

ROCK RIDGE INTERCHANGE PROTOCOL VERSION 1

**AN ISO 9660:1988 COMPLIANT APPROACH
TO PROVIDING ADEQUATE CD-ROM SUPPORT
FOR POSIX FILE SYSTEM SEMANTICS**

**IEEE CDROM FILE SYSTEMS FORMAT
WORKING GROUP**

Revision 1.10

cdfsf@ymi.COM

DRAFT STANDARD

(This page should be replaced with a blank page for back of title)

(This page should be replaced with the Table of Contents)

(This page should be replaced with a blank page for back of TOC)

(This page should be replaced with the List of Tables)

(This page should be replaced with a blank page for back of Tables)

1. PREFACE

1.1 Purpose and Scope

Producers and users of POSIX compliant systems and software have faced a significant, yet artificial, barrier to their effectively using CD-ROM technology for software distribution and information publishing -- ISO 9660:1988 alone provides inadequate support for delivery of POSIX file system information. The Rock Ridge Group was formed to generate a proposed standard for utilizing the System Use Areas provided by the ISO 9660 standard to record complete POSIX file system semantics. This proposal utilizes the System Use Sharing Protocol for recording this information. The IEEE CDROM File Systems Format group is formalizing the Rock Ridge standard for ISO adoption, incorporating changes to support sparse files and coexistence with other CDROM formats.

1.2 Summary of Sections

Section 1	Contains this preface.
Section 2	Contains an overview of the Rock Ridge Interchange Protocol.
Section 3	Contains an overview of the notation used in this document.
Section 4	Contains the Rock Ridge Interchange Protocol.
Section 5	Contains the bibliography.

2. OVERVIEW

The Rock Ridge Interchange Protocol (RRIP) specifies an extension to the ISO 9660 format for CD-ROM which enables the recording of POSIX File System semantics. The RRIP utilizes the System Use Sharing Protocol (SUSP) to specify the definition of a set of System Use Fields for this purpose.

The RRIP specifies the definition of a set of System Use Fields for recording:

- uid, gid, and permissions
- file mode bits, file types, setuid, setgid, and sticky bit
- file links
- device nodes
- symbolic links
- POSIX file names
- reconstruction of deep directories
- time stamps

3. TERMINOLOGY AND NOTATION

It is assumed that the RRIP is being recorded within an ISO 9660:1988 compliant volume using the System Use Sharing Protocol (SUSP:1991A). Unless defined herein, or otherwise specified, terms shall be as defined in ISO 9660:1988 or SUSP:1991A.

The following notation is used in this document.

3.1 Decimal and Hexadecimal Notation

Numbers in decimal notation are represented by decimal digits, namely 0 to 9.

Numbers in hexadecimal notation are represented by hexadecimal digits, namely 0 to 9 and A to F, shown in parentheses. E.g. the hexadecimal number 7F will be written as (7F).

3.2 Binary powers of 2

Powers of two may be represented by abbreviation. The abbreviation K represents the value 1024. E.g. the decimal number 4096 may be written as 4K. Likewise, the abbreviation M represents 1048576.

3.3 File Naming Conventions

In all fields defined in ISO 9660:1988, the character set to be used shall be as specified in ISO 9660. The character set to be used in the System Use Fields defined herein shall depend upon whether the fields are recorded in a directory tree associated with a Primary Volume Descriptor or a Supplementary Volume Descriptor.

3.3.1 Primary Volume Descriptor File Naming Convention

Within a directory tree identified by a Primary Volume Descriptor of an ISO 9660 volume, the character set used in the System Use Fields defined for the RRIP shall be the ISO Standard 8859-1:1987, as specified in the X/Open Portability Guide Issue 3, XSI Supplementary Definitions. For general portability across most POSIX compatible systems, the 7-bit U.S. ASCII character set should be used.

The POSIX File Naming Convention states that the filename may be composed of any character except slash (2F) and the null byte (00). The special filename, dot (2E), refers to the directory specified by its predecessor. The special filename dot-dot (2E)(2E), refers to the parent directory of its predecessor directory.

For maximum portability across implementations conforming to the POSIX Standard, filenames should only contain the following characters:

(41) thru (5A)	'A' thru 'Z'
(61) thru (7A)	'a' thru 'z'

(30) thru (39)	'0' thru '9'
(2E)	period
(5F)	underscore
(2D)	hyphen

Upper and lower case letters shall retain their unique identities between conforming implementations.

3.3.2 Supplementary Volume Descriptor File Naming Convention

Within a directory tree identified by a Supplementary Volume Descriptor of an ISO 9660 volume, the character set used in the System Use Fields defined for the RRIP shall be the coded graphic character set(s) identified by the escape sequence(s) in the Supplementary Volume Descriptor (c-characters, section 7.4.2, ISO 9660:1988).

3.4 Reader Makes Right

Receiving systems which are capable of interpreting the System Use Fields defined herein, but which cannot handle the full extent of the file naming convention utilized by this specification may have to restrict themselves to the use of the ISO 9660 compliant file names stored in the File Identifier field of the ISO 9660 directory structure.

Alternatively, the developer of the receiving system may choose to provide file name conversion capability. Any such system must provide unique names for all unique files on the disc.

