and country-code information for all CTOC table entries. When used to describe an extra entry field, the field contains information for that specific CTOC table entry.  The low-order word of the field contains one of the following code page values:  Zero Use standard ISO 8859/1 code page. This is identical to code page 1004 without code points defined in hex columns 0, 1, 8, and 9.  CTOC_CHARSET_CODEPAGE (0x0000nnnn) Use code page 0xnnnn, where 0xnnnn is the 16-bit code page number. For example, 0x00000352 for OS/2

code page 850, or 0x000004E4 for Windows 3.1 code page 1252.

The high-order word contains one of the following country codes:

Zero

Ignore this field.

Country code

See “Country Codes,” earlier in this chapter, for a list of currently defined country codes.

CTOC\_EFU\_LANGUAGE

The field contains language and dialect information for the **achName** field. These values override any values specified in a CSET chunk.

When used to describe an extra header field, the field contains language information for all CTOC table entries. When used to describe an extra entry field, the field contains information for that specific CTOC table entry.

The low-order word of the field contains one of the following language codes:

Zero

Ignore this field.

Language code

See “Language and Dialect Codes,” earlier in this chapter, for a list of currently defined language codes.

The high-order word of the field contains one of the following dialect codes:

Zero

Ignore this field.

Dialect code

See “Language and Dialect Codes,” earlier in this chapter, for a list of currently defined dialect codes.

CTOC\_EFU\_COMPRESSPARAM0  
(0x05) through  
CTOC\_EFU\_COMPRESSPARAM9  
(0x14)

Specifies a compression parameter. See “Compression of Compound File Elements,” later in this chapter.

---

## Compression of Compound File Elements

Compound file elements can be compressed. The **dwCompressTech** field of a CTOC table entry contains a FOURCC compression technique identifier for the corresponding compound file element. If the field is zero, the compound file element is not compressed.

The definition of a specific compression technique may specify that either the entire compound file element is compressed, or that some specific subset, for example one or more RIFF chunks, is compressed.

The **dwUncompressSize** field contains the number of bytes that the compound file element will occupy in memory after decompression. If the compound file element is not compressed, this field contain the same value as the **dwSize** field, which identifies the file size of the compound file element.

Compression techniques may demand extra header fields or extra entry fields for decompression parameters. Compression technique identifiers, and any new entry fields corresponding to decompression technique parameters, must be unique. See “Registering Multimedia Formats” in Chapter 1, “Overview of Multimedia Specifications,” for information on registering compression techniques.

## **Compound File Element Group (CGRP) Chunk**

The actual elements of data referenced by the CTOC chunk are stored in a compound file Element Group (CGRP) chunk. The CGRP chunk contains all the compound file elements, concatenated together into one contiguous block of data. Some of the elements in the CGRP chunk might be unused, if the element was marked for deletion or was altered and stored elsewhere within the CGRP chunk.

Elements within the CGRP chunk are of arbitrary size and can appear in a specific or arbitrary order, depending upon the file format definition. Each element is identified by a corresponding CTOC table entry.

Using the standard RIFF notation, the CGRP chunk is defined as follows:

```
<CGRP-chunk> → CGRP([<compound file element>]...)
```

## **Placement of the CTOC and CGRP Chunks**

The specific file format definition can specify which of the two chunks appear first the data file. Generally, the CTOC chunk is placed at the front of the file to reduce the seek and read times required to access it. During authoring time, an application might place the CTOC chunk at the end of the file, so it can be expanded as elements are added to the CGRP chunk.

## Chapter 3

# Multimedia File Formats

---

This chapter describes the multimedia file formats. Most of these file formats are based on the Resource Interchange File Format (RIFF), described in Chapter 2.

This chapter describes the following file formats:

- Bundle File Format (BND)
- Device Independent Bitmap File Format (DIB)
- RIFF DIB File Format (RDIB)
- Musical Instrument Digital Interface File Format (MIDI)
- RIFF MIDI File Format (RMID)
- Palette File Format (PAL)
- Rich Text Format (RTF)
- Waveform Audio File Format (WAVE)

## Bundle File Format

The Bundle (BND) format contains a series of RIFF chunks or other multimedia files. The BND file is defined as follows:

```
<BND-file> → RIFF( 'BND' <CTOC-chunk> <CGRP-chunk> )
```

The **<CTOC-chunk>** and **<CGRP-chunk>** formats are defined in “Compound File Structure,” in Chapter 2, “Resource Interchange File Format.”

Each compound file element must be capable of standing alone as an independent file. An element may not be a random chunk (except the RIFF chunk, indicating a RIFF file) or random binary data (unless the binary data is supposed to be treated as a file).

## Device Independent Bitmap File Format

The Device Independent Bitmap (DIB) format represents bitmap images in a device-independent manner. Bitmaps can be represented at 1, 4, and 8 bits per pixel, with a palette containing colors

represented in 24 bits. Bitmaps can also be represented at 24 bits per pixel without a palette and in a run-length encoded format.

This documentation describes three types of DIB files:

- Windows version 3.0 device-independent bitmap files
- OS/2 Presentation Manager version 1.2 device-independent bitmap files
- RIFF device-independent bitmap files

The Windows 3.0 and Presentation Manager 1.2 DIBs are similar, so they are discussed together.

## Overview of DIB Structure

Windows 3.0 and Presentation Manager 1.2 DIB files consist of the following sequence of data structures:

- A file header
- A bitmap information header
- A color table
- An array of bytes that defines the bitmap bits

The following sections describe each of these structures.

## Bitmap File Header

The bitmap file header contains information about the type, size, and layout of a device-independent bitmap (DIB) file. In both the Windows 3.0 and Presentation Manager 1.2 DIBs, it is defined as a BITMAPFILEHEADER data structure:

```
typedef struct tagBITMAPFILEHEADER {  
    WORD    bfType;  
    DWORD   bfSize;  
    WORD    bfReserved1;  
    WORD    bfReserved2;  
    DWORD   bfOffBits;  
} BITMAPFILEHEADER;
```

The following table describes the fields.

---

<b>Field</b>	<b>Description</b>
<b>bfType</b>	Specifies the file type. It must consist of the character sequence BM (WORD value 0x4D42).
<b>bfSize</b>	Specifies the file size in bytes.
<b>bfReserved1</b>	Reserved. Must be set to zero.
<b>bfReserved2</b>	Reserved. Must be set to zero.

**bfOffBits** Specifies the byte offset from the BITMAPFILEHEADER structure to the actual bitmap data in the file.

---

## Bitmap Information Header

The BITMAPINFO and BITMAPCOREINFO data structures define the dimensions and color information for Windows 3.0 and Presentation Manager 1.2 DIBs, respectively. They are defined as follows:

---

<b>Windows 3.0 DIB</b>	<b>Presentation Manager 1.2 DIB</b>
<pre>typedef struct tagBITMAPINFO {     BITMAPINFOHEADER bmiHeader;     RGBQUAD bmiColors[1]; } BITMAPINFO;</pre>	<pre>typedef struct _BITMAPCOREINFO {     BITMAPCOREHEADER bmciHeader;     RGBTRIPLE bmciColors[1]; } BITMAPCOREINFO;</pre>

---

These structures are essentially alike, and this section discusses them simultaneously. Each field name for the Windows BITMAPINFO structure is followed by the corresponding field name for the Presentation Manager BITMAPCOREINFO 1.2 structure, in parentheses.

The following table describes the fields.

---

<b>Windows (PM) Field</b>	<b>Description</b>
<b>bmiHeader</b> <b>(bmciHeader)</b>	Specifies information about the dimensions and color format of the DIB. The BITMAPINFOHEADER and BITMAPCOREHEADER data structures are described in the next section.
<b>bmiColors</b> <b>(bmciColors)</b>	Specifies the DIB color table. The RGBQUAD and RGBTRIPLE data structures are described in “Bitmap Color Table,” later in this chapter.

---

## Information Header Structures

The BITMAPINFOHEADER and BITMAPCOREHEADER structures contain information about the dimensions and color format of Windows 3.0 and Presentation Manager 1.2 DIBs, respectively. They are defined as follows:

---

### Windows 3.0 DIB

```
typedef struct tagBITMAPINFOHEADER {
    DWORD biSize;
    DWORD biWidth;
    DWORD biHeight;
    WORD biPlanes;
    WORD biBitCount;
    DWORD biCompression;
    DWORD biSizeImage;
    DWORD biXPelsPerMeter;
    DWORD biYPelsPerMeter;
    DWORD biClrUsed;
    DWORD biClrImportant;
} BITMAPINFOHEADER;
```

### Presentation Manager 1.2 DIB

```
typedef struct tagBITMAPCOREHEADER {
    DWORD bcSize;
    WORD bcWidth;
    WORD bcHeight;
    WORD bcPlanes;
    WORD bcBitCount;
} BITMAPCOREHEADER;
```

---

Because these structures are essentially alike, except for the added fields in the Windows 3.0 structure, this section discusses them simultaneously. Each field name for the Windows structure is followed by the corresponding field name for the Presentation Manager structure, in parentheses.

### Common Fields

The following fields are present in both the Windows 3.0 and Presentation Manager 1.2 formats:

---

Windows (PM) Field	Description
<b>biSize (bcSize)</b>	Specifies the number of bytes required by the BITMAPINFOHEADER structure. You can use this field to distinguish between Windows 3.0 and Presentation Manager 1.2 DIBs.
<b>biWidth (bcWidth)</b>	Specifies the width of the DIB in pixels.
<b>biHeight (bcHeight)</b>	Specifies the height of the DIB in pixels.
<b>biPlanes (bcPlanes)</b>	Specifies the number of planes for the target device. Must be set to 1.
<b>wBitCount (bcBitCount)</b>	Specifies the number of bits-per-pixel. See “Interpreting the Color Table,” later in this section, for more information.

---

## Windows Fields

The following fields are present only in the Windows 3.0 BITMAPINFOHEADER structure:

---

Field	Description								
<b>biCompression</b>	<p>Specifies the type of compression for a compressed bitmap. It can be one of the following values:</p> <hr/> <table><thead><tr><th>Value</th><th>Meaning</th></tr></thead><tbody><tr><td>BI_RGB</td><td>Specifies that the bitmap is not compressed.</td></tr><tr><td>BI_RLE4</td><td>Specifies a run-length encoded format for bitmaps with 4 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by two word-length color indexes.</td></tr><tr><td>BI_RLE8</td><td>Specifies a run-length encoded format for bitmaps with 8 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by a color-index byte.</td></tr></tbody></table> <hr/> <p>See “Windows 3.0 Bitmap Compression Formats” later in this document for information about the encoding schemes.</p>	Value	Meaning	BI_RGB	Specifies that the bitmap is not compressed.	BI_RLE4	Specifies a run-length encoded format for bitmaps with 4 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by two word-length color indexes.	BI_RLE8	Specifies a run-length encoded format for bitmaps with 8 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by a color-index byte.
Value	Meaning								
BI_RGB	Specifies that the bitmap is not compressed.								
BI_RLE4	Specifies a run-length encoded format for bitmaps with 4 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by two word-length color indexes.								
BI_RLE8	Specifies a run-length encoded format for bitmaps with 8 bits-per-pixel. The compression format is a two-byte format consisting of a count byte followed by a color-index byte.								
<b>biSizeImage</b>	Specifies the size in bytes of the image.								
<b>biXPelsPerMeter</b>	Specifies the horizontal resolution in pixels per meter of the target device for the bitmap. An application can use this value to select a bitmap from a resource group that best matches the characteristics of the current device.								
<b>biYPelsPerMeter</b>	Specifies the vertical resolution in pixels per meter of the target device for the bitmap.								
<b>biClrUsed</b>	<p>Specifies the number of color values in the color table actually used by the bitmap. Possible values follow.</p> <hr/> <table><thead><tr><th>Value</th><th>Result</th></tr></thead><tbody><tr><td>0</td><td>Bitmap uses the maximum number of colors corresponding to the value of the <b>wBitCount</b> field.</td></tr><tr><td>Nonzero</td><td>If the <b>wBitCount</b> value is less than 24, the <b>biClrUsed</b> value indicates the actual number of colors which the graphics engine or device driver will access.  If the <b>wBitCount</b> value is 24, the <b>biClrUsed</b> value indicates the size of the reference color table used to optimize performance of Windows color palettes.</td></tr></tbody></table> <hr/> <p>If the bitmap is a “packed” bitmap (that is, a bitmap in which the bitmap array immediately follows the BITMAPINFO header and which is referenced by a single pointer), the <b>biClrUsed</b> field must be set to 0 or to the actual size of the color table. See “Interpreting the Color Table,” later in this section, for more information on how this field affects the interpretation of the color table.</p>	Value	Result	0	Bitmap uses the maximum number of colors corresponding to the value of the <b>wBitCount</b> field.	Nonzero	If the <b>wBitCount</b> value is less than 24, the <b>biClrUsed</b> value indicates the actual number of colors which the graphics engine or device driver will access.  If the <b>wBitCount</b> value is 24, the <b>biClrUsed</b> value indicates the size of the reference color table used to optimize performance of Windows color palettes.		
Value	Result								
0	Bitmap uses the maximum number of colors corresponding to the value of the <b>wBitCount</b> field.								
Nonzero	If the <b>wBitCount</b> value is less than 24, the <b>biClrUsed</b> value indicates the actual number of colors which the graphics engine or device driver will access.  If the <b>wBitCount</b> value is 24, the <b>biClrUsed</b> value indicates the size of the reference color table used to optimize performance of Windows color palettes.								
<b>biClrImportant</b>	Specifies the number of color indexes that are considered important for displaying the bitmap. If this value is 0, then all colors are								

---



important.

---

## Bitmap Color Table

The color table is a collection of 24-bit RGB values. There are as many entries in the color table as there are colors in the bitmap. The color table isn't present for bitmaps with 24 color bits because each pixel is represented by 24-bit RGB values in the actual bitmap data area.

### Color Table Structure

The color table for Windows 3.0 and Presentation Manager 1.2 DIBs consists of an array of RGBQUAD and RGBTRIPLE structures, respectively. These structures are defined as follows:

---

#### Windows 3.0 DIB

```
typedef struct tagRGBQUAD {
    BYTE rgbBlue;
    BYTE rgbGreen;
    BYTE rgbRed;
    BYTE rgbReserved;
} RGBQUAD;
```

#### Presentation Manager 1.2 DIB

```
typedef struct tagRGBTRIPLE {
    BYTE rgbtBlue;
    BYTE rgbtGreen;
    BYTE rgbtRed;
} RGBTRIPLE;
```

---

Because these structures are essentially alike, this section discusses them simultaneously. Each field name for the Windows RGBQUAD structure is followed by the corresponding field name for the Presentation Manager RGBTRIPLE structure, in parentheses.

