

19-June-1975

Revised 15-June-1977

Revised 15-April-1981

Copyright (C) 1975, 1977, 1981
Digital Equipment Corporation, Maynard, Mass.

The material included in this functional specification, including but not limited to instruction times and operating speeds, is for information purposes only. All such material is subject to change without notice. Consequently Digital Equipment Corporation makes no claim and shall not be liable for its accuracy.

This software is furnished under a license for use only on a single computer system and may be copied only with the inclusion of the above copyright notice. This software, or any other copies thereof, may not be provided or otherwise made available to any other person except for use on such system and to one who agrees to these license terms. Title to and ownership of the software shall at all times remain in Digital Equipment Corporation.

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation.

Digital Equipment Corporation assumes no responsibility for the use or reliability of its software on equipment which is not supplied by Digital Equipment Corporation.

1.0 Scope

This document is a specification of the on-media structure that is used by Files-11. Files-11 is a general purpose file structure which is intended to be the standard file structure for all medium to large PDP-11 systems. Small systems such as RT-11 have been specifically excluded because the complexity of Files-11 would impose too great a burden on their simplicity and small size.

This document describes structure level 1 of Files-11, also referred to as ODS-1 (on-disk structure version 1). This has been implemented on the RSX-11 family, (RSX-11M, RSX-11M-PLUS, IAS, and RSX-11D) and on VMS. This document describes the final level of functionality for ODS-1. Structure level 2 (ODS-2) has been implemented on VMS and is the basis for all new disk structure enhancements.

1.1 Summary of revisions made to this specification

1. Expanded File Characteristics to include most ODS-2 options.
2. Corrected H.FPRO to H.DFPR.
3. Added new fields to home block for date and count of home block modifications.
4. Added Single Directory Support description and home block field.
5. Added field in home block for pack serial number (H.PKSR).
6. Added description of modified storage control block format to support large disks.
7. Restricted maximum number of blocks supported on a volume to 1,044,480.
8. Restricted ODS-1 to one block "clusters".
9. Restricted ODS-1 to single volume structures.
10. Clarified and expanded references to operating system support and relationship to ODS-2.
11. Removed RMS-11 definitions, to be provided in separate specification common to ODS-1 and ODS-2.

2.0 Medium

Files-11 is a structure which is imposed on a medium. That medium must have certain properties, which are described in the following section. Generally speaking, block addressable storage devices such as disks and Dectape are suitable for Files-11; hence Files-11 structured media are generically referred to as disks.

2.1 Volume

The basic medium that carries a Files-11 structure is referred to as a volume. A volume (also often referred to as a unit) is defined as an ordered set of logical blocks. A logical block is an array of 512 8-bit bytes. The logical blocks in a volume are consecutively numbered from 0 to n-1, where the volume contains n logical blocks. The number assigned to a logical block is called its logical block number, or LBN. Files-11 is theoretically capable of describing volumes up to 232 blocks in size. In practice, a volume should be at least 100 blocks in size to be useful; current implementations of Files-11 will handle volumes up to 224 blocks.

The logical blocks of a volume must be randomly addressable. The volume must also allow transfers of any length up to 65k bytes, in multiples of four bytes. When a transfer is longer than 512 bytes, consecutively numbered logical blocks are transferred until the byte count is satisfied. In other words, the volume can be viewed as a partitioned array of bytes. It must allow reads and writes of arrays of any length less than 65k bytes, provided that they start on a logical block boundary and that the length is a multiple of four bytes. When only part of a block is written, the contents of the remainder of that logical block will be undefined.

2.2 Volume Sets

This section is of historical interest only. ODS-1 does not and will not support volume sets. A volume set is a collection of related units that are normally treated as one logical device in the usual