In general, whenever there is recorded a (potentially) system-specific identifier or numerical value, accomplishing any necessary modifications or mapping of these are the responsibility of the receiving system. This document provides for an Application Programming Interface (API) to support this function.

4. ROCK RIDGE INTERCHANGE PROTOCOL

The Rock Ridge Interchange Protocol (RRIP) utilizes System Use Areas provided by ISO 9660:1988. The System Use Area of the Directory Record is used to record the POSIX file system information. The System Use Sharing Protocol is used for recording information in each of these areas.

4.1 System Use Fields Provided by this Specification

The Rock Ridge Interchange Protocol defines the following fundamental System Use Fields:

"PX"	POSIX file attributes
"PN"	POSIX device modes
"SL"	Symbolic link
"NM"	Alternate name
"CL"	Child link
"PL"	Parent link
"RE"	Relocated directory
"TF"	Time stamp(s) for a file
"SF"	File data is in sparse file format

Additionally, this specification provides required identification information to be recorded in an "ER" (Extensions Reference) System Use Field for the purpose of identifying discs on which the Rock Ridge Interchange Protocol is implemented.

4.1.1 Description of the PX System Use Field

Recording of the "PX" System Use Field in the System Use Area of each ISO 9660 directory record shall be mandatory. No more than one "PX" System Use Field shall appear in (all) the System Use Area(s) for a single directory record.

If the file type in a directory record is of type directory, then the POSIX File Mode Field ([BP 4] in this section) and File Flags (ISO 9660 Format section 9.1.6) should both indicate such, with the exception of relocated directories, indicated by a "CL" field (section 3.1.5.1), for which the ISO file flags shall indicate a normal file, but the POSIX File Mode Field shall indicate a directory.

The format of the "PX" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "PX" type System Use Field. The bytes in this field shall be (50)(58) ("PX").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "PX" System Use Field. The number in this field shall be 36 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.

- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "PX" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 to BP 12 - POSIX File Mode" has the same meaning as the `st_mode` defined in the header `sys/stat.h` by the IEEE Std 1003.1-1988. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3. The valid mask values for this field are combinations of the following:

Octal Value	Mnemonic	Meaning
0000400	<code>S_IRUSR</code>	read permission (owner)
0000200	<code>S_IWUSR</code>	write permission (owner)
0000100	<code>S_IXUSR</code>	execute permission (owner)
0000040	<code>S_IRGRP</code>	read permission (group)
0000020	<code>S_IWGRP</code>	write permission (group)
0000010	<code>S_IXGRP</code>	execute permission (group)
0000004	<code>S_IROTH</code>	read permission (other)
0000002	<code>S_IWOTH</code>	write permission (other)
0000001	<code>S_IXOTH</code>	execute permission (other)
0004000	<code>S_ISUID</code>	set user ID on execution
0002000	<code>S_ISGID</code>	set group ID on execution
0002000	<code>S_ENFMT</code>	enforced file locking (shared w/ set group ID)
0001000	<code>S_ISVTX</code>	save swapped text even after use
0170000	<code>S_IFMT</code>	type of file
0140000	<code>S_IFSOCK</code>	socket
0120000	<code>S_IFLNK</code>	symbolic link
0100000	<code>S_IFREG</code>	regular
0060000	<code>S_IFBLK</code>	block special
0020000	<code>S_IFCHR</code>	character special
0040000	<code>S_IFDIR</code>	directory
0010000	<code>S_IFIFO</code>	pipe or FIFO

- [5] "BP 13 to BP 20 - POSIX File Links" has the same meaning as the `st_nlink` defined in the header `sys/stat.h` by the IEEE Std 1003.1-1988. If the file type described by the directory record is a directory, the value in this field shall equal the number of directories in the directory described by this directory record (i.e. any directory record which has the "directory" bit set, including the "." and ".." records). Otherwise, it shall be 1. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.
- [6] "BP 21 to BP 28 - POSIX File User ID" has the same meaning as the `st_uid` defined in the header `sys/stat.h` by the IEEE Std 1003.1-1988. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.

- [7] "BP 29 to BP 36 - POSIX File Group ID" has the same meaning as the `st_gid` defined in the header `sys/stat.h` by the IEEE Std 1003.1-1988. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.

TABLE 1. PX System Use Field - Version 1

'P' (BP1)	'X' (BP2)	LENGTH (BP3)	1 (BP4)	FILE MODE (BP5 to BP12)	LINKS (BP13 to BP20)
--------------	--------------	-----------------	------------	----------------------------	-------------------------

USER ID (BP21 to BP28)	GROUP ID (BP29 to BP36)
---------------------------	----------------------------

4.1.2 Description of the PN System Use Field

This field is mandatory if the file type recorded in the "PX" File Mode field for the given directory record indicates a character or block device. This field, if present, should be ignored for all other file types. No more than one "PN" System Use Field shall appear in (all) the System Use Area(s) for a single directory record.

If the receiving system records device numbers as 32-bit numbers, only the "Dev_t Low" field shall be used. If the receiving system records device numbers as 64-bit numbers, the "Dev_t High" and "Dev_t Low" fields shall be concatenated to make one 64-bit number.

The format of the "PN" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "PN" type System Use Field. The bytes in this field shall be (50)(4E) ("PN").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "PN" System Use Field. The number in this field shall be 20 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "PN" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.