### Order of Colors

The colors in the table should appear in order of importance. This can help a device driver render a bitmap on a device that cannot display as many colors as there are in the bitmap. If the DIB is in Windows 3.0 format, the driver can use the **biClrImportant** field of the BITMAPINFOHEADER structure to determine which colors are important.

### Field Descriptions

The RGBQUAD (RGBTRIPLE) structure contains the following fields:

---

#### Windows (PM) Field

#### Description

<b>rgbBlue (rgbtBlue)</b>	Specifies the blue intensity.
<b>rgbGreen (rgbtGreen)</b>	Specifies the green intensity.
<b>rgbRed (rgbtRed)</b>	Specifies the red intensity.
<b>rgbReserved</b> (no PM equivalent)	Not used. Must be set to 0.

---

## Locating the Color Table

An application can use the **biSize** (**bcSize**) field of the BITMAPINFOHEADER (BITMAPCOREHEADER) structure to locate the color table. Each of the following statements assigns the pColor variable the byte offset of the color table from the beginning of the file:

```
// Windows 3.0 DIB
pColor = (LPSTR)pBitmapInfo + (WORD)pBitmapInfo->biSize

// Presentation Manager 1.2 DIB
pColor = (LPSTR)pBitmapCoreInfo + (WORD)pBitmapCoreInfo->bcSize
```

## Interpreting the Color Table

The **biSize** (**bcSize**) field of the BITMAPINFOHEADER (BITMAPCOREHEADER) structure specifies how many bits define each pixel and specifies the maximum number of colors in the bitmap. Its value affects your interpretation of the color table.

The **biSize** (**bcSize**) field can have any of the following values:

---

Value	Meaning
1	The bitmap is monochrome, and the color table contains two entries. Each bit in the bitmap array represents a pixel. If the bit is clear, the pixel is displayed with the color of the first entry in the color table. If the bit is set, the pixel has the color of the second entry in the table.
4	The bitmap has a maximum of 16 colors. Each pixel in the bitmap is represented by a four-bit index into the color table.  For example, if the first byte in the bitmap is 0x1F, then the byte represents two pixels. The first pixel contains the color in the second table entry, and the second pixel contains the color in the 16th table entry.
8	The bitmap has a maximum of 256 colors. Each pixel in the bitmap is represented by a byte-sized index into the color table. For example, if the first byte in the bitmap is 0x1F, then the first pixel has the color of the thirty-second table entry.
24	The bitmap has a maximum of $2^{24}$ colors. The <b>bmiColors</b> ( <b>bmciColors</b> ) field is NULL, and each three bytes in the bitmap array represent the relative intensities of red, green, and blue, respectively, of a pixel.

---

### Note on Windows DIBs

For Windows 3.0 DIBs, the field of the BITMAPINFOHEADER structure specifies the number of color indexes in the color table actually used by the bitmap. If the **biClrUsed** field is set to 0, the bitmap uses the maximum number of colors corresponding to the value of the field.

## Bitmap Data

The bits in the array are packed together, but each line of pixels, or scan line, must be zero-padded to end on a LONG boundary. When the bitmap is in memory, segment boundaries can appear anywhere in the bitmap. The origin of the bitmap is the lower-left corner. The following section discusses compression formats for the Windows 3.0 bitmap data.

## Windows 3.0 Bitmap Compression Formats

Windows supports run-length encoded formats for compressing 4- and 8-bit bitmaps. Compression reduces the disk and memory storage required for the bitmap. The following sections describe the compression formats.

### Compression of 8-Bit-Per-Pixel DIBs

When the **biCompression** field is set to BI\_RLE8, the bitmap is compressed using a run-length encoding format for an 8-bit bitmap. This format uses two modes:

RDIB' form is defined as follows, using the standard RIFF form definition notation:

```
<RDIB-form> → RIFF ( 'RDIB' data( <DIB-data> ) )
```

The **<DIB-data>** format is defined in “Device Independent Bitmap File Format,” earlier in this chapter.

## Extended RDIB Format

The extended RDIB format, designed to incorporate enhancements such as compression, is defined as follows:

```
<RDIB-form> → RIFF( 'RDIB'
    <bmhd-ck> // Bitmap header chunk
    [ <pal-file> | // Internal palette chunk
      <XPAL-ck> ] // External palette chunk
    <bitmap-data> ) // Bitmap data
```

The **<pal-file>** chunk can be any of the palette-file formats discussed in “Palette File Format,” later in this chapter. The **<bmhd-ck>**, **<XPAL-chunk>**, and **<bitmap-data>** are described in the following sections.

## Bitmap Header Chunk

The **<bmhd-ck>** bitmap header chunk is defined as follows:

```
<bmhd-chunk> → bmhd( struct {
    DWORD dwMemSize; // If dwPelFormat is 'data', only these
    DWORD dwPelFormat; // four fields are present.
    WORD wTransType;
```

```

        DWORD   dwTransVal;
        DWORD   dwHdrSize;           // Fields from dwHdrSize forward match
        DWORD   dwWidth;            // the Windows BITMAPINFOHEADER
        DWORD   dwHeight;           // structure, though some fields can
        WORD    dwPlanes;           // contain new values.
        WORD    dwBitCount;
        DWORD   dwCompression;
        DWORD   dwSizeImage;
        DWORD   dwXPelsPerMeter;
        DWORD   dwYPelsPerMeter;
        DWORD   dwClrUsed;
        DWORD   dwClrImportant;
    } )

```

If the **dwCompression** field equals BI\_RGB or BI\_RLE8 or BI\_RLE4, then the extended RDIB has the same bitmap format as a simple RDIB.

Each pixel format defines the orientation, or position of the bitmap origin. Windows bitmaps (identified by a value of 'data' in the **dwPixelFormat** field) have the origin at the bottom left. By default, the other formats have the origin at the top left.

Field	Description																																																																				
<b>dwMemSize</b>	<p>Equal to the size of the bitmap bits if the bits are uncompressed. For RDIBs with <b>dwPixelFormat</b> equal to 'data,' <b>dwMemSize</b> has one of the following values:</p> <table border="1"> <thead> <tr> <th>Image Type</th> <th>Field Value</th> </tr> </thead> <tbody> <tr> <td>Non-RLE</td> <td>Same as <b>dwSizeImage</b> value</td> </tr> <tr> <td>8-bit RLE</td> <td>Size as an uncompressed, 8-bit image</td> </tr> <tr> <td>4-bit RLE</td> <td>Size as an uncompressed, 4-bit image</td> </tr> </tbody> </table>	Image Type	Field Value	Non-RLE	Same as <b>dwSizeImage</b> value	8-bit RLE	Size as an uncompressed, 8-bit image	4-bit RLE	Size as an uncompressed, 4-bit image																																																												
Image Type	Field Value																																																																				
Non-RLE	Same as <b>dwSizeImage</b> value																																																																				
8-bit RLE	Size as an uncompressed, 8-bit image																																																																				
4-bit RLE	Size as an uncompressed, 4-bit image																																																																				
<b>dwPixelFormat</b>	<p>Specifies a FOURCC code defining the pixel format of the bitmap data. The bitmap data is stored in a chunk (or chunks) that has the same chunk ID as is contained in <b>dwPixelFormat</b>. The compression scheme and pixel depth of the bitmap data are recorded in the <b>dwCompression</b> and <b>dwBitCount</b> fields. The current bitmap data values are as follows:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Bitmap Data Location and Format</th> </tr> </thead> <tbody> <tr> <td>'data'</td> <td>Bitmap data is stored in a 'data' chunk using the format defined for Windows 3.0 device independent bitmaps (DIBs). An application can display the bitmap properly even if the fields after (and including) <b>dwMemSize</b> are ignored.</td> </tr> <tr> <td>'palb'</td> <td>Bitmap data is stored in a 'palb' chunk. The pixel format is one of the Windows 3.0 RGB palettized formats (1 to 8 bpp, depending on the value of the <b>dwBitCount</b> field).</td> </tr> <tr> <td>'rgbb'</td> <td>Bitmap data is stored in a 'rgbb' chunk. Pixel format is packed, unpalettized RGB represented at 16, 24, or 32 bits per pixel. The following shows the ordering of the RGB bits for each pixel-depth value. The first extra bit (if present) is the high-order bit.</td> </tr> </tbody> </table> <table border="1"> <thead> <tr> <th><b>dwBitCount</b></th> <th>Extra</th> <th>Red</th> <th>Green</th> <th>Blue</th> </tr> </thead> <tbody> <tr> <td>1</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>2</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>3</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>4</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>5</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>6</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>7</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>8</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>16</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>24</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>32</td> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table>	Value	Bitmap Data Location and Format	'data'	Bitmap data is stored in a 'data' chunk using the format defined for Windows 3.0 device independent bitmaps (DIBs). An application can display the bitmap properly even if the fields after (and including) <b>dwMemSize</b> are ignored.	'palb'	Bitmap data is stored in a 'palb' chunk. The pixel format is one of the Windows 3.0 RGB palettized formats (1 to 8 bpp, depending on the value of the <b>dwBitCount</b> field).	'rgbb'	Bitmap data is stored in a 'rgbb' chunk. Pixel format is packed, unpalettized RGB represented at 16, 24, or 32 bits per pixel. The following shows the ordering of the RGB bits for each pixel-depth value. The first extra bit (if present) is the high-order bit.	<b>dwBitCount</b>	Extra	Red	Green	Blue	1					2					3					4					5					6					7					8					16					24					32				
Value	Bitmap Data Location and Format																																																																				
'data'	Bitmap data is stored in a 'data' chunk using the format defined for Windows 3.0 device independent bitmaps (DIBs). An application can display the bitmap properly even if the fields after (and including) <b>dwMemSize</b> are ignored.																																																																				
'palb'	Bitmap data is stored in a 'palb' chunk. The pixel format is one of the Windows 3.0 RGB palettized formats (1 to 8 bpp, depending on the value of the <b>dwBitCount</b> field).																																																																				
'rgbb'	Bitmap data is stored in a 'rgbb' chunk. Pixel format is packed, unpalettized RGB represented at 16, 24, or 32 bits per pixel. The following shows the ordering of the RGB bits for each pixel-depth value. The first extra bit (if present) is the high-order bit.																																																																				
<b>dwBitCount</b>	Extra	Red	Green	Blue																																																																	
1																																																																					
2																																																																					
3																																																																					
4																																																																					
5																																																																					
6																																																																					
7																																																																					
8																																																																					
16																																																																					
24																																																																					
32																																																																					

15	1	5	5	5
16	0	5	6	5
24	0	8	8	8
32	8	8	8	8

'yuvb' Bitmap data is stored in a 'yuvb' chunk. Pixel format is packed, unpalettized YUV. The exact pixel format is currently undefined. By the time this draft is final, the pixel format will be defined similarly to the 'rgbb' definition.

### wTransType

Specifies the type of transparency representation, if any, used for this image. This is normally used for either image overlay applications, where one image may be visually on top of another, and all pels of the transparency color should not be drawn. Examples include sprites, clip art and motion video overlay. Wherever the transparency color occurs in the picture, the background should be visible.

This information is stored with the image, so that multiple images that use the same color map may all have different transparency color. There are 5 different values for the transparency variable. These are:

Value	Result
BITT_NONE (0x0000)	No pels are considered transparent in this image.
BITT_MAPINDEX (0x0001)	One of the color map/palette entries should be considered the transparency color. All instances of this pel should NOT be drawn, and the existing background should be allowed to show through.
BITT_SINGLECOLOR (0x0002)	A single RGB or YUV value is considered transparent and should not be drawn.
BITT_BITPLANE (0x0003)	An individual bit plane is considered transparent, and all pels that have that bit or bits "on" should not be drawn.
BITT_MULTILEVEL (0x0004)	A set of bits indicate multiple levels of transparency or opacity. This is usually used with 32-bit RGB, where the high 8 bits indicate transparency.

### dwTransVal

These bytes allow the image definition to indicate the exact information about the transparent color. The information is dependent on the value of the **wTransType** as follows:

wTransType	dwTransVal Contents
BITT_NONE	Not used.
BITT_MAPINDEX	Specifies a palette index, either 0 through 16 or 0 through 255, depending on the number of palette entries.
BITT_SINGLECOLOR	Specifies an RGB or YUV value (2 to 4 bytes in size, depending on the pixel format

specified by **dwPelFormat**). All pels that match **dwTransVal** should be considered transparent.

**BITT\_BITPLANE** Specifies a bit mask identifying the bits used to indicate a transparent pel. Any pel that has this set of bits set is totally transparent. This allows multiple colors to be considered transparent. This method works for palettized images; in this case, the value refers to a map entry that is considered transparent.

**BITT\_MULTILEVEL** Specifies bits to use for transparency levels. These bits act as a mask on every pel, and each pel can be matched to the mask to determine the transparency level for the pel.

For example, if **dwTransVal** has value 0xFF000000, then there are 256 levels of transparency. Each pel can be evaluated against the mask. If the pel has a value FFxxxxxx, then it is fully transparent. If the pel has a value 00xxxxxx, then it is fully visible. If the pel has a value 7Fxxxxxx, then the pel is half visible.

**dwHdrSize** Specifies the size of the data portion of the **<bmhdr>** chunk. This is always 40, the size of the BITMAPINFOHEADER structure.

**dwWidth** Specifies the width of the DIB in pixels.

**dwHeight** Specifies the height of the DIB in pixels.

**wPlanes** Specifies the number of planes. This value is normally 1, but it can be 3 or 4 for 24-bit RGB and 32-bit RGB images, respectively. In a multiplane DIB, each color component (for example, red, green, and blue) is stored as a separate plane, and each plan is stored in a separate bitmap data chunk. For example, in a 3-plane, 24-bit 'rgbb' bitmap, the red colors are stored in one 'rgbb' chunk, the green colors in a second 'rgbb' chunk, and the blue colors in a third 'rgbb' chunk.

Allowing the separate RGB planes to be compressed independently can dramatically improve the compression ratio. The **wPlanes** value must be 1 if **dwPelFormat** equals 'data'.

**wBitCount** Specifies the number of bits per pixel. If the **dwPelFormat** field equals 'data', this field must contain values compatible with the Windows 3.0 DIB definition.

**dwCompression** Specifies the type of compression for a compressed bitmap. It can be one of the following values:

Value	Meaning
BI_NONE (0xFFFF0000)	Specifies that the bitmap is not compressed. Pixel values are <i>not</i> padded to four-byte boundaries.
BI_RGB (0x00000000)	Specifies that the bitmap is an uncompressed, 1-, 4-, 8-, or a 24-bit image. For 24-bit images, the palette is optional. Bitmap bits are represented as defined by Windows 3.0 for BI_RGB DIBs. The <b>dwPelFormat</b> field must be set to 'data'.

BI_RLE8 (0x00000001)	Specifies a run-length encoded, compressed bitmap (as defined by Windows 3.0 BI_RLE8 DIBs). The palette is required. The <b>dwPelsFormat</b> field must be set to 'data'.
BI_RLE4 (0x00000002)	Specifies a run-length encoded, compressed bitmap (as defined by Windows 3.0 BI_RLE4 DIBs). The palette is required. The <b>dwPelsFormat</b> field must be set to 'data'.
BI_PACK (0xFFFF0001)	Specifies a simple PACKBITS byte compression scheme consisting of one-byte counts followed by byte data, in the form:  <pre>&lt;count byte n&gt;&lt;data byte 1&gt;&lt;data byte 2&gt;...&lt;data byte n&gt; &lt;count byte n&gt;&lt;data byte to repeat&gt;</pre> <p>The high-order bit of the count byte <i>n</i> is a decision bit:</p>

---

<i>n</i> Value	Data Representation
$n < 0x80$	A run of $n+1$ non-repeating bytes follows.
$n > 0x80$	Data byte is repeated $(n-0x80+1)$ times.
$n = 0x80$	Reserved.

---

BI_TRANS (0xFFFF0002)	Specifies transitional compression, using a table of byte transitions or sequences. See "Transitional Compression," following this table.
BI_CCC (0xFFFF0003)	Specifies CCC compression, a method involving encoding each 4-by-4 block of the image using two colors. See "CCC Compression," following this table.
BI_JPEGN (0xFFFF0004)	To be defined later, when the ISO completes the official specification.

---

<b>dwSizeImage</b>	Specifies the size in bytes of the compressed image.
<b>dwXPelsPerMeter</b>	Specifies the horizontal resolution in pixels per meter of the target device for the bitmap. An application can use this value to select a bitmap from a resource group that best matches the characteristics of the current device. This field is set to zero if unused.
<b>dwYPelsPerMeter</b>	Specifies the vertical resolution in pixels per meter of the target device for the bitmap. This field is set to zero if unused.
<b>dwClrUsed</b>	Specifies the number of palette entries actually used by the bitmap. Possible values follow.

---

Value	Result
0	Bitmap uses the maximum number of colors corresponding to the value of the <b>wBitCount</b> field.
Nonzero	If the <b>wBitCount</b> is less than 24, <b>dwClrUsed</b> specifies the actual number of colors which the

---

graphics engine or device driver will access.

If the **wBitCount** field is set to 24, **dwClrUsed** specifies the size of the reference color table used to optimize performance of Windows color palettes.

---

**dwClrImportant**

Specifies the number of color indexes that are considered important for displaying the bitmap. If this value is 0, then all colors are important.

---

## Transitional Compression

If the **dwCompression** field is set to BI\_TRANS, the data is transitionally compressed using a table of byte transitions or sequences. Values in the data indicate a table position to start at, and the table provides continuing references to other table positions. Transitional compression applies only to eight-bit data, either from an eight-bit palettized image or from a multi-plane image in which each color component is represented in eight bits.

The table consists of up to 256 16-byte rows at the beginning of the data section of the object. Nibbles (half-bytes) in the data section indicate an offset into a table row, at which location is stored the actual byte value. The actual value then becomes the row applicable to the next data nibble. The transitional encoding scheme is described more fully in a separate IBM document.

In transitional compression, the data section is a two-part compound object having the following items:

- A transition table
- The compressed image data

The transition table consists of an integer indicating the table size in bytes and a table of 16-byte rows. The first byte in each row is a row number and the next 15 are transition values. Rows are in descending sequence. The image is compressed according to the following rules:

- Data is in nibbles (half-bytes) or in nibble-pairs (successive half-bytes which may cross a byte boundary).
- The first byte is a nibble-pair. It is the first byte of the image and also the first row number.
- Following a nibble-pair is a series of transition nibbles (1-15) ended by a terminator (0). Each transition nibble indicates an offset in the current row at which the next byte in the image is found; this value is also the next row number.
- The terminator indicates that the next image byte is not in the table, but instead in the following nibble-pair. This value is also the next row number.
- If the picture has an odd number of nibbles (i.e., it ends in the first half of the last byte), an extra zero nibble is included.

## CCC Compression

TBD.



## Palette Chunk

A PLT chunk represents a color table and consists of a valid PAL file. The PAL file format is defined in “Palette File Format,” later in this chapter.

## External Palette Chunk

Instead of a PLT chunk, an RDIB may contain an XPLT chunk, which indicates that the bitmap's palette is stored outside the bitmap. The palette might be stored in a separate file or as a separate compound file element. The XPLT chunk indicates the name and location of the external palette chunk and is defined as follows:

```
<XPLT-chunk> → XPLT(<fccLocation:FOURCC> <szPaletteName::ZSTR>)
```

The **fccLocation** contains one of the following FOURCC values specifying the location of the external chunk:

---

<b>fccLocation Value</b>	<b>Chunk Location</b>
‘full’	Palette is located in an external file, and the <b>szPaletteName</b> value specifies a complete filename with path.
‘file’	Palette is located in an external file, and the <b>szPaletteName</b> value specifies a filename without path.
‘elem’	Palette is located in the same compound file containing the DIB. The <b>szPaletteName</b> value specifies the name of the compound file element.

---

The **szPaletteName** consists of a null-terminated string (ZSTR) containing the name of the external chunk containing the palette.

## Bitmap Data Chunk

The <bitmap-data> contains bitmap data in the format specified by the **biPixelFormat** field of the <bmhd-chunk>.

# MIDI and RIFF MIDI File Formats

The Musical Instrument Digital Interface (MIDI) file format represents a Standard MIDI File, as defined by the MIDI Manufacturers Association. A MIDI file contains commands instructing instruments to play specific notes and perform other operations.

The specifications for MIDI and MIDI files can be obtained from the following organization:

International MIDI Association (IMA)  
5316 W. 57th Street  
Los Angeles, CA 90056  
(213) 649-6434.

The 'RMID' format consists of a standard MIDI file enclosed in a RIFF chunk. Enclosing the MIDI file in a 'RIFF' chunk allows the file to be consistently identified; for example, an 'INFO' list can be included in the file.

The 'RMID' form is defined as follows, using the standard RIFF form definition:

```
<RMID-form> → RIFF ('RMID' data( <MIDI-data> ))
```

The <MIDI-data> is equivalent to a Standard MIDI File.

## Palette File Format

The Palette (PAL) File Format represents a logical palette, which is a collection of colors represented as RGB values. There are two types of PAL formats:

- A simple PAL format
- An extended PAL format

### Simple PAL Format

The simple PAL format is defined as follows:

```
RIFF('PAL' data( <palette:LOGPALETTE> ))
```

LOGPALETTE is the Windows 3.0 logical palette structure, defined as follows:

```
typedef struct tagLOGPALETTE {  
    WORD        palVersion;  
    WORD        palNumEntries;  
    PALETTEENTRY palPalEntry[];  
} LOGPALETTE;
```

The LOGPALETTE structure fields are as follows:

---

Field	Description
-------	-------------

---

---

<b>palVersion</b>	Specifies the Windows version number for the structure.
<b>palNumEntries</b>	Specifies the number of palette color entries.
<b>palPalEntry[]</b>	Specifies an array of PALETTEENTRY data structures that define the color and usage of each entry in the logical palette.

---

The colors in the palette entry table should appear in order of importance. This is because entries earlier in the logical palette are most likely to be placed in the system palette.

The PALETTEENTRY data structure specifies the color and usage of an entry in a logical color palette. The structure is defined as follows:

```
typedef struct tagPALETTEENTRY {
    BYTE    peRed;
    BYTE    peGreen;
    BYTE    peBlue;
    BYTE    peFlags;
} PALETTEENTRY;
```

The PALETTEENTRY structure fields are as follows:

the bitmap's palette is stored outside the bitmap. The palette might be stored in a separate file or as a separate compound file element. The XPLT chunk indicates the name and location of the external palette chunk and is defined as follows:

```
<XPLT-chunk> → XPLT(<fccLocation:FOURCC> <szPaletteName::ZSTR>)
```

The **fccLocation** contains one of the following FOURCC values specifying the location of the external chunk:

---

<b>fccLocation Value</b>	<b>Chunk Location</b>
'full'	Palette is located in an external file, and the <b>szPaletteName</b> value specifies a complete filename with path.
'file'	Palette is located in an external file, and the <b>szPaletteName</b> value specifies a filename without path.
'elem'	Palette is located in the same compound file containing the DIB. The <b>szPaletteName</b> value specifies the name of the compound file element.

---

The **szPaletteName** consists of a null-terminated string (ZSTR) containing the name of the external chunk containing the palette.

## Bitmap Data Chunk

The <bitmap-data> contains bitmap data in the format specified by the **biPelFormat** field of the <bmhd-chunk>.

# MIDI and RIFF MIDI File Formats

The Musical Instrument Digital Interface (MIDI) file format represents a Standard MIDI File, as defined by the MIDI Manufacturers Association. A MIDI file contains commands instructing instruments to play specific notes and perform other operations.

The specifications for MIDI and MIDI files can be obtained from the following organization:

International MIDI Association (IMA)  
5316 W. 57th Street  
Los Angeles, CA 90056  
(213) 649-6434.

The 'RMID' format consists of a standard MIDI file enclosed in a RIFF chunk. Enclosing the MIDI file in a 'RIFF' chunk allows the file to be consistently identified; for example, an 'INFO' list can be included in the file.

The 'RMID' form is defined as follows, using the standard RIFF form definition:

```
<RMID-form> → RIFF ('RMID' data( <MIDI-data> ))
```

The <MIDI-data> is equivalent to a Standard MIDI File.

## Palette File Format

The Palette (PAL) File Format represents a logical palette, which is a collection of colors represented as RGB values. There are two types of PAL formats:

- A simple PAL format
- An extended PAL format

### Simple PAL Format

The simple PAL format is defined as follows:

```
RIFF('PAL' data( <palette:LOGPALETTE> ))
```

LOGPALETTE is the Windows 3.0 logical palette structure, defined as follows:

```
typedef struct tagLOGPALETTE {  
    WORD        palVersion;  
    WORD        palNumEntries;  
    PALETTEENTRY palPalEntry[];  
} LOGPALETTE;
```

The LOGPALETTE structure fields are as follows:

---

Field	Description
-------	-------------

---

---

<b>palVersion</b>	Specifies the Windows version number for the structure.
<b>palNumEntries</b>	Specifies the number of palette color entries.
<b>palPalEntry[]</b>	Specifies an array of PALETTEENTRY data structures that define the color and usage of each entry in the logical palette.

---

The colors in the palette entry table should appear in order of importance. This is because entries earlier in the logical palette are most likely to be placed in the system palette.

The PALETTEENTRY data structure specifies the color and usage of an entry in a logical color palette. The structure is defined as follows:

```
typedef struct tagPALETTEENTRY {
    BYTE    peRed;
    BYTE    peGreen;
    BYTE    peBlue;
    BYTE    peFlags;
} PALETTEENTRY;
```

The PALETTEENTRY structure fields are as follows:

---

<b>Field</b>	<b>Description</b>
<b>peRed</b>	Specifies the intensity of red for the palette entry color.
<b>peGreen</b>	Specifies the intensity of green for the palette entry color.
<b>peBlue</b>	Specifies the intensity of blue for the palette entry color.
<b>peFlags</b>	Specifies how the palette entry is to be used.

---

## Extended PAL Format

The extended PAL format includes the following:

- A palette-header chunk
- A data chunk containing an RGB palette (consisting of a LOGPALETTE structure) or some other palette type, including YUV and XYZ palettes.

For an RGB palette, the extended PAL format is represented as follows:

```
RIFF('PAL' plth( <palette-header> ) data( <LOGPALETTE-data> ))
```

For a YUV palette, the extended PAL format is represented as follows:

```
RIFF('PAL' plth( <palette-header> ) yuvp( <YUV-LOGPALETTE-data> ))
```

Both the <LOGPALETTE-data> and <YUV-LOGPALETTE-data> use the Windows 3.0 LOGPALETTE structure, described in “Simple PAL Format,” earlier in this section. The <YUV-LOGPALETTE-data> contains YUV values instead of RGB values.

The ‘plth’ chunk is defined as follows:

```

<plth-ck> → PLT( struct {
    DWORD    dwMapType;
    WORD     wWhite;           // Fields from this point on are
    WORD     wBlack;          // optional. If they are included
    WORD     wBorder;         // but not used, set them to 0xFFFF.
    WORD     wRegisteredMap;
    WORD     wCustomBase;     // If an application encounters a
    WORD     wCustomCnt;      // 'PLT' chunk smaller than shown
    WORD     wRsvBase;        // here, it should treat the missing
    WORD     wRsvCount;       // fields as unused.
    WORD     wArtBase;
    WORD     wArtCnt;
    WORD     wNumIntense;
} )

```

The structure fields are described in the following:

Field	Description								
<b>dwMapType</b>	<p>FOURCC code specifying the type of palette. Currently, the following palette types are identified:</p> <table border="1"> <thead> <tr> <th>Code</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>'data'</td> <td>Specifies an RGB palette. Data chunk contains a LOGPALETTE structure.</td> </tr> <tr> <td>'yuvp'</td> <td>Specifies a YUV palette. Data chunk contains a YUV palette.</td> </tr> <tr> <td>'xyzp'</td> <td>Specifies an XYZ palette. Data chunk contains a XYZ palette.</td> </tr> </tbody> </table>	Code	Description	'data'	Specifies an RGB palette. Data chunk contains a LOGPALETTE structure.	'yuvp'	Specifies a YUV palette. Data chunk contains a YUV palette.	'xyzp'	Specifies an XYZ palette. Data chunk contains a XYZ palette.
Code	Description								
'data'	Specifies an RGB palette. Data chunk contains a LOGPALETTE structure.								
'yuvp'	Specifies a YUV palette. Data chunk contains a YUV palette.								
'xyzp'	Specifies an XYZ palette. Data chunk contains a XYZ palette.								
<b>wWhite</b> <b>wBlack</b>	Specify palette-map indices corresponding to the closest value of white and black. These identify the pair of colors with the best contrast for use in cursors, calibration, etc. These values are usually changed if the palette changes. Ignore these fields if they contain 0xFFFF.								
<b>wBorder</b>	Specifies the index of the palette entry to be used for any display-border regions, if supported by the display device. Ignore this field if it contains 0xFFFF.								
<b>wRegisteredMap</b>	<p>Specifies how many palette entries correspond to a registered color map. Registered entries are stored at the front of the palette. Ignore this field if it contains 0xFFFF.</p> <p>Registered map entries are always stored at the beginning of the palette, so <b>wRegisteredMap</b> also indicates the index of the first custom color in the palette. Registered color maps include predefined palettes for general use, forest/nature, or seascides. Currently defined values are the following:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>PAL_UNREGISTERED (0xFFFF)</td> <td>Color map does not contain colors from a registered color map.</td> </tr> <tr> <td>PAL_VGA (0x0000)</td> <td>Color map contains the standard 16 VGA colors.</td> </tr> <tr> <td>PAL_AVC198 (0x0001)</td> <td></td> </tr> </tbody> </table>	Value	Description	PAL_UNREGISTERED (0xFFFF)	Color map does not contain colors from a registered color map.	PAL_VGA (0x0000)	Color map contains the standard 16 VGA colors.	PAL_AVC198 (0x0001)	
Value	Description								
PAL_UNREGISTERED (0xFFFF)	Color map does not contain colors from a registered color map.								
PAL_VGA (0x0000)	Color map contains the standard 16 VGA colors.								
PAL_AVC198 (0x0001)									