- [4] "BP 5 to BP 12 - Dev_t High" shall contain as a 32-bit number the high order 32 bits of the device number. If the receiving system records device numbers as 32-bit numbers this field shall be zero and ignored. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.
- [5] "BP 13 to BP 20 - Dev_t Low" shall contain as a 32-bit number the low order 32-bits of the device number. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.

TABLE 2. PN System Use Field - Version 1

'P' (BP1)	'N' (BP2)	20 (BP3)	1 (BP4)	DEV_T HIGH (BP5 to BP12)	DEV_T LOW (BP13 to BP20)
--------------	--------------	-------------	------------	-----------------------------	-----------------------------

4.1.3 Description of the SL System Use Field

The purpose of the "SL" System Use Field is to store the content of a symbolic link. This System Use Field is mandatory if the file type recorded in the "PX" File Mode field for the given directory record indicates a symbolic link. For other file types, this System Use Field should be ignored. If the receiving system does not support symbolic links the system should invoke "Reader-Makes-Right".

If the file type recorded in the "PX" File Mode field for the given directory record indicates a symbolic link, the directory record shall point to a file, the contents of which are not specified by this document.

If more than one "SL" System Use Field is recorded in the System Use Area(s) for a single directory record, the Component Area (see section 4.1.3.1 below) of each should be concatenated together, in the order in which they were recorded, until a CONTINUE flag with value zero is encountered (see [4] below), to obtain the entire set of Component Records for the symbolic link.

The method of recording is system independent. Under reader makes right, the receiving system is responsible for supplying appropriate values and/or notations for the component delimiter and special cases provided for below.

The format of the "SL" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "SL" type System Use Field. The bytes in this field shall be (53)(4C) ("SL").
- [2] "BP 3 - Length (LEN_SL)" shall specify as an 8-bit number the length in bytes of the "SL" System Use Field. The number in this field shall be 5 plus the length of the Component Area recorded in this "SL" field. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "SL" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 - Flags" shall contain bit field flags numbered 0 to 7 starting with the least significant bit as follows:

Position	Name	Interpretation if set to 1
0	CONTINUE	This Symbolic Link continues in next "SL"
all others	RESERVED	value must be 0

- [5] "BP 6 to LEN_SL - Component Records" shall contain Component Records (described below).

TABLE 3. SL System Use Field - Version 1

'S' (BP1)	'L' (BP2)	LENGTH (BP3)	1 (BP4)	FLAGS (BP5)	COMPONENT RECORDS (BP6 to LEN_SL)
--------------	--------------	-----------------	------------	----------------	--------------------------------------

4.1.3.1 Description of the SL System Use Field Component Record

Within a "SL" System Use Field, each component of the pathname shall be recorded as one or more component records. A component does not contain the component separator (/ in POSIX). Recording a single component of a symbolic link may require multiple Component Records. If the component is greater than 255 bytes or will not fit into the current System Use Area or Continuation Area more than one Component Record will be recorded for the component. Multiple Component Records, specifying one or more separate components of the symbolic link may be recorded in the Component Area of a single "SL" field.

If a single Component Record is used to record a single component of a symbolic link, the CONTINUE flag must be set to zero. If multiple Component Records are used to record a single component of a symbolic link, the CONTINUE flag must be set to one in each Component Record except the last and zero in the last Component Record recording the given component.

Component Records shall be recorded contiguously within each Component Area, starting in the first byte of each Component Area. The last Component Record in the Component Area of an "SL" field may be continued in the Component Area of the next recorded "SL" field.

Each Component Record shall have the following format:

- [A] "BP 1 - Component Flags" shall contain bit field flags numbered 0 to 7 starting with the least significant bit as follows:

Position	Name	Interpretation if set to 1
0	CONTINUE	Component recorded in this "SL" continues in next "SL" Component Record
1	CURRENT	Component refers to the current directory (. in POSIX)
2	PARENT	Component refers to the parent of the current directory (.. in POSIX)
3	ROOT	Component refers to the root of the current directory tree for this process (/ in POSIX)
4	VOLROOT	Component refers to the directory the current CD-ROM volume is mounted on
5	HOST	The local host name should be inserted as the value of the current component
all others	RESERVED	value must be 0

Bits 1 - 7 are mutually exclusive.

- [B] "BP 2 - Component Length (LEN_CP)" shall specify as an 8-bit number the length in bytes of the (portion of) the component recorded in the current Component Record. This length shall not include the Component Record Flags byte or Length byte. If any of the 2^1 thru 2^5 bits are set, the value of this field shall be zero. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [C] "BP 3 to 2 + LEN_CP - Component" shall contain (the portion of) the component recorded in the current Component Record. The content of this field shall be recorded according to section 3.2 above.

TABLE 4. SL System Use Field - Component Record

COMP_FLAGS (BP1)	COMP_LEN (BP2)	COMPONENT (BP3 to 2+LEN_CP)
---------------------	-------------------	--------------------------------

4.1.4 Description of the NM System Use Field

The purpose of the "NM" System Use Field is to store the content of an Alternate Name to support POSIX-style or other names. This System Use Field is optional. If no "NM" field(s) are recorded for a specific directory record, the ISO 9660 file identifier shall be used.