<b>wCustomBase</b>	Specifies the index of the first custom color of the palette. The beginning of the palette contains the entries of the registered map, so <b>wCustomBase</b> also indicates the number of entries in the registered palette. Map entries starting with <b>wCustomBase</b> comprise additional custom colors used in the bitmap. Ignore this value if <b>wRegisteredPalette</b> is PAL_UNREGISTERED, or if <b>wCustomBase</b> contains 0xFFFF.
<b>wCustomCnt</b>	Specifies the number of custom colors in the palette. Ignore this value if <b>wRegisteredPalette</b> is PAL_UNREGISTERED, or if this field contains 0xFFFF.
<b>wRsvBase</b>	Specifies the index of the first reserved color of the palette. Reserved colors are those reserved for menus, text, and other screen elements. Reserved colors must be stored contiguously. Ignore this field if it contains 0xFFFF.
<b>wRsvCnt</b>	Specifies the number of reserved entries. Ignore this field if it contains 0xFFFF.
<b>wArtBase</b>	<p>Specifies the index of the first art color of the palette. Art colors are colors used for text and drawing. Art colors consist of a number of hues, each of which has multiple intensities. The various intensities are used for anti-aliasing, a method of using different shades of a color to improve the quality of images displayed on low-resolution devices.</p> <p>For example, if the first art color is red anti-aliased to black with three intensities, the first three entries in the palette would be dark red, medium red, and bright red. The art colors constitute an array, and all hues have the same number of intensities. The user can set both the number of hues and the number of intensities. Ignore these fields if they contain 0xFFFF.</p>
<b>wArtCnt</b>	Specifies the number of art colors. Ignore this field if it contains 0xFFFF.
<b>wNumIntense</b>	Specifies the number of palette entries reserved for the anti-aliased levels of a given art color. This field must be present if <b>wArtBase</b> is present. Ignore this field if it contains 0xFFFF.

---

## Rich Text Format (RTF)

The Rich Text Format (RTF) is a standard method of encoding formatted text and graphics using only 7-bit ASCII characters. Formatting includes different font sizes, faces, and styles, as well as paragraph alignment, justification, and tab control.

RTF is described in the *Microsoft Word Technical Reference: For Windows and OS/2*, published by Microsoft Press.

# Waveform Audio File Format (WAVE)

This section describes the Waveform format, which is used to represent digitized sound.

The WAVE form is defined as follows. Programs must expect (and ignore) any unknown chunks encountered, as with all RIFF forms. However, **<fmt-ck>** must always occur before **<wave-data>**, and both of these chunks are mandatory in a WAVE file.

```
<WAVE-form> →
    RIFF( 'WAVE'
        <fmt-ck>                // Format
        [<fact-ck>]            // Fact chunk
        [<cue-ck>]             // Cue points
        [<playlist-ck>]        // Playlist
        [<assoc-data-list>]    // Associated data list
        <wave-data> )          // Wave data
```

The WAVE chunks are described in the following sections.

## WAVE Format Chunk

The WAVE format chunk **<fmt-ck>** specifies the format of the **<wave-data>**. The **<fmt-ck>** is defined as follows:

```
<fmt-ck> →      fmt(  <common-fields>
                    <format-specific-fields> )

<common-fields> →
    struct
    {
        WORD    wFormatTag;           // Format category
        WORD    wChannels;           // Number of channels
        DWORD   dwSamplesPerSec;     // Sampling rate
        DWORD   dwAvgBytesPerSec;    // For buffer estimation
        WORD    wBlockAlign;         // Data block size
    }
```

The fields in the **<common-fields>** chunk are as follows:

---

Field	Description
<b>wFormatTag</b>	A number indicating the WAVE format category of the file. The content of the <b>&lt;format-specific-fields&gt;</b> portion of the 'fmt' chunk, and the interpretation of the waveform data, depend on this value.  You must register any new WAVE format categories. See "Registering Multimedia Formats" in Chapter 1, "Overview of Multimedia Specifications," for information on registering WAVE format categories.  "Wave Format Categories," following this section, lists the currently defined WAVE format categories.
<b>wChannels</b>	The number of channels represented in the waveform data, such as 1 for mono or 2 for stereo.
<b>dwSamplesPerSec</b>	The sampling rate (in samples per second) at which each channel



should be played.

**dwAvgBytesPerSec**

The average number of bytes per second at which the waveform data should be transferred. Playback software can estimate the buffer size using this value.

**wBlockAlign**

The block alignment (in bytes) of the waveform data. Playback software needs to process a multiple of **wBlockAlign** bytes of data at a time, so the value of **wBlockAlign** can be used for buffer alignment.

---

The <format-specific-fields> consists of zero or more bytes of parameters. Which parameters occur depends on the WAVE format category—see the following section for details. Playback software should be written to allow for (and ignore) any unknown <format-specific-fields> parameters that occur at the end of this field.

## WAVE Format Categories

The format category of a WAVE file is specified by the value of the **wFormatTag** field of the ‘fmt’ chunk. The representation of data in <wave-data>, and the content of the <format-specific-fields> of the ‘fmt’ chunk, depend on the format category.

The currently defined open non-proprietary WAVE format categories are as follows:

---

<b>wFormatTag Value</b>	<b>Format Category</b>
WAVE_FORMAT_PCM (0x0001)	Microsoft Pulse Code Modulation (PCM) format

---

The following are the registered proprietary WAVE format categories:

<b>wFormatTag Value</b>	<b>Format Category</b>
IBM_FORMAT_MULAW (0x0101)	IBM mu-law format
IBM_FORMAT_ALAW (0x0102)	IBM a-law format
IBM_FORMAT_ADPCM (0x0103)	IBM AVC Adaptive Differential Pulse Code Modulation format

The following sections describe the Microsoft WAVE\_FORMAT\_PCM format.

## Pulse Code Modulation (PCM) Format

If the **wFormatTag** field of the <fmt-ck> is set to WAVE\_FORMAT\_PCM, then the waveform data consists of samples represented in pulse code modulation (PCM) format. For PCM waveform data, the <format-specific-fields> is defined as follows:

```
<PCM-format-specific> →
    struct
    {
        WORD    wBitsPerSample;        // Sample size
    }
```

The **wBitsPerSample** field specifies the number of bits of data used to represent each sample of each channel. If there are multiple channels, the sample size is the same for each channel.

For PCM data, the **wAvgBytesPerSec** field of the 'fmt' chunk should be equal to the following formula rounded up to the next whole number:

$$wChannels \times wBitsPerSecond \times \frac{wBitsPerSample}{8}$$

The **wBlockAlign** field should be equal to the following formula, rounded to the next whole number:

$$wChannels \times \frac{wBitsPerSample}{8}$$

### Data Packing for PCM WAVE Files

In a single-channel WAVE file, samples are stored consecutively. For stereo WAVE files, channel 0 represents the left channel, and channel 1 represents the right channel. The speaker position mapping for more than two channels is currently undefined. In multiple-channel WAVE files, samples are interleaved.

The following diagrams show the data packing for a 8-bit mono and stereo WAVE files:

<b>Sample 1</b>	<b>Sample 2</b>	<b>Sample 3</b>	<b>Sample 4</b>
Channel 0	Channel 0	Channel 0	Channel 0

### Data Packing for 8-Bit Mono PCM

Sample 1		Sample 2	
Channel 0 (left)	Channel 1 (right)	Channel 0 (left)	Channel 0 (right)

### Data Packing for 8-Bit Stereo PCM

The following diagrams show the data packing for 16-bit mono and stereo WAVE files:

Sample 1		Sample 2	
Channel 0 low-order byte	Channel 0 high-order byte	Channel 0 low-order byte	Channel 0 high-order byte

### Data Packing for 16-Bit Mono PCM

Sample 1			
Channel 0 (left) low-order byte	Channel 0 (left) high-order byte	Channel 1 (right) low-order byte	Channel 1 (right) high-order byte

### Data Packing for 16-Bit Stereo PCM

#### Data Format of the Samples

Each sample is contained in an integer  $i$ . The size of  $i$  is the smallest number of bytes required to contain the specified sample size. The least significant byte is stored first. The bits that represent the sample amplitude are stored in the most significant bits of  $i$ , and the remaining bits are set to zero.

For example, if the sample size (recorded in **nBitsPerSample**) is 12 bits, then each sample is stored in a two-byte integer. The least significant four bits of the first (least significant) byte is set to zero.

The data format and maximum and minimum values for PCM waveform samples of various sizes are as follows:

Sample Size	Data Format	Maximum Value	Minimum Value
One to eight bits	Unsigned integer	255 (0xFF)	0
Nine or more bits	Signed integer $i$	Largest positive value of $i$	Most negative value of $i$

For example, the maximum, minimum, and midpoint values for 8-bit and 16-bit PCM waveform data are as follows:

Format	Maximum Value	Minimum Value	Midpoint Value
--------	---------------	---------------	----------------

8-bit PCM	255 (0xFF)	0	128 (0x80)
16-bit PCM	32767 (0x7FFF)	-32768 (-0x8000)	0

### Examples of PCM WAVE Files

Example of a PCM WAVE file with 11.025 kHz sampling rate, mono, 8 bits per sample:

```
RIFF( 'WAVE'      fmt(1, 1, 11025, 11025, 1, 8)
      data( <wave-data> ) )
```

Example of a PCM WAVE file with 22.05 kHz sampling rate, stereo, 8 bits per sample:

```
RIFF( 'WAVE'      fmt(1, 2, 22050, 44100, 2, 8)
      data( <wave-data> ) )
```

Example of a PCM WAVE file with 44.1 kHz sampling rate, mono, 20 bits per sample:

```
RIFF( 'WAVE'      INFO(INAM("O Canada"Z))
      fmt(1, 1, 44100, 132300, 3, 20)
      data( <wave-data> ) )
```

## Storage of WAVE Data

The **<wave-data>** contains the waveform data. It is defined as follows:

```
<wave-data> →      { <data-ck> | <data-list> }
<data-ck> →      data( <wave-data> )
<wave-list> →      LIST( 'wavl' {      <data-ck> |      // Wave
samples                                     <silence-ck> }... ) // Silence
<silence-ck> →      slnt( <dwSamples:DWORD> ) // Count of
// silent samples
```

**Note:** The ‘slnt’ chunk represents silence, not necessarily a repeated zero volume or baseline sample. In 16-bit PCM data, if the last sample value played before the silence section is a 10000, then if data is still output to the D to A converter, it must maintain the 10000 value. If a zero value is used, a click may be heard at the start and end of the silence section. If play begins at a silence section, then a zero value might be used since no other information is available. A click might be created if the data following the silent section starts with a nonzero value.

## FACT Chunk

The **<fact-ck>** fact chunk stores important information about the contents of the WAVE file. This chunk is defined as follows:

```
<fact-ck> →      fact( <dwFileSize:DWORD> ) // Number of samples
```

The “fact” chunk is required if the waveform data is contained in a “wavl” LIST chunk and for all compressed audio formats. The chunk is not required for PCM files using the “data” chunk format.

The "fact" chunk will be expanded to include any other information required by future WAVE formats. Added fields will appear following the **<dwFileSize>** field. Applications can use the chunk size field to determine which fields are present.

## Cue-Points Chunk

The **<cue-ck>** cue-points chunk identifies a series of positions in the waveform data stream. The **<cue-ck>** is defined as follows:

```

<cue-ck> →      cue(      <dwCuePoints:DWORD>          // Count of cue points
                    <cue-point>... )                // Cue-point table

<cue-point> →   struct {
                    DWORD   dwName;
                    DWORD   dwPosition;
                    FOURCC  fccChunk;
                    DWORD   dwChunkStart;
                    DWORD   dwBlockStart;
                    DWORD   dwSampleOffset;
                }

```

The **<cue-point>** fields are as follows:

---

Field	Description
<b>dwName</b>	Specifies the cue point name. Each <b>&lt;cue-point&gt;</b> record must have a unique <b>dwName</b> field.
<b>dwPosition</b>	Specifies the sample position of the cue point. This is the sequential sample number within the play order. See "Playlist Chunk," later in this document, for a discussion of the play order.
<b>fccChunk</b>	Specifies the name or chunk ID of the chunk containing the cue point.
<b>dwChunkStart</b>	Specifies the file position of the start of the chunk containing the cue point. This is a byte offset relative to the start of the data section of the 'wav' LIST chunk.
<b>dwBlockStart</b>	Specifies the file position of the start of the block containing the position. This is a byte offset relative to the start of the data section of the 'wav' LIST chunk.
<b>dwSampleOffset</b>	Specifies the sample offset of the cue point relative to the start of the block.

---

## Examples of File Position Values

The following table describes the <cue-point> field values for a WAVE file containing multiple 'data' and 'slnt' chunks enclosed in a 'wavl' LIST chunk:

Cue Point Location	Field	Value
In a 'slnt' chunk	<b>fccChunk</b>	FOURCC value 'slnt'.
	<b>dwChunkStart</b>	File position of the 'slnt' chunk relative to the start of the data section in the 'wavl' LIST chunk.
	<b>dwBlockStart</b>	File position of the data section of the 'slnt' chunk relative to the start of the data section of the 'wavl' LIST chunk.
	<b>dwSampleOffset</b>	Sample position of the cue point relative to the start of the 'slnt' chunk.
In a PCM 'data' chunk	<b>fccChunk</b>	FOURCC value 'data'.
	<b>dwChunkStart</b>	File position of the 'data' chunk relative to the start of the data section in the 'wavl' LIST chunk.
	<b>dwBlockStart</b>	File position of the cue point relative to the start of the data section of the 'wavl' LIST chunk.
	<b>dwSampleOffset</b>	Zero value.
In a compressed 'data' chunk	<b>fccChunk</b>	FOURCC value 'data'.
	<b>dwChunkStart</b>	File position of the start of the 'data' chunk relative to the start of the data section of the 'wavl' LIST chunk.
	<b>dwBlockStart</b>	File position of the enclosing block relative to the start of the data section of the 'wavl' LIST chunk. The software can begin the decompression at this point.
	<b>dwSampleOffset</b>	Sample position of the cue point relative to the start of the block.

The following table describes the <cue-point> field values for a WAVE file containing a single 'data' chunk:

Cue Point Location	Field	Value
Within PCM data	<b>fccChunk</b>	FOURCC value 'data'.
	<b>dwChunkStart</b>	Zero value.
	<b>dwBlockStart</b>	Zero value.
	<b>dwSampleOffset</b>	Sample position of the cue point relative to the start of the 'data' chunk.

In a compressed 'data' chunk	<b>fccChunk</b>	FOURCC value 'data'.
	<b>dwChunkStart</b>	Zero value.
	<b>dwBlockStart</b>	File position of the enclosing block relative to the start of the 'data' chunk. The software can begin the decompression at this point.
	<b>dwSampleOffset</b>	Sample position of the cue point relative to the start of the block.

---

## Playlist Chunk

The **<playlist-ck>** playlist chunk specifies a play order for a series of cue points. The **<playlist-ck>** is defined as follows:

```

<playlist-ck> →   plst(
                   <dwSegments:DWORD>           // Count of play segments
                   <play-segment>... )          // Play-segment table

<play-segment> →  struct {
                   WORD  dwName;
                   WORD  dwLength;
                   WORD  dwLoops;
                   }

```

The **<play-segment>** fields are as follows:

---

Field	Description
<b>dwName</b>	Specifies the cue point name. This value must match one of the names listed in the <b>&lt;cue-ck&gt;</b> cue-point table.
<b>dwLength</b>	Specifies the length of the section in samples.
<b>dwLoops</b>	Specifies the number of times to play the section.

---

## Associated Data Chunk

The **<assoc-data-list>** associated data list provides the ability to attach information like labels to sections of the waveform data stream. The **<assoc-data-list>** is defined as follows:

```

<assoc-data-list> →  LIST( 'adtl'
                          <labl-ck>           // Label
                          <note-ck>           // Note
                          <ltxt-ck>           // Text with
data length           <file-ck> )           // Media file

<labl-ck> →         labl( <dwName:DWORD>
                          <data:ZSTR> )

<note-ck> →        note( <dwName:DWORD>
                          <data:ZSTR> )

```

```

<ltxt-ck> →          ltxt(  <dwName:DWORD>
                        <dwSampleLength:DWORD>
                        <dwPurpose:DWORD>
                        <wCountry:WORD>
                        <wLanguage:WORD>
                        <wDialect:WORD>
                        <wCodePage:WORD>
                        <data:BYTE>... )

<file-ck> →          file(  <dwName:DWORD>
                        <dwMedType:DWORD>
                        <fileData:BYTE>...)

```

## Label and Note Information

The ‘labl’ and ‘note’ chunks have similar fields. The ‘labl’ chunk contains a label, or title, to associate with a cue point. The ‘note’ chunk contains comment text for a cue point. The fields are as follows:

Field	Description
<b>dwName</b>	Specifies the cue point name. This value must match one of the names listed in the <cue-ck> cue-point table.
<b>data</b>	Specifies a NULL-terminated string containing a text label (for the ‘labl’ chunk) or comment text (for the ‘note’ chunk).

## Text with Data Length Information

The “ltxt” chunk contains text that is associated with a data segment of specific length. The chunk fields are as follows:

Field	Description
<b>dwName</b>	Specifies the cue point name. This value must match one of the names listed in the <cue-ck> cue-point table.
<b>dwSampleLength</b>	Specifies the number of samples in the segment of waveform data.
<b>dwPurpose</b>	Specifies the type or purpose of the text. For example, <b>dwPurpose</b> can specify a FOURCC code like ‘scrp’ for script text or ‘capt’ for close-caption text.
<b>wCountry</b>	Specifies the country code for the text. See “Country Codes” in Chapter 2, “Resource Interchange File Format,” for a current list of country codes.
<b>wLanguage, wDialect</b>	Specify the language and dialect codes for the text. See “Language and Dialect Codes” in Chapter 2, “Resource Interchange File Format,” for a current list of language and dialect codes.
<b>wCodePage</b>	Specifies the code page for the text.



## Embedded File Information

The 'file' chunk contains information described in other file formats (for example, an 'RDIB' file or an ASCII text file). The chunk fields are as follows:

---

<b>Field</b>	<b>Description</b>
<b>dwName</b>	Specifies the cue point name. This value must match one of the names listed in the <cue-ck> cue-point table.
<b>dwMedType</b>	Specifies the file type contained in the <b>fileData</b> field. If the <b>fileData</b> section contains a RIFF form, the <b>dwMedType</b> field is the same as the RIFF form type for the file.  This field can contain a zero value.
<b>fileData</b>	Contains the media file.

---

## Chapter 4

# Media Control Interface

---

The Media Control Interface (MCI) is a high-level command control interface to multimedia devices and resource files. MCI provides applications with device-independent capabilities for controlling audio and visual peripherals. Your application can use MCI to control any multimedia device, including audio playback and recording, as well as videodisc and videotape players.

MCI provides a standard command set for playing and recording multimedia devices and resource files. Developers creating multimedia applications are encouraged to use this high-level command interface rather than the low-level functions specific to each platform. The MCI command set acts as a platform-independent layer that sits between multimedia applications and the underlying system software.

The command set is extensible in two ways:

- Developers can incorporate new multimedia devices and file formats in the MCI command set by creating new MCI drivers to interpret the commands.
- New commands and command options can be added to support special features or functions required by new multimedia devices or file formats.

## MCI Command Strings

Using MCI, an application can control multimedia devices using simple command strings like open, play, and close. The MCI commands provide a generic interface to different multimedia devices, reducing the number of commands a developer needs to learn. A multimedia application might even accept MCI commands from an end user and pass them unchanged to the MCI driver, which parses the command and performs the appropriate action.

A set of basic commands is supported by all MCI devices. Developers can also define MCI commands and command options specific to a particular multimedia device or file format. These device-specific commands and command options are needed only when the basic command set does not support a feature specific to the device or file format.

## Example of MCI Command Use

The following example shows a series of MCI commands that play track 6 of an audio compact disc:

```
open cdaudio
set cdaudio time format tmsf
play cdaudio from 6 to 7
close cdaudio
```

The next example shows a similar series of MCI commands that play the first 10,000 samples of a waveform audio file:

```
open c:\mmdata\purplefi.wav type waveaudio alias finch
set finch time format samples
play finch from 1 to 10000 wait
close finch
```

Notice the following:

- The same basic commands (**open**, **play**, and **close**) are used with both devices.
- The **open** command for the “waveaudio” device includes a filename specification. The “waveaudio” device is a compound device (one associated with a media element), while the “cdaudio” device is a simple device (one without an associated media element).
- The **set** commands both specify time formats, but the time format options for the “cdaudio” device are different from those used with the “waveaudio” device.
- The parameters used with the **from** and **to** flags are appropriate to the respective device. For the “cdaudio” device, the parameters specify a range of tracks; for the “waveaudio” device, the parameters specify a range of samples.

## Categories of MCI Command Strings

MCI command strings divide into the following categories:

- *System commands* are interpreted directly by MCI rather than being relayed to a device.
- *Required commands* are recognized by all MCI devices. If a device does not support a required command, it can return “unsupported function” in response to the message.
- *Basic commands* are optional commands. If a device uses a basic command, it must respond to all options for that command. If a device does not use a basic command, it can return “unrecognized command” in response to the message.
- *Extended commands* are specific to a device type or device class; for example, videodisc players. These commands contain both unique commands and extensions to the required and basic commands.

## Command Syntax Conventions

This chapter uses the following documentation conventions:

---

<b>Convention</b>	<b>Description</b>
<b>bold</b>	MCI command or flag keyword.
<i>italics</i>	Command parameter to be replaced with a valid string, number, or rectangle specification.
“quotes”	Parameter text to be typed exactly as shown.
[ brackets ]	Optional flags or parameters

---

## System Commands

The following list summarizes the system commands. MCI supports these commands directly rather than passing them to MCI devices.

---

<b>Message</b>	<b>Description</b>
<b>sound</b>	Play system sounds defined in a system setup file.
<b>sysinfo</b>	Returns information about MCI devices.

---

## Required Commands

The following list summarizes the required commands. All devices recognize these messages. If a device does not support a required command, it can return “unsupported function” in response to the message.

---

<b>Message</b>	<b>Description</b>
<b>capability</b>	Obtains the capabilities of a device.
<b>close</b>	Closes the device.
<b>info</b>	Obtains textual information from a device.
<b>open</b>	Initializes the device.
<b>status</b>	Returns various status information from the device.

---

## Basic Commands

The following list summarizes the basic commands. MCI devices are not required to recognize these commands. If the device does not recognize a basic command, it can return “unrecognized command” in response to the message.

---

<b>Message</b>	<b>Description</b>
<b>load</b>	Recalls data from a disk file.
<b>pause</b>	Stops playing.
<b>play</b>	Starts transmitting output data.
<b>record</b>	Starts recording input data.
<b>resume</b>	Resumes playing or recording from a paused state.
<b>save</b>	Saves data to a disk file.
<b>seek</b>	Seeks forward or backward.
<b>set</b>	Sets the operating state of the device.
<b>status</b>	Obtains status information about the device. (The flags for this command supplement the flags for the command in the required command group.)
<b>stop</b>	Stops playing.

---

## Extended Commands

MCI devices can have additional commands or extend the definition of the required and basic commands. While some extended commands only apply to a specific device driver most of them apply to all devices of a particular type. For example, the MIDI sequencer command set extends the **set** command to add time formats needed by MIDI sequencers. You can find descriptions of extended commands in the command tables in this chapter.

### Extended Commands Reserved for Future Use

The following commands can be defined as extended commands. With the exception of the **delete** command, they are not currently defined for any MCI devices.

---

<b>Message</b>	<b>Description</b>
<b>copy</b>	Copies data to the Clipboard. Parameters and flags for this message vary according to the selected device.
<b>cut</b>	Moves data from the MCI element to the Clipboard. Parameters and flags for this message vary according to the selected device.
<b>delete</b>	Removes data from the MCI element. Parameters and flags for this message vary according to the selected device.

---

## Creating a Command String

There are three components associated with each command string: the command, the name or ID of the device receiving the command, and the command arguments. A command string has the following form:

```
command device_name arguments
```

These components contain the following information:

- The *command* includes a command from the system, required, basic, or extended command set. Examples of commands include **open**, **close**, and **play**.
- The *device\_name* designates the target of the *command*. MCI accepts the names of MCI device types and names of media elements for the *device\_name*. An example of a device name is “cdAudio”.
- The *arguments* specify the flags and parameters used by the *command*. Flags are key words recognized by the MCI command, and parameters are variables associated with the MCI command or flag. Parameters specify variable data values such as filenames, track or frame numbers, or speed values. You can use the following data types for the parameters in a string command:
  - Strings—String data types can be delimited by leading and trailing white space or by matching quotation marks. If MCI encounters a single (unmatched) quotation mark, it ignores the quotation mark. To embed a quote in string, use two quotes (“”). To specify an empty string, you can use double quotes (“”) for the string.
  - Signed long integers—Signed long integer data types are delimited by leading and trailing white space. Unless otherwise specified, integers can be positive or negative. If using negative integers, do not embed white space between the negative sign and the first digit.
  - Rectangle—Rectangle data types are an ordered list of four signed integer values. White space delimits this data type as well as separates each integer in the list.

For example, the **play** command uses the arguments “**from position to position**” to specify starting and ending points for the playback. The **from** and **to** arguments are flags, and the two *position* values are parameters.

For example, the following command string instructs the CD audio player “cdaudio” to play from the start of the waveform to position 500:

```
play cdaudio from 0 to 500
```

Unspecified command arguments assume a default value. For example, if the flag *from* was unspecified in the previous example, the audio player would start playing at the current position.

## About MCI Device Types

Your application identifies an MCI device by specifying an MCI *device type*. A device type indicates the physical type of device. The following table lists the currently defined MCI device types:

---

<b>Device Type</b>	<b>Description</b>
cdaudio <sup>1</sup>	CD audio player
dat	Digital audio tape player
digitalvideo	Digital video in a window (not GDI based)
other	Undefined MCI device
scanner	Image scanner
sequencer <sup>1</sup>	MIDI sequencer
vcr	Videotape recorder or player
videodisc <sup>1</sup>	Videodisc player
waveaudio <sup>1</sup>	Audio device that plays digitized waveform files

---

<sup>1</sup>An extended command set is provided for these devices.

If you have a particular device type installed more than once, the device type names in the system setup file have integers appended to them. This creates unique names for each MCI device type entry. For example, if the “cdaudio” device type is installed twice, the names “cdaudio1” and “cdaudio2” are used to create unique names for each occurrence of the device type. Each name usually refers to a different CD audio player in the system.

## Using MCI Command Strings

The tables at the end of this chapter describe command strings for the MCI devices. The following sections describe commonly used command strings.

### Opening a Device

Before using a device, you must initialize it with the **open** command. The number of devices you can have open depends on the amount of available memory. The **open** command has the following syntax:

```
open device_name [shareable] [type device_type] [alias alias ]
```

The parameters for the **open** command are:

---

<b>Parameters</b>	<b>Description</b>
-------------------	--------------------

---

---

<i>device_name</i>	Specifies the destination device or MCI element name (filename).
<b>shareable</b>	Allows applications to share a common device or device element.
<b>type</b> <i>device_type</i>	Specifies the device when the <i>device_name</i> refers to an MCI element.
<b>alias</b> <i>alias</i>	Specifies an alternate name for the device.

---

MCI classifies device drivers as *compound* and *simple*. Compound device drivers use a *device element*—a media element associated with a device—during operation. For most compound device drivers, the device element is the source or destination data file. For file elements, the element name references a file and its path.

Simple device drivers do not require a device element for playback. For example, compact disc audio device drivers are simple device drivers.

## Opening Simple Devices

Simple devices require only the *device\_name* for operation. You don't need to provide any additional information (such as a name of a data file) to open these devices. For these devices, substitute the name of a device type obtained from the system setup file. For example, you can open a videodisc device with the following command:

```
open videodisc1
```

## Opening Compound Devices

There are three ways to open a compound device:

- By specifying just the device type
- By specifying both the element name and the device type
- By specifying just the element name

To determine the capabilities of a device, you can open a device by specifying only the device type. When opened this way, most compound devices will let you determine their capabilities and close them. For example, you can open the sequencer with the following command:

```
open sequencer
```

To associate a device element with a particular device, you must specify the element name and device type. In the **open** command, substitute the element name for the *device\_name*, add the **type** flag, and substitute the name of the device you want to use for *device\_type*. This combination lets your application specify the MCI device it needs to use. For example, you can open a device element of the waveaudio device with the following command:

```
open right.wav type waveaudio
```

To associate a default MCI device with a device element, you can specify just an element name. In this case, MCI uses the filename extension of the element name to select the device type.