If more than one "NM" System Use Field appears in (all) the System Use Area(s) for a single directory record, the contents ([5] below) of each should be concatenated together, in the order in which they were recorded, until a CONTINUE flag with value zero is encountered (see [4] below), to obtain the entire Alternate Name.

"NM" System Use Fields recorded for the ISO 9660 directory records with names (00) and (01), used to designate the current and parent directories, respectively, should be ignored. Instead, the receiving system should convert these names to the appropriate receiving system-dependent designations for the current and parent directories.

No sorting of the directory records by Alternate Names is specified by the RRIP, nor can one necessarily be imposed by originating systems or assumed by receiving systems. The ISO 9660 specifies a sorting order based upon the ISO 9660 file identifier (see ISO 9660:1988, section 9.3).

The format of the "NM" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "NM" type System Use Field. The bytes in this field shall be (4E)(4D) ("NM").
- [2] "BP 3 - Length (LEN_NAM)" shall specify as an 8-bit number the length in bytes of the "NM" System Use Field. The number in this field shall be 5 plus the length (of the portion) of the Alternate Name recorded in this "NM" field. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "NM" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.

- [4] "BP 5 - Flags" shall contain bit field flags numbered 0 to 7 starting with the least significant bit as follows:

Position	Name	Interpretation if set to 1
0	CONTINUE	Alternate Name continues in next "NM" field
1	CURRENT	Alternate Name refers to the current directory (. in POSIX)
2	PARENT	Alternate Name refers to the parent of the current directory (.. in POSIX)
3	RESERVED	value must be 0
4	RESERVED	value must be 0
5	HOST	of the current CD-ROM volume The local host name should be inserted as the value of the Alternate Name
all others	RESERVED	value must be 0

Bits 1 - 7 are mutually exclusive.

- [5] "BP 6 to LEN_NAM - Alternate Name" shall contain (a portion of) the content of the Alternate Name. The content of this field shall be recorded according to section 3.2 above.

TABLE 5. NM System Use Field - Version 1

'N' (BP1)	'M' (BP2)	LENGTH (BP3)	1 (BP4)	FLAGS (BP5)	ALTERNATE NAME (BP6 to LEN_NAM)
--------------	--------------	-----------------	------------	----------------	------------------------------------

4.1.5 System Use Fields for Handling Deep Directory Trees

The ISO 9660:1988 mandates directory depths of no more than eight levels. Deeper directories must be reorganized to be recorded under the ISO 9660. The RRIP includes definitions of three System Use Fields to support logical reconstruction of deep directory trees while retaining complete ISO 9660 compliance.

For each specific directory, either all three of the following fields must be appropriately recorded, or none shall be recorded.

Table 9 and Table 10 at the end of this section have graphical examples of Deep Directory Trees.

4.1.5.1 Description of the CL System Use Field

The purpose of the "CL" System Use Field is to record the new location of a directory which has been relocated. The field contains the Logical Block number of the Logical Sector in which the first directory record of the moved Directory is stored.

The "CL" System Use Field is optional. If recorded, a "CL" System Use Field shall be recorded in the System Use Area of a ISO 9660 directory record which describes a file which has the same name as, and occupies the original position in the ISO 9660 directory tree of, the moved Directory. No more than one "CL" System Use Field shall appear in (all) the System Use Area(s) for a single directory record.

Except for the ISO 9660 name, the Alternate Name (recorded in an "NM" System Use Field), and the new location of the Directory, all other information stored in the directory for this file should be ignored. The contents of this file are not specified by this document. All attributes of the moved Directory shall be recorded in the first directory record ("dot" entry) of the moved Directory in its new location.

The format of the "CL" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "CL" type System Use Field. The bytes in this field shall be (43)(4C) ("CL").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "CL" System Use Field. The number in this field shall be 12 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "CL" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 to BP 12 - Location of Child Directory" shall specify as a 32-bit number the Logical Block Number of the first Logical Block allocated to the moved directory. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.

TABLE 6. CL System Use Field - Version 1

'C' (BP1)	'L' (BP2)	12 (BP3)	1 (BP4)	LOC of CHILD DIRECTORY (BP5 to BP12)
--------------	--------------	-------------	------------	---

4.1.5.2 Description of the PL System Use Field

The purpose of the "PL" System Use Field is to record the location of the original parent Directory of a Directory which has been relocated. The field contains the Logical Block number of the Logical Sector in which the first directory record of the original parent Directory of said moved Directory is stored.

For each moved Directory which is recorded using a "CL" System Use Field, a corresponding "PL" System Use Field is required. The "PL" System Use Field shall be recorded in the System Use Area of the second directory record ("dot-dot" entry) of the moved Directory. No more than one "PL" System Use Field shall appear in (all) the System Use Area(s) for a single directory record.

The format of the "PL" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "PL" type System Use Field. The bytes in this field shall be (50)(4C) ("PL").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "PL" System Use Field. The number in this field shall be 12 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "PL" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 to BP 12 - Location of Parent Directory" shall specify as a 32-bit number the Logical Block Number of the first Logical Block allocated to the original parent directory of the moved directory. This field shall be recorded according to ISO 9660:1988 Format section 7.3.3.

TABLE 7. PL System Use Field - Version 1

'P' (BP1)	'L' (BP2)	12 (BP3)	1 (BP4)	LOC of PARENT DIRECTORY (BP5 to BP12)
--------------	--------------	-------------	------------	--

4.1.5.3 Description of the RE System Use Field

The purpose of the "RE" System Use Field is to indicate to a receiving system which can understand the System Use Fields "CL" and "PL" that the directory record in which this "RE" System Use Field is recorded has been relocated from another position in the original directory tree.

An "RE" System Use Field shall not be recorded unless a corresponding "CL" System Use Field is recorded. If recorded, a "RE" System Use Field shall be recorded in the System Use Area of the directory record which describes the moved Directory in the new parent directory of the moved Directory.

The format of the "RE" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "RE" type System Use Field. The bytes in this field shall be (52)(45) ("RE").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "RE" System Use Field. The number in this field shall be 4 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "RE" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.

TABLE 8. RE System Use Field - Version 1

'R' (BP1)	'E' (BP2)	4 (BP3)	1 (BP4)
--------------	--------------	------------	------------

TABLE 9. Deep Directory Relocation

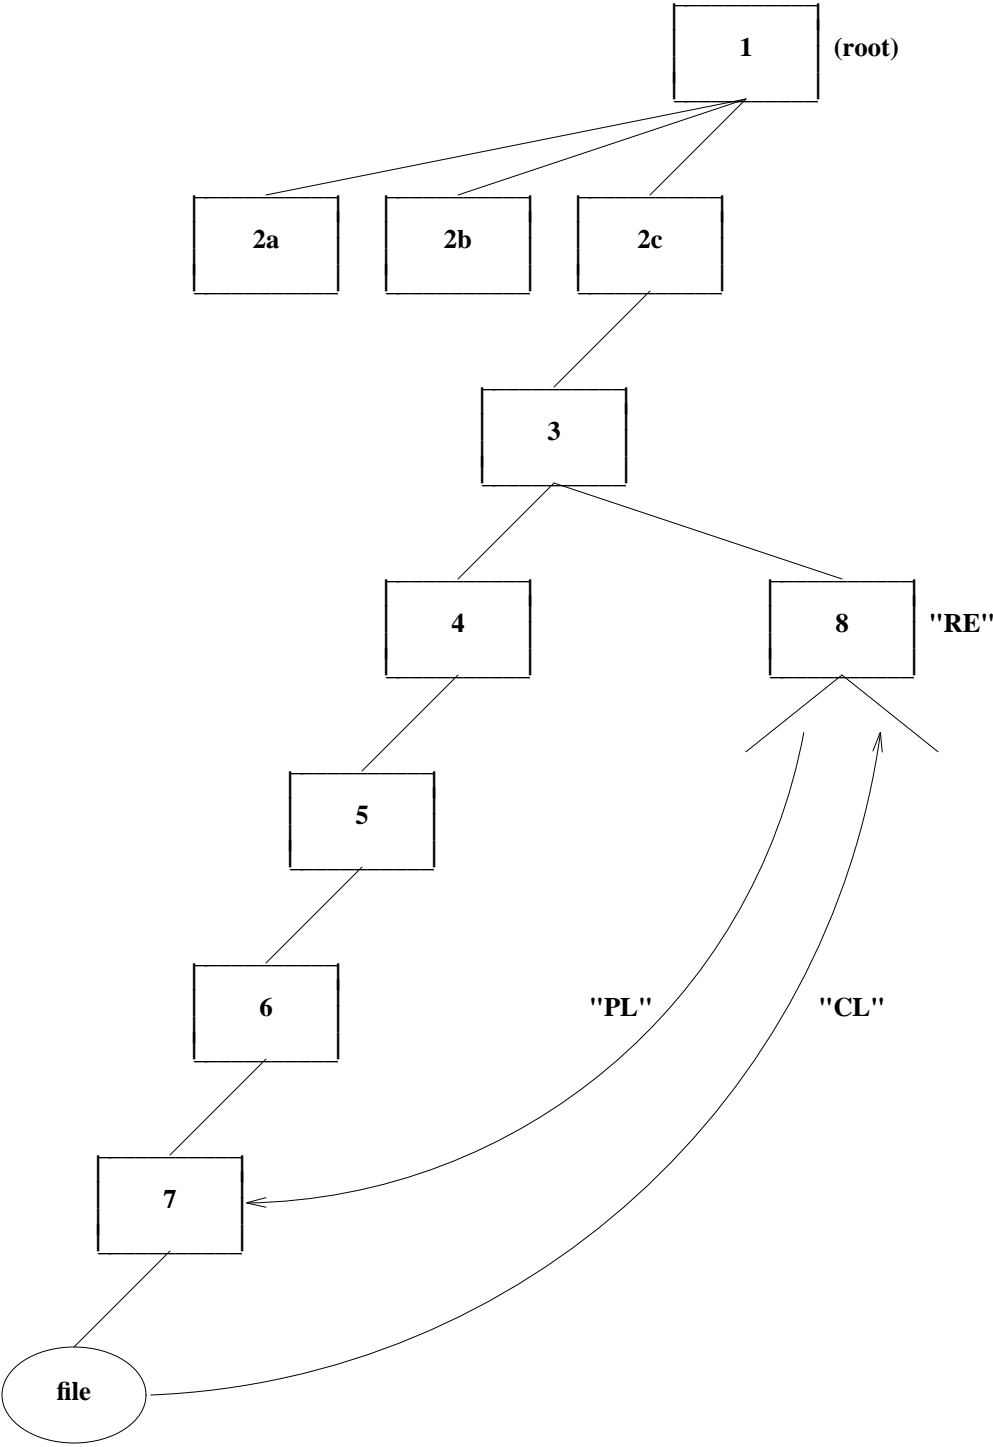
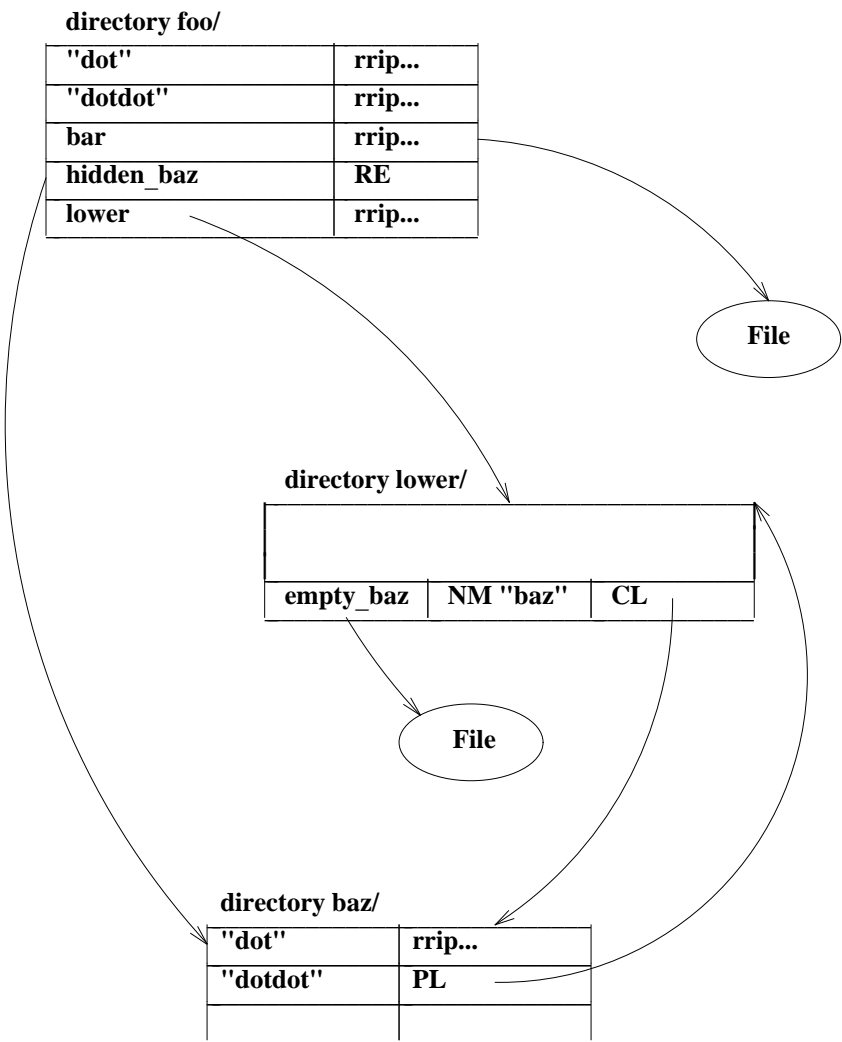