## Using the Shareable Flag

The **shareable** flag lets multiple applications or tasks concurrently access the same device (or element) and device instance.

If your application opens a device or device element without the **shareable** flag, no other application can access it simultaneously. If your application opens a device or device element as shareable, other applications can also access it by also opening it as shareable. The shared device or device element gives each application the ability to change the parameters governing the operating state of the device or device element. Each time that a device or device element is opened as shareable, a unique device ID is returned (even though the device IDs refer to the same instance)

If you make a device or device element shareable, your application should not make any assumptions about the state of a device. When working with shared devices, your application might need to compensate for changes made by other applications using the same services.

If a device can service only one application or task it will fail an open with the **shareable** flag.

While most compound device elements are not shareable, you can open multiple elements (where each element is unique), or you can open a single element multiple times. If you open a single file element multiple times, MCI creates an independent instance for each open device. Each file element opened within a task must have a unique name. The **alias** flag described in the next section lets you use a unique name for each element.

## Using the Alias Flag

The **alias** flag specifies an alternate name for the given device. The alias provides a shorthand notation for compound devices with lengthy pathnames. If your application creates a device alias, it must use the alias rather than the device name for all subsequent references.

## Opening New Device Elements

To create a new device element for a task such as capturing a sound using waveform recording, specify **new** as a *device\_name*. MCI does not save a new file element until you save it with the **save** command. When creating a new file, you must include a device alias with the **open** command. The following commands open a new waveaudio device element, start and stop recording, save the file element, and close the device element:

```
open new type waveaudio alias capture
record capture
stop capture
save capture orca.wav
close capture
```

## Closing a Device

The **close** command releases access to a device or device element. To help MCI manage the devices, your application must explicitly close each device or device element when it is finished with it.

## Shortcuts and Variations for MCI Commands

The MCI string interface lets you use several shortcuts when working with MCI devices.

### Using All as a Device Name

You can specify **all** as a *device\_name* for any command that does not return information. When you specify **all**, the command is sent to all devices opened by your application. For example, “close all” closes all open devices and “play all” starts playing all devices opened by the task. Because MCI sends the commands to each device, there is a delay between when the first device receives the command and when the last device receives the command.

### Combining the Device Type and Device Element Name

You can eliminate the **type** flag in the **open** command if you combine the device type with the device element name. MCI recognizes this combination when you use the following syntax:

*device\_type!element\_name*

The exclamation mark separates the device type from the element name. The following example opens the right.wav element with the waveaudio device:

```
open waveaudio!right.wav
```

### Automatic Open

If MCI cannot identify the *device\_name* as an already open device, MCI tries to automatically open the specified device. Automatic open does not let your application specify the **type** flag. If the device type is not supplied, MCI determines the device type from the element (filename) extensions listed in the system setup file. If you want to use a specific device, you can combine the device type name with the device element name using the exclamation mark.

Only the command-string interface supports automatic open. Automatic open will fail for device-specific commands. For example, a command to unlock the front panel of a videodisc player will fail an automatic open because this capability is specific to the particular videodisc player.

A device that was opened using the automatic open feature will not respond to a command that uses **all** as a device name.

### Automatic Close

MCI automatically closes any device automatically opened using the command-string interface. MCI closes a device when the command completes, when you abort the command, when you request notification with a subsequent command, or when MCI detects a failure.

## Using Wait and Notify Flags

Normally, MCI commands return to the user immediately, even if it takes several minutes to complete the action initiated by the command. For example, after a VCR device receives a rewind command, it returns before the tape has finished rewinding. You can use either of the following required MCI flags to modify this default behavior:

---

<b>Flag</b>	<b>Description</b>
<b>notify</b>	Directs the device to send an MM_MCINOTIFY message to a window when the requested action is complete.
<b>wait</b>	Directs the device to wait until the requested action is complete before returning to the application.

---

## Using the Notify Flag

The **notify** flag directs the device to post an MM\_MCINOTIFY message when the device completes an action. Your application must have a window procedure to process the MM\_MCINOTIFY message for notification to have any effect. While the results of a notification are application-dependent, the application's window procedure can act upon four possible conditions associated with the notify message:

- Notification will occur when the notification conditions are satisfied.
- Notification can be superseded.
- Notification can be aborted.
- Notification can fail.

A successful notification occurs when the conditions required for initiating the callback are satisfied and the command completed without interruption.

A notification is superseded when the device has a notification pending and you send it another notify request. When a notification is superseded, MCI resets the callback conditions to correspond to the notify request of the new command.

A notification is aborted when you send a new command that prevents the callback conditions set by a previous command from being satisfied. For example, sending the stop command cancels a notification pending for the “play to 500” command. If your command interrupts a command that has a notification pending, and your command also requests notification, MCI will abort the first notification immediately and respond to the second notification normally.

A notification fails if a device error occurs while a device is executing the MCI command. For example, MCI posts this message when a hardware error occurs during a play command.

## Obtaining Information From MCI Devices

Every device responds to the **capability**, **status**, and **info** commands. These commands obtain information about the device. For example, your application can determine if a videodisc requires a device element using the following command:

```
capability videodisc compound file
```

For most videodisc devices, this example would return **false**. The flags listed for the required and basic commands provide a minimum amount of information about a device. Many devices supplement the required and basic flags with extended flags to provide additional information about the device.

When you request information with the **capability**, **status**, or **info** command, the argument list can contain only one flag requesting information. The string interface can only return one string or value in response to a **capability**, **status**, or **info** command.

## The Play Command

The **play** command starts playing a device. Without any flags, the **play** command starts playing from the current position and plays until the command is halted or until the end of the media or file is reached. For example, “play cdaudio” starts playing an audio disc from the position where it was stopped.

Most devices support the **play** command also support the **from** and **to** flags. These flags indicate the position at which the device should start and stop playing. For example, “play cdaudio from 0” plays the audio disc from the beginning of the first track. The units assigned to the position value depend on the device. For example, the position is normally specified in frames for CAV videodiscs, and milliseconds for digital audio.

As an extended command, devices add flags to use the capabilities of a particular device. For example, the **play** command for videodisc players adds the flags **fast**, **slow**, **reverse**, and **scan**.

## Stop, Pause, and Resume Commands

The **stop** command suspends the playing or recording of a device. Many devices include the basic command **pause**, which also suspends these sessions. The difference between **stop** and **pause** depends on the device. Usually **pause** suspends operation but leaves the device ready to resume playing or recording immediately.

Using **play** or **record** to restart a device will reset the **to** and **from** positions specified before the device was paused or stopped. Without the **from** flag, these commands reset the start position to the current position. Without the **to** flag, they reset the end position to the end of the media. If you want to continue playing or recording but want to stop at a position previously specified, use the **to** flag with these commands and repeat the position value.

Some devices include the **resume** command to restart a paused device. This command does not change the **to** and **from** positions specified with the **play** or **record** command, which preceded the pause command.

# MCI System Commands

The following commands are interpreted directly by MCI. The remaining command tables list commands interpreted by the devices.

---

Command	Description
<b>sound</b>	The device name of this command specifies a sound defined in a system setup file.. If it is not found, MCI uses a system default sound.
<b>sysinfo</b> <i>item</i>	Obtains MCI system information. One of the following <i>items</i> modifies <b>sysinfo</b> :
<b>installname</b>	Returns the name used to install the device.
<b>quantity</b>	Returns the number of MCI devices of the type specified by the device-name field. The device-name field must contain a standard MCI device type. Any digits after the name are ignored. The special device name <b>all</b> returns the total number of MCI devices in the system.
<b>quantity open</b>	Returns the number of open MCI devices of the type specified by the device name. The device name must be a standard MCI device type. Any digits after the name are ignored. The special device name <b>all</b> returns the total number of MCI devices in the system that are open.
<b>name</b> <i>index</i>	Returns the name of an MCI device. The <i>index</i> ranges from 1 to the number of devices of that type. If <b>all</b> is specified for the device name, <i>index</i> ranges from 1 to the total number of devices in the system.
<b>name</b> <i>index</i> <b>open</b>	Returns the name of an open MCI device. The <i>index</i> ranges from 1 to the number of devices of that type. If <b>all</b> is specified for the device name, <i>index</i> ranges from 1 to the total number of devices in the system.

---

# Required Commands for All Devices

The following commands are recognized by all devices. Extended commands can add other options to these commands. A list of the errors common to all the commands follows the required command table.

---

Command	Description
<b>capability</b> <i>item</i>	Requests information about a particular capability of a device. While other capabilities are defined for specific devices and device types, the following <i>items</i> are always available:  <b>can eject</b> Returns <b>true</b> if the device can eject the media. <b>can play</b> Returns <b>true</b> if the device can play. <b>can record</b> Returns <b>true</b> if the device supports recording. <b>can save</b> Returns <b>true</b> if the device can save data. <b>compound device</b> Returns <b>true</b> if the device requires an element name. <b>device type</b> Returns one of the following: <b>audio tape</b> <b>cdaudio</b> <b>digital audio tape</b> <b>scanner</b> <b>sequencer</b> <b>videodisc</b> <b>videotape</b> <b>waveaudio</b> <b>has audio</b> Returns <b>true</b> if the device supports audio playback. <b>has video</b> Returns <b>true</b> if the device supports video. <b>uses files</b> Returns <b>true</b> if the element of a compound device is a file pathname.
<b>close</b>	When sent to a simple device, closes the device. When sent to a compound device element, closes the element and any resources associated with it. MCI unloads a device when it is no longer being used.
<b>info</b> <i>item</i>	Fills a user-supplied buffer with a NULL-terminated string containing textual information. One of the following <i>item</i> modifies <b>info</b> :  <b>product</b> Returns a description of the hardware associated with a device. This usually includes the manufacturer and model information.
<b>open</b> <i>items</i>	Initializes the device. The following optional <i>items</i> modify <b>open</b> :  <b>alias</b> <i>device alias</i> Specifies an alternate name for the given device. If specified, it must be used for subsequent references.  <b>shareable</b> Initializes the device or element as shareable. Subsequent attempts to open it fail unless you specify <b>shareable</b> in both the original and subsequent <b>open</b> commands. MCI returns an error

		if it is already open and not shareable.
	<b>type</b> <i>device type</i>	Specifies the compound device that controls a device element. As an alternative to <b>type</b> , MCI can use the filename extension entries to select the device based on the extension used by the device element.
<b>status</b> <i>item</i>	Obtains status information for the device. One of the following <i>items</i> modifies <b>status</b> :	
	<b>mode</b>	Returns the current mode of the device.  Commonly supported standard modes are: <b>not ready</b> , <b>paused</b> , <b>playing</b> , <b>stopped</b> , <b>open</b> , <b>recording</b> , and <b>seeking</b>
	<b>ready</b>	Returns <b>true</b> if the device is ready.

---

## Basic Commands for Specific Device Types

In addition to the commands described previously, each device supports a set of commands specific to its device type. Where possible, these type-specific commands are identical between types. When type-specific commands are common to multiple devices, they are considered basic commands. For example, the basic **play** command is identical for videodisc and videotape players. Other basic commands are listed in the following table. Although these commands are optional for a device, if a command is used it must recognize all options listed in this table. The options generally provide for a minimum set of capabilities, but some devices may return “unsupported function” if an option is used which clearly doesn’t apply.

---

Command	Description
<b>load</b> <i>item</i>	Load a device element from disk. The following optional <i>item</i> modifies <b>load</b> :  <i>filename</i> Specifies the source path and file.
<b>pause</b>	Pauses playing or recording.
<b>play</b> <i>items</i>	Start playing the device. The following optional <i>items</i> modify <b>play</b> :  <b>from</b> <i>position</i> Specifies the position to start and stop playing. <b>to</b> <i>position</i> If <b>from</b> is omitted, the play starts from the current position; if <b>to</b> is omitted, the play stops at the end of the media.

<b>record</b> <i>items</i>	Start recording data. All data recorded after a file is opened is discarded if the file is closed without saving it. The following optional <i>items</i> modify <b>record</b> :
<b>insert</b>	Specifies that new data is added to the device element at the current position.
<b>from</b> <i>position</i> <b>to</b> <i>position</i>	Specifies the positions to start and stop recording. If <b>from</b> is omitted, the device starts recording at the current position; if <b>to</b> is omitted, the device records until a <b>stop</b> or <b>pause</b> command is received.
<b>overwrite</b>	Specifies that new data will replace data in the device element.
	The default recording mode ( <b>insert</b> or <b>overwrite</b> ) depends on the specific device. Each device should define a default recording mode.
<b>resume</b>	Resumes playing or recording following a pause.
<b>save</b> <i>item</i>	Saves the MCI element. The following optional <i>item</i> modifies <b>save</b> :
<i>filename</i>	Specifies the destination path and file.
<b>seek</b> <i>item</i>	Moves to the specified position and stops. One of the following is required for <i>item</i> :
<b>to</b> <i>position</i>	Specifies the position to stop the seek.
<b>to start</b>	Seeks to the start of the media or device element.
<b>to end</b>	Seeks to the end of the media or device element.
<b>set</b> <i>items</i>	Sets the various control <i>items</i> :
<b>audio all off</b> <b>audio all on</b>	Enables or disables audio output
<b>audio left off</b> <b>audio left on</b>	Enables or disables output to the left audio channel.
<b>audio right off</b> <b>audio right on</b>	Enables or disables output to the right audio channel.
<b>door closed</b>	Loads the media and closes the door if possible.
<b>door open</b>	Opens the door and ejects the tray if possible.
<b>time format</b> <b>milliseconds</b>	Sets time format to milliseconds. All position information is this format after this command. You can abbreviate milliseconds as <b>ms</b> .
<b>video off</b> <b>video on</b>	Enables or disables video output.
<b>status</b> <i>item</i>	Obtains status information for the device. One of the following <i>items</i> modifies <b>status</b> :
<b>current track</b>	Returns the current track.
<b>length</b>	Returns the total length of the segment.



<b>length track</b> <i>track_number</i>	Returns the length of the serial track specified by <i>track_number</i> .
<b>number of tracks</b>	Returns the number of tracks on the media.
<b>position</b>	Returns the current position.
<b>position track</b> <i>track_number</i>	Returns the position of the start of the track specified by <i>track_number</i> .
<b>start position</b>	Returns the starting position of the media or device element.
<b>time format</b>	Returns the time format.
<b>stop</b>	Stops the device.