TABLE 10. Detailed Deep Directory Relocation



4.1.6 Description of the TF System Use Field

The purpose of the "TF" System Use Field is to allow the recording of a complete set of time stamps related to a file. This System Use Field shall be optional. If this field does not exist, the POSIX `st_atime`, `st_ctime` and `st_mtime` should have the same value as Recording Date and Time Field of the ISO 9660:1988 directory record. If both the "TF" System Use Field and the XAR are present, the time attributes stored in these two areas must be consistent. If only the XAR is present, the `st_atime` should have the same value as the Recording Date and Time Field of the ISO 9660 directory record.

Multiple "TF" fields may be recorded, using any combination of time stamps and time formats, but each individual time stamp may be recorded only once per directory record.

The format of the "TF" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "TF" type System Use Field. The bytes in this field shall be (54)(46) ("TF").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "TF" System Use Field. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "TF" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 - Flags" shall contain bit field flags numbered 0 to 7 starting with the least significant bit as follows:

Position	Name	Interpretation if set to 1
0	CREATION	Creation time recorded
1	MODIFY	Modification time recorded
2	ACCESS	Last Access time recorded
3	ATTRIBUTES	Last Attribute Change time recorded
4	BACKUP	Last Backup time recorded
5	EXPIRATION	Expiration time recorded
6	EFFECTIVE	Effective time recorded
7	LONG_FORM	ISO 9660 17-byte time format used

If the `LONG_FORM` bit is set to one, all time stamps in this "TF" System Use Field shall be recorded using the format specified in Section 8.4.26.1 of ISO 9660:1988. If the `LONG_FORM` bit is set to zero, all time stamps in this "TF" System Use Field shall be recorded using the format specified in Section 9.1.5 of ISO 9660:1988.