---

# CD Audio (Redbook) Commands

The CD audio command set provides a common method for playing CD audio sequences. CD audio devices support the following core set of commands:

---

Command	Description
<b>capability</b> <i>item</i>	Requests information about the capabilities of the CD audio device. One of the following <i>items</i> is required: <ul style="list-style-type: none"><li><b>can eject</b> Returns <b>true</b> if the CD audio device can eject the media.</li><li><b>can play</b> Returns <b>true</b> if the CD audio device can play the media.</li><li><b>can record</b> Returns <b>false</b>.</li><li><b>can save</b> Returns <b>false</b>.</li><li><b>compound device</b> Returns <b>false</b>.</li><li><b>device type</b> Returns <b>CDaudio</b>.</li><li><b>has audio</b> Returns <b>true</b>.</li><li><b>has video</b> Returns <b>false</b>.</li><li><b>uses files</b> Returns <b>false</b>.</li></ul>
<b>close</b>	Closes the device.
<b>info</b> <i>item</i>	Fills a user-supplied buffer with a NULL-terminated string containing textual information. One of the following optional <i>item</i> modifies <b>info</b> : <ul style="list-style-type: none"><li><b>product</b> Returns the product name and model of the current audio device.</li></ul>
<b>open</b> <i>items</i>	Initializes the device. MCI reserves <code>cdaudio</code> for the compact disc audio device type. The following optional <i>items</i> modify <b>open</b> : <ul style="list-style-type: none"><li><b>alias</b> <i>device_alias</i> Specifies an alternate name for the given device. If specified, it must also be used for subsequent references.</li><li><b>shareable</b> Initializes the device as shareable. Subsequent attempts to open it fail unless you specify <b>shareable</b> in both the original and subsequent <b>open</b> commands. MCI returns an error if it is already open and not shareable.</li></ul>
<b>pause</b>	Pauses playing.
<b>play</b> <i>items</i>	Starts playing audio. The following optional <i>items</i> modify <b>play</b> : <ul style="list-style-type: none"><li><b>from</b> <i>position</i> to <i>position</i> Specifies the position to start and stop playing.</li></ul>
<b>resume</b>	Resumes playing from a paused state.
<b>seek</b> <i>item</i>	Moves to the specified location on the disc. If already playing or recording, the device is stopped. One of the following <i>items</i> modifies <b>seek</b> : <ul style="list-style-type: none"><li><b>to</b> <i>position</i> Specifies the destination position for the seek.</li></ul>

		If it is greater than the length of the disc, an out-of-range error is returned.
	<b>to start</b>	Specifies to seek to the start of the audio data on the CD.
	<b>to end</b>	Specifies to seek to the end of the audio data on the CD.
<b>set items</b>		Sets the various control <i>items</i> :
	<b>audio all off</b> <b>audio all on</b>	Enables or disables audio output..
	<b>audio left off</b> <b>audio left on</b>	Enables or disables output to the left audio channel.
	<b>audio right off</b> <b>audio right on</b>	Enables or disables output to the right audio channel.
	<b>door closed</b>	Retracts the tray and closes the door if possible.
	<b>door open</b>	Opens the door and ejects the tray if possible.
	<b>time format milliseconds</b>	Sets the time format to milliseconds. All position information is this format after this command. You can abbreviate milliseconds as <b>ms</b> .
	<b>time format msf</b>	Sets the time format to <i>mm:ss:ff</i> , where <i>mm</i> is minutes, <i>ss</i> is seconds, and <i>ff</i> is frames. All position information is in this format after this command. On input, <i>ff</i> can be omitted if 0, and <i>ss</i> can be omitted if both it and <i>ff</i> are 0. These fields have the following maximum values:  Minutes 99 Seconds 59 Frames 74
	<b>time format tmsf</b>	Sets the time format to <i>tt:mm:ss:ff</i> where “ <i>tt</i> ” is tracks, “ <i>mm</i> ” is minutes, “ <i>ss</i> ” is seconds, and “ <i>ff</i> ” is frames. All position information is in this format after this command. On input “ <i>ff</i> ” can be omitted if 0, “ <i>ss</i> ” can be omitted if both it and “ <i>ff</i> ” are 0, and “ <i>mm</i> ” can be omitted if it, “ <i>ss</i> ” and “ <i>ff</i> ” are 0. These fields have the following maximum values:  Tracks 99 Minutes 99 Seconds 59 Frames 74
<b>status item</b>		Obtains status information for the device. One of the following <i>items</i> modifies <b>status</b> :
	<b>current track</b>	Returns the current track.
	<b>length</b>	Returns the total length of the disc.
	<b>length track</b> <i>track_number</i>	Returns the length of the track specified by <i>track_number</i> .
	<b>media present</b>	Returns <b>true</b> if the CD is inserted in the drive; otherwise, it returns <b>false</b> .

<b>mode</b>	Returns <b>not ready</b> , <b>open</b> , <b>paused</b> , <b>playing</b> , <b>seeking</b> , or <b>stopped</b> for the current mode of the drive.
<b>number of tracks</b>	Returns the number of tracks on the CD.
<b>position</b>	Returns the current position.
<b>position track</b> <i>track_no</i>	Returns the starting position of the track specified by <i>track_no</i> .
<b>ready</b>	Returns <b>true</b> if the drive is ready.
<b>start position</b>	Returns the starting position of the CD.
<b>time format</b>	Returns the current time format.
<b>stop</b>	Stops playing.

---

# MIDI Sequencer Commands

The MIDI sequencer supports the following set of commands:

---

Command	Description
<b>capability</b> <i>item</i>	Requests additional information about the capabilities of the MIDI sequencer. One of the following <i>items</i> is required: <ul style="list-style-type: none"><li><b>can eject</b> Returns <b>false</b>.</li><li><b>can play</b> Returns <b>true</b> if the sequencer can play.</li><li><b>can record</b> Returns <b>true</b> if the sequencer can record MIDI data.</li><li><b>can save</b> Returns <b>true</b> if the sequencer can save MIDI data.</li><li><b>compound device</b> Generally returns <b>true</b>; most sequencers are compound devices..</li><li><b>device type</b> Returns <b>sequencer</b>.</li><li><b>has audio</b> Returns <b>true</b>.</li><li><b>has video</b> Returns <b>false</b>.</li><li><b>uses files</b> Returns <b>true</b>.</li></ul>
<b>close</b>	Closes the sequencer element and the port and file associated with it.
<b>info</b> <i>item</i>	Fills a user-supplied buffer with a NULL-terminated string containing textual information. One of the following optional <i>item</i> modifies <b>info</b> : <ul style="list-style-type: none"><li><b>product</b> Returns the product name of the current MIDI sequencer.</li></ul>
<b>open</b> <i>items</i>	Initializes the sequencer. The following optional <i>items</i> modify <b>open</b> : <ul style="list-style-type: none"><li><b>alias</b> <i>device_alias</i> Specifies an alternate name for the sequencer element. If specified, it must also be used for subsequent references.</li><li><b>shareable</b> Initializes the sequencer element as shareable. Subsequent attempts to open it fail unless you specify shareable in both the original and subsequent <b>open</b> commands. MCI returns an invalid device error if it is already open and not shareable.</li><li><b>type</b> <i>device_type</i> MCI reserves sequencer for the MIDI sequencer device type. As an alternative to <b>type</b>, MCI can use the element filename extension entries to select the sequencer.</li></ul>
<b>pause</b>	Pauses playing.
<b>play</b> <i>items</i>	Starts playing the sequencer. The following optional <i>items</i> modify <b>play</b> : <ul style="list-style-type: none"><li><b>from</b> <i>position</i> Specifies the positions to start and stop playing. If <b>from</b> is omitted, play starts at the current position; if <b>to</b> is omitted, play stops at</li></ul>

the end of the file.

**record items**

Starts recording MIDI data. All data recorded after a file is opened is discarded if the file is closed without saving it. The following optional *items* modify **record**:

- insert** Specifies that new data is added to the device element.
- from position**  
**to position** Specifies the positions to start and stop recording. If **from** is omitted, the device starts recording at the current position; if **to** is omitted, the device records until a **stop** or **pause** command is received.
- overwrite** Specifies that new data will replace data in the device element.

**resume**

Resumes playing or recording following a pause.

**save item**

Saves the MCI element. The following item *modifies save*:

- filename* The *filename* specifies the destination path and file.

**seek item**

Moves to the specified position in the file. One of the following *items* is required:

- to position** Specifies the final position for the seek.
- to start** Specifies to seek to the start of the sequence.
- to end** Specifies to seek to the end of the sequence.

**set items**

Sets the various control *items*:

- audio all off**  
**audio all on** Enables or disables audio output..
- audio left off**  
**audio left on** Enables or disables output to the left audio channel.
- audio right off**  
**audio right on** Enables or disables output to the right audio channel.
- master MIDI** Sets the MIDI sequencer as the synchronization source. Synchronization data is sent in MIDI format.
- master none** Inhibits the sequencer from sending synchronization data.
- master SMPTE** Sets the MIDI sequencer as the synchronization source. Synchronization data is sent in SMPTE format.
- offset time** Sets the SMPTE offset *time* in colon form (*hours:minutes:seconds:frames*). The offset is the beginning time of a SMPTE based sequence.
- port port\_number** Sets the MIDI port receiving the MIDI messages. This command will fail if the port you are trying to open is being used by another application.
- port mapper** Sets the MIDI mapper as the port receiving the MIDI messages. This command will fail if

the MIDI mapper or a port it needs is being used by another application.

<b>port none</b>	Disables the sending of MIDI messages..
<b>slave file</b>	Sets the MIDI sequencer to use file data as the synchronization source. This is the default.
<b>slave MIDI</b>	Sets the MIDI sequencer to use incoming data MIDI for the synchronization source. The sequencer recognizes synchronization data with the MIDI format.
<b>slave none</b>	Sets the MIDI sequencer to ignore synchronization data.
<b>slave SMPTE</b>	Sets the MIDI sequencer to use incoming MIDI data for the synchronization source. The sequencer recognizes synchronization data with the SMPTE format.
<b>tempo</b> <i>tempo_value</i>	Sets the tempo of the sequence according to the current time format. For a ppqn-based file, the <i>integer</i> is interpreted as beats per minute. For a SMPTE-based file, the <i>integer</i> is interpreted as frames per second.
<b>time format milliseconds</b>	Sets time format to milliseconds. All position information is specified as milliseconds following this command. The sequence file sets the default format to ppqn or SMPTE. You can abbreviate milliseconds as <b>ms</b> .
<b>time format song pointer</b>	Sets time format to song pointer (sixteenth notes). This can only be performed for a sequence of division type ppqn.
<b>time format SMPTE 24</b>	Sets time format to SMPTE 24 frame rate. All position information is specified in SMPTE format following this command. The sequence file sets the default format to ppqn or SMPTE.
<b>time format SMPTE 25</b>	Sets time format to SMPTE 25 frame rate. All position information is specified in SMPTE format following this command. The sequence file sets the default format to ppqn or SMPTE.
<b>time format SMPTE 30</b>	Sets time format to SMPTE 30 frame rate. All position information is specified in SMPTE format following this command. The sequence file sets the default format to ppqn or SMPTE.
<b>time format SMPTE 30 drop</b>	Sets time format to SMPTE 30 drop frame rate. All position information is specified in SMPTE format following this command. The sequence file sets the default format to ppqn or SMPTE.

**status item**

Obtains status information for the MIDI sequencer. One of the following *items* modifies **status**:

<b>current track</b>	Returns the current track number.
<b>division type</b>	Returns one of the following file division type: PPQN, SMPTE 24 frame, SMPTE 25

	frame, SMPTE 30 drop frame, or SMPTE 30 frame. Use this information to determine the format of the MIDI file, and the meaning of tempo and position information.
<b>length</b>	Returns the length of a sequence in the current time format. For ppqn files, this will be song pointer units. For SMPTE files, this will be in colon form ( <i>hours:minutes:seconds:frames</i> ).
<b>length track</b> <i>track_number</i>	Returns the length of a sequence using the current time format. For ppqn files, this will be song pointer units. For SMPTE files, this will be in colon form ( <i>hours:minutes:seconds:frames</i> ).
<b>master</b>	Returns <b>midi</b> , <b>none</b> , or <b>smppte</b> depending on the type of synchronization set.
<b>media present</b>	The sequencer returns <b>true</b> .
<b>mode</b>	Returns <b>not ready</b> , <b>paused</b> , <b>playing</b> , <b>seeking</b> , or <b>stopped</b> .
<b>number of tracks</b>	Returns the number of tracks.
<b>offset</b>	Returns the offset of a SMPTE-based file. The time is returned in colon form ( <i>hours:minutes:seconds:frames</i> ). The offset is the starting time of a SMPTE based sequence.
<b>port</b>	Returns the MIDI port number assigned to the sequence.
<b>position</b>	Returns the current position of a sequence in the current time format. For ppqn files, this will be song pointer units. For SMPTE files, this will be in colon form ( <i>hours:minutes:seconds:frames</i> ).
<b>position track</b> <i>track_number</i>	Returns the current position of the track specified by <i>track_number</i> in the current time format. For ppqn files, this will be song pointer units. For SMPTE files, this will be in colon form ( <i>hours:minutes:seconds:frames</i> ).
<b>ready</b>	Returns <b>true</b> if the device is ready.
<b>slave</b>	Returns <b>file</b> , <b>midi</b> , <b>none</b> , or <b>smppte</b> depending on the type of synchronization set.
<b>start position</b>	Returns the starting position of the media or device element.
<b>tempo</b>	Returns the current tempo of a sequence in the current time format. For files with ppqn format, the tempo is in beats per minute. For files with SMPTE format, the tempo is in frames per second.
<b>time format</b>	Returns the time format.
<b>stop</b>	Stops playing.