- [4+N] "BP 6+(X*(N-1)) to 5+(X*N) Time Stamp" shall contain the Nth time stamp indicated in [4] as being recorded, starting with the 0th bit and working sequentially through the list of recordable time stamps. The `LONG_FORM` bit does not indicate the presence or absence of any time stamp. The value of X in the expression above shall be 17 if the `LONG_FORM` bit is set to 1, and 7 otherwise.

The recorded time for each of the time stamps recorded in this field shall be local time. if recorded, CREATION, Creation Time, has the same meaning as in ISO 9660:1988 Format section 9.5.4.

If recorded, MODIFY, File Modification Date and Time, has the same meaning as in ISO 9660:1988 Format section 9.5.5. This field shall be used by the `st_mtime` for POSIX conformance.

If recorded, ACCESS, File Last Access Date and Time, shall specify the date and time of the day at which the information in the file was last accessed. This field shall be used by the `st_atime` for POSIX conformance.

If recorded, ATTRIBUTES, Last Attribute Change Time, shall be used by the `st_ctime` field for POSIX conformance.

If recorded, BACKUP, Last Backup Time, shall provide a time stamp for the most recent backup of this file. The utilization of this information is not restricted by this specification.

If recorded, EXPIRE, File Expiration Date and Time, has the same meaning as in ISO 9660:1988 Format section 9.5.6.

If recorded, EFFECT File Effective Date and Time" has the same meaning as in ISO 9660:1988 Format section 9.5.7.

TABLE 11. TF System Use Field - Version 1

'T' (BP1)	'F' (BP2)	LENGTH (BP3)	1 (BP4)	FLAGS (BP5)	TIME STAMPS (BP6 to LENGTH)
--------------	--------------	-----------------	------------	----------------	--------------------------------

4.1.7 Description of the SF System Use Field

The purpose of the "SF" System Use Field is to indicate that the file identified by the current directory record is stored as a "sparse" file, and to provide additional information which is necessary to retrieve the file contents. This System Use Field (and sparse file encoding) is optional. No more than one "SF" System Use Field shall appear in (all) the System Use Area(s) for a single directory record.

The "SF" field is specifically designed to provide support for the encoding and delivery of Unix-style sparse files in a platform-independent manner. The sparse file encoding allocates physical blocks only if the block actually contains non-zero data.

The format of the "SF" System Use Field is as follows:

- [1] "BP 1 to BP 2 - Signature Word" shall indicate that the System Use Field is a "SF" type System Use Field. The bytes in this field shall be (53)(46) ("SF").
- [2] "BP 3 - Length" shall specify as an 8-bit number the length in bytes of the "SF" System Use Field. The number in this field shall be 12 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [3] "BP 4 - System Use Field Version" shall specify as an 8-bit number an identification of the version of the "SF" System Use Field. The number in this field shall be 1 for this version. This field shall be recorded according to ISO 9660:1988 Format section 7.1.1.
- [4] "BP 5 to BP 12 - Virtual File Size" shall specify as a 32-bit number the apparent size of the recorded data file, i.e. the size of the file as reported on the originating (Unix) file system, or the offset from the beginning of the file to the last application-addressable byte in the file. This field shall be recorded according to the ISO 9660:1988 Format section 7.3.3.

TABLE 12. SF System Use Field - Version 1

'S' (BP1)	'F' (BP2)	LENGTH (BP3)	1 (BP4)	VIRTUAL FILE SIZE (BP5 to BP12)
--------------	--------------	-----------------	------------	------------------------------------

4.1.7.1 Encoding and Recording of Sparse Files

Sparse Files are encoded within the File Section as specified in section 6.4.4.2 of the ISO 9660:1988. The directory record Data Length as specified in section 9.1.4 of the ISO 9660:1988 shall specify the length of the file section, including the SF Header Block, all Index Blocks, and the sparse file data.

The initial (number 0) 2K byte block of the File Section shall be an SF Header Block.

The format of the SF Header Block is identical to the format of the SF System Use Field. All unused bytes in the SF header block shall be set to binary zero (0).

The second (number 1) 2K byte block of the File Section shall be the first Index Block. Each Index Block of an the encoded file shall hold 256 table entries.

Each Table Entry of an Index Block shall be eight bytes, recording a 32 bit number as specified in section 7.3.3 of the ISO 9660:1988. The value of each Table Entry is interpreted as a set of bit fields numbered 0 to 31 starting with the least significant bit as follows:

Bits	Name	Interpretation if set to 1
0-23	BLOCK	24 bit logical block number. If the TABLE bit is set, this number is the 2K byte block offset from the first block of the File Section to

		the next block of table entries.
		If the TABLE bit is not set, this number is the 256 byte block offset from the first block of the File Section to the first block of data referenced by this Table Entry.
		If the ZERO bit is set, this block number must be zero.
24-29	RESERVED	value must be 0.
30	TABLE	This Table Entry's BLOCK pointer references a 2K block of 256 Table Entries.
31	ZERO	This Table Entry specifies a file extent of data containing only zeros.

The TABLE and ZERO bits are mutually exclusive.

Each Table Entry of the first Index Block corresponds to the high-order byte of a 32 bit, linearly addressed file position. Any Table Entry in the first Index Block which has the TABLE bit set, has a block pointer which refers to a logical 2K block (relative to the File Section) containing a second tier Index Block. Any Table Entry with the ZERO bit set represents a logical file region of 16M bytes of zeros. If the TABLE bit and the ZERO bit are both zero, then the BLOCK field of the Table Entry is a pointer to the first logical block of a contiguous 16M byte region of data.

Each Table Entry of a second tier Index Block corresponds to the second most significant byte of a 32 bit, linearly addressed file position. Any Table Entry in a second tier Index Block which has the TABLE bit set, has a block pointer which refers to a logical 2K block (relative to the File Section) containing a third tier Index Block. Any Table Entry with the ZERO bit set represents a logical file region of 64K bytes of zeros. If the TABLE bit and the ZERO bit are both zero, then the BLOCK field of the Table Entry is a pointer to the first logical block of a contiguous 64K byte region of data.

Each Table Entry of a third tier Index Block corresponds to the third most significant byte of a 32 bit, linearly addressed file position. The block pointer refers to a logical 256 byte block (relative to the File Section) which contains 256 bytes of actual data. A Table Entry with the ZERO bit set represents a logical file region of 256 bytes of zeros. No Table Entry in a third tier Index Block may have the TABLE bit set. If the TABLE bit and the ZERO bit are both zero, then the BLOCK field of the Table Entry is a pointer to the first logical block of a contiguous 256 byte region of data.

The positions of these 256 byte data blocks shall be numbered with data block number 0 being coincident with the first 256 bytes of block 0 of the encoded file. Thus the first 2K byte block of the file, which actually holds the high-order byte addressing table, would consume data block positions 0 to 3, and data block 495 ($= 123 \times 4 + 3$) would be bytes 768 to 1023 located in the 123rd physical 2K byte block of the encoded file. Though the data may be efficiently recorded in sequentially numbered blocks, ordered according to increasing address of the recorded data, this is not required.

The second and third tier index blocks (if needed) may be recorded in sequentially numbered blocks, after the first index block, but before the first data block, although this is not required.

4.2 Required Recording and Consistency

The "PX" System Use Fields shall be recorded in every directory record. All recorded instances of the "PX" and "TF" System Use Fields in directory records referring to a single directory must be consistent.

4.3 Specification of the ER System Use Field Values for RRIP

The Extension Version number for the version of the RRIP defined herein shall be 1. The content of the Extension Identifier field shall be "RRIP_1991A". The Identifier Length shall be 10.

The recommended content of the Extension Descriptor shall be "THE ROCK RIDGE INTERCHANGE PROTOCOL PROVIDES SUPPORT FOR POSIX FILE SYSTEM SEMANTICS." The corresponding Description Length is 84.

The recommended content of the Extension Source shall be "PLEASE CONTACT DISC PUBLISHER FOR SPECIFICATION SOURCE. SEE PUBLISHER IDENTIFIER IN PRIMARY VOLUME DESCRIPTOR FOR CONTACT INFORMATION." The corresponding Source Length is 135.

5. BIBLIOGRAPHY

- ISO 9660:1988 - Information Processing - Volume and file structure of CD-ROM for information interchange
- IEEE Standard Portable Operating System Interface for Computer Environments, IEEE std 1003.1-1988 (a.k.a POSIX Standard), New York, New York:IEEE
- System Use Sharing Protocol, IEEE CDFS Rev 1.10

CONTENTS

1. PREFACE	1
1.1 Purpose and Scope	1
1.2 Summary of Sections	1
2. OVERVIEW	3
3. TERMINOLOGY AND NOTATION	5
3.1 Decimal and Hexadecimal Notation	5
3.2 Binary powers of 2	5
3.3 File Naming Conventions	5
3.3.1 Primary Volume Descriptor File Naming Convention	5
3.3.2 Supplementary Volume Descriptor File Naming Convention	6
3.4 Reader Makes Right	6
4. ROCK RIDGE INTERCHANGE PROTOCOL	7
4.1 System Use Fields Provided by this Specification	7
4.1.1 Description of the PX System Use Field	7
4.1.2 Description of the PN System Use Field	9
4.1.3 Description of the SL System Use Field	10
4.1.3.1 Description of the SL System Use Field Component Record	11
4.1.4 Description of the NM System Use Field	13
4.1.5 System Use Fields for Handling Deep Directory Trees	14
4.1.5.1 Description of the CL System Use Field	15
4.1.5.2 Description of the PL System Use Field	16
4.1.5.3 Description of the RE System Use Field	17
4.1.6 Description of the TF System Use Field	20
4.1.7 Description of the SF System Use Field	21
4.1.7.1 Encoding and Recording of Sparse Files	22
4.2 Required Recording and Consistency	23
4.3 Specification of the ER System Use Field Values for RRIP	24
5. BIBLIOGRAPHY	26

LIST OF TABLES

TABLE 1. PX System Use Field - Version 1	9
TABLE 2. PN System Use Field - Version 1	10
TABLE 3. SL System Use Field - Version 1	11
TABLE 4. SL System Use Field - Component Record	12
TABLE 5. NM System Use Field - Version 1	14
TABLE 6. CL System Use Field - Version 1	15
TABLE 7. PL System Use Field - Version 1	16
TABLE 8. RE System Use Field - Version 1	17
TABLE 9. Deep Directory Relocation	18
TABLE 10. Detailed Deep Directory Relocation	19
TABLE 11. TF System Use Field - Version 1	21
TABLE 12. SF System Use Field - Version 1	22