---



# Videodisc Player Commands

Videodisc players support the following core set of commands:

---

Command	Description
<b>capability</b> <i>item</i>	Reports the capabilities of the device. The device should report capabilities according to the type of disc (CAV or CLV) inserted in the drive. If no disc is inserted, the device should assume CAV. One of the following optional <i>items</i> modifies <b>capability</b> : <ul style="list-style-type: none"><li><b>can eject</b> Returns <b>true</b> if the device can eject the media.</li><li><b>can play</b> Returns <b>true</b> if the device supports playing.</li><li><b>can record</b> Returns <b>true</b> if the video device can record.</li><li><b>can reverse</b> Returns <b>true</b> if the device can play in reverse, <b>false</b> otherwise. This is always <b>false</b> if a CLV disc is inserted.</li><li><b>can save</b> Returns <b>false</b>.</li><li><b>compound device</b> Returns <b>false</b>.</li><li><b>device type</b> Returns <b>videodisc</b>.</li><li><b>fast play rate</b> Returns the standard fast play rate of the player in frames per second. Returns 0 if the device cannot play fast.</li><li><b>has audio</b> Returns <b>true</b> if the videodisc player has audio.</li><li><b>has video</b> Returns <b>true</b>.</li><li><b>media type</b> Returns <b>CAV</b>, <b>CLV</b>, or <b>other</b>, depending on the type of videodisc.</li><li><b>normal play rate</b> Returns the normal play rate in frames per second. Returns 0 for CLV discs.</li><li><b>slow play rate</b> Returns the standard slow play rate in frames per second. Returns 0 if the device cannot play slow.</li><li><b>uses files</b> Returns <b>false</b>.</li></ul>
<b>close</b>	Closes the device.
<b>escape</b> <i>item</i>	Sends custom information to a device. The following <i>item</i> modifies <b>escape</b> : <ul style="list-style-type: none"><li><i>string</i> Specifies the custom information sent to the device.</li></ul>
<b>info</b> <i>item</i>	Fills a user-supplied buffer with a NULL-terminated string containing textual information. The following optional <i>item</i> modifies <b>info</b> : <ul style="list-style-type: none"><li><b>product</b> Returns the product name of the device that the peripheral is controlling.</li></ul>
<b>open</b> <i>items</i>	Initializes the device. MCI reserves video disc for the videodisc device type. The following optional <i>items</i> modify <b>open</b> : <ul style="list-style-type: none"><li><b>alias device_ alias</b> Specifies an alternate name for the given</li></ul>

		device. If specified, it must also be used for subsequent references.
	<b>shareable</b>	Initializes the device as shareable. Subsequent attempts to open it fail unless you specify <b>shareable</b> in both the original and subsequent <b>open</b> commands. MCI returns an invalid device error if it is already open and not shareable.
<b>pause</b>		Stops playing. If a CAV disc is playing, it also freezes the video frame. If a CLV disc is playing, the player is stopped.
<b>play items</b>		Starts playing. The following optional <i>items</i> modify <b>play</b> :
	<b>fast</b>	Indicates that the device should play faster or slower than normal. To determine the exact speed on a particular player, use the <b>status speed</b> command. To specify the speed more precisely, use the <b>fps</b> flag. <b>Slow</b> applies only to CAV discs.
	<b>slow</b>	
	<b>from position</b> <b>to position</b>	Specifies the positions to start and stop playing. Positions are in frames for CAV discs and in seconds for CLV discs, unless <b>chapter</b> is also used (in which case, the position is given in chapters). If <b>from</b> is omitted, play starts at the current position; if <b>to</b> is omitted, the play stops at the end of the disc.
	<b>reverse</b>	Sets the play direction to backwards. This applies only to CAV discs.
	<b>scan</b>	Indicates the play speed is as fast as possible, possibly with audio disabled. This applies only to CAV discs.
	<b>speed integer</b>	Specifies the rate of play. Currently supported speed values are measured in frames per second, which is the default. This applies only to CAV discs.
<b>resume</b>		Resumes playing.
<b>seek item</b>		Searches using fast forward or fast reverse with video and audio off. The following optional <i>items</i> modify <b>seek</b> :
	<b>reverse</b>	Indicates the seek direction on CAV discs is backwards. This modifier is invalid if <b>to</b> is specified.
	<b>to position</b>	Specifies the end position to stop the seek. If <b>to</b> is not specified, the seek continues until the end of the media is reached.
	<b>to start</b>	Specifies to seek to the start of the disc.
	<b>to end</b>	Specifies to seek to the end of the disc.
<b>set items</b>		Sets the various control <i>items</i> :
	<b>audio all off</b> <b>audio all on</b>	Enables or disables audio output.
	<b>audio left off</b> <b>audio left on</b>	Enables or disables output to the left audio channel.

	<b>audio right off</b>	Enables or disables output to the right audio channel.
	<b>audio right on</b>	
	<b>door open</b>	Opens the door and ejects the tray, if possible.
	<b>door closed</b>	Retracts the tray and closes the door, if possible.
	<b>time format frames</b>	Sets the position format to frames on CAV discs. All position information is specified in this format following this command. This is the default for CAV discs.
	<b>time format hms</b>	Sets position format to <i>h:mm:ss</i> where <i>h</i> is hours, <i>mm</i> is minutes, and <i>ss</i> is seconds. All position information is specified in this format following this command. On input, <i>h</i> may be omitted if 0, and <i>mm</i> may be omitted if both it and <i>h</i> are 0. This is the default for CLV discs.
	<b>time format milliseconds</b>	Sets the position format to milliseconds. All position information is specified in this format following this command. You can abbreviate milliseconds as <b>ms</b> .
	<b>time format track</b>	Sets the position format to tracks (chapters). All position information is specified in this format following this command.
	<b>video on</b>	Turns the video on or off.
	<b>video off</b>	
<b>spin item</b>		Starts the disc spinning or stops the disc from spinning. One of the following <i>items</i> modifies <b>status</b> :
	<b>down</b>	Stops the disc from spinning.
	<b>up</b>	Starts the disc spinning.
<b>status item</b>		Obtains status information for the device. One of the following <i>items</i> modifies <b>status</b> :
	<b>current track</b>	Returns the current track (chapter) number.
	<b>disc size</b>	Returns either 8 or 12 to indicate the size of the loaded disc in inches.
	<b>forward</b>	Returns <b>true</b> if the play direction is forward or if the device is not playing; <b>false</b> if the play direction is backward.
	<b>length</b>	Returns the total length of the segment.
	<b>length track</b> <i>track_number</i>	Returns the length of the track (chapter) specified by <i>track_number</i> .
	<b>media present</b>	Returns <b>true</b> if a disc is inserted in the device, <b>false</b> otherwise.
	<b>media type</b>	Returns either <b>CAV</b> , <b>CLV</b> , or <b>other</b> depending on the type of videodisc.
	<b>mode</b>	Returns <b>not ready</b> , <b>opened</b> , <b>paused</b> , <b>parked</b> , <b>playing</b> , <b>scanning</b> , <b>seeking</b> , or <b>stopped</b> .
	<b>number of tracks</b>	Returns the number of tracks (chapters) on the media.

<b>position</b>	Returns the current position.
<b>position track</b> <i>track_number</i>	Returns the position of the start of the track (chapter) specified by <i>track_number</i> .
<b>ready</b>	Returns <b>true</b> if the device is ready.
<b>side</b>	Returns 1 or 2 to indicate which side of the disc is loaded.
<b>speed</b>	Returns the speed in frames per second.
<b>start position</b>	Returns the starting position of the disc.
<b>time format</b>	Returns the time format.

**step** *items*

Step the play one or more frames forward or backward. The default action is to step one frame forward. The **step** command applies only to CAV discs. The following *items* modifies **step**:

<b>by frames</b>	Specifies the number of <i>frames</i> to step. If a negative value is used, the <b>reverse</b> flag is ignored.
<b>reverse</b>	Step backward.

**stop** Stop playing.

---

# Waveform Audio Commands

Waveform audio drivers must support the following core set of commands:

Command	Description
<b>capability</b> <i>item</i>	Requests additional information about the capabilities of the waveform audio driver. One of the following <i>items</i> modify <b>capability</b> : <ul style="list-style-type: none"><li><b>can eject</b> Returns <b>false</b>.</li><li><b>can play</b> Returns <b>true</b> if the device can play. The wave audio device returns true if an output device is available.</li><li><b>can record</b> Returns <b>true</b> if the waveform driver can record. The waveform audio device returns <b>true</b> if an input device is available.</li><li><b>can save</b> Returns <b>true</b> if the wave audio device can save data.</li><li><b>compound device</b> Generally returns <b>true</b>; most waveform audio devices are compound devices.</li><li><b>device type</b> Returns <b>waveaudio</b>.</li><li><b>has audio</b> Returns <b>true</b></li><li><b>has video</b> Returns <b>false</b>.</li><li><b>inputs</b> Returns the total number of input devices.</li><li><b>outputs</b> Returns the total number of output devices.</li><li><b>uses files</b> Returns <b>true</b>.</li></ul>
<b>close</b>	Closes the device element and any resources associated with it.
<b>cue</b> <i>item</i>	Prepares for playing or recording. The <b>cue</b> command does not have to be issued prior to playing or recording. However, depending on the device, it might reduce the delay associated with the <b>play</b> or <b>record</b> command. This command fails if playing or recording is in progress. The <i>item</i> is one of the following: <ul style="list-style-type: none"><li><b>input</b> Prepares for recording.</li><li><b>output</b> Prepares for playing. This is the default.</li></ul>
<b>delete</b> <i>items</i>	Deletes a data segment from the MCI element. The following optional <i>items</i> modify <b>delete</b> : <ul style="list-style-type: none"><li><b>from</b> <i>position</i> <b>to</b> <i>position</i> Specifies the positions to start and stop deleting data. If <b>from</b> is omitted, the deletion starts at the current position; if <b>to</b> is omitted, the deletion stops at the end of the file or waveform.</li></ul>
<b>info</b> <i>item</i>	Fills a user-supplied buffer with a NULL-terminated string containing textual information. One of the following <i>items</i> modifies <b>info</b> : <ul style="list-style-type: none"><li><b>file</b> Returns the current filename.</li><li><b>product</b> Returns the product name of the current audio</li></ul>

		output device.
	<b>input</b>	Returns the product name of the current waveform input device or <b>none</b> if no device is set.
	<b>output</b>	Returns the product name of the current waveform output device or <b>none</b> if no device is set.
<b>open items</b>		Initializes the device. The following <i>items</i> are optional:
	<b>alias</b> <i>device_alias</i>	Specifies an alternate name for the given device. If specified, it must also be used the alias for references.
	<b>buffer</b> <i>buffer_size</i>	Sets the size in seconds of the buffer used by the wave audio device. The default size of the buffer is set when the wave audio device is installed or configured. Typically, the buffer size is set to 4 seconds.
	<b>shareable</b>	Initializes the device element as shareable. Subsequent attempts to open it fail unless you specify <b>shareable</b> in both the original and subsequent <b>open</b> commands. MCI returns an error if it is already open and not shareable.
	<b>type</b> <i>device_type</i>	Specifies the compound device used to control a device element. MCI reserves waveaudio for the waveform audio device type. As an alternative to <b>type</b> , MCI can use the element filename extension entries to select the controlling device
<b>pause</b>		Pauses playing or recording.
<b>play items</b>		Starts playing audio. The following optional <i>items</i> modify <b>play</b> :
	<b>from</b> <i>position</i> <b>to</b> <i>position</i>	Specifies the positions to start and stop playing. If <b>from</b> is omitted, play starts at the current position; if <b>to</b> is omitted, play stops at the end of the file or waveform.
<b>record items</b>		Starts recording audio. All data recorded after a file is opened is discarded if the file is closed without saving it. The following optional <i>items</i> modify <b>record</b> :
	<b>insert</b>	Specifies that new data is added to the device element.
	<b>from</b> <i>position</i> <b>to</b> <i>position</i>	Specifies the positions to start and stop recording. If <b>from</b> is omitted, the device starts recording at the current position; if <b>to</b> is omitted, the device records until a <b>stop</b> or <b>pause</b> command is received.
	<b>overwrite</b>	Specifies that new data will replace data in the device element.
<b>resume</b>		Resumes playing or recording following a pause.
<b>save item</b>		Saves the MCI element in its current format. The following item modifies <b>save</b> :
	<i>filename</i>	Specifies the file and pathname used to save

data.

<b>seek item</b>	Moves to the specified location in the file. Playback or recording is stopped after the seek. One of the following items modify <b>seek</b> :
<b>to position</b>	Specifies the stop position.
<b>to start</b>	Specifies to seek to the first sample.
<b>to end</b>	Specifies to seek to the last sample.
<b>set items</b>	Sets the following control <i>items</i> :
<b>alignment integer</b>	Sets the alignment of data blocks. The file is saved in the new format.
<b>any input</b>	Use any input that supports the current format when recording. This is the default.
<b>any output</b>	Use any output that supports the current format when playing. This is the default.
<b>audio all off</b> <b>audio all on</b>	Enables or disables audio output.
<b>audio left off</b> <b>audio left on</b>	Enables or disables output to the left audio channel.
<b>audio right off</b> <b>audio right on</b>	Enables or disables output to the right audio channel.
<b>bitspersample</b> <i>bit_count</i>	Sets the number of bits per sample played or recorded. The file is saved in this format.
<b>bytespersec</b> <i>byte_rate</i>	Sets the average number of bytes per second played or recorded. The file is saved in this format.
<b>channels</b> <i>channel_count</i>	Sets the channel count for playing and recording. The file is saved in this format.
<b>format tag tag</b>	Sets the format type for playing and recording. The file is saved in this format.
<b>format tag pcm</b>	Sets the format type to PCM for playing and recording. The file is saved in this format.
<b>input integer</b>	Sets the audio channel used as the input.
<b>output integer</b>	Sets the audio channel used as the output.
<b>samplespersec</b> <i>integer</i>	Sets the sample rate for playing and recording. The file is saved in this format.
<b>time format bytes</b>	Sets the time format to bytes. All position information is specified as bytes following this command.
<b>time format milliseconds</b>	Sets the time format to milliseconds. All position information is specified as milliseconds following this command. You can abbreviate milliseconds as <b>ms</b> .
<b>time format samples</b>	Sets the time format to samples. All position information is specified as samples following this command.
<b>status item</b>	Obtains status information for the device. One of the following <i>items</i> modifies <b>status</b> :

<b>alignment</b>	Returns the block alignment of data in bytes.
<b>bitspersample</b>	Returns the bits per sample.
<b>bytespersec</b>	Returns the average number of bytes per second played or recorded.
<b>channels</b>	Returns the number of channels set (1 for mono, 2 for stereo).
<b>current track</b>	Returns the index of the current track.
<b>format tag</b>	Returns the format tag.
<b>input</b>	Returns the currently set input. If no input is set, the error returned indicates that any device can be used.
<b>length</b>	Returns the total length of the waveform.
<b>length track</b> <i>track_number</i>	Returns the length of the waveform track.
<b>level</b>	Returns the current audio sample value.
<b>media present</b>	Returns <b>true</b> .
<b>mode</b>	Returns <b>not ready, paused, playing, stopped, recording, or seeking</b> .
<b>number of tracks</b>	Returns the number of tracks (chapters).
<b>output</b>	Returns the currently set output. If no output is set, the error returned indicates that any device can be used.
<b>position</b>	Returns the current position.
<b>position track</b> <i>track_number</i>	Returns the position of the track specified by <i>track_number</i> .
<b>ready</b>	Returns <b>true</b> if the device is ready.
<b>samplespersec</b>	Returns the number of samples per second played or recorded.
<b>start position</b>	Returns the starting position of the waveform data.
<b>time format</b>	Returns the current time format.
<b>stop</b>	Stops playing or recording.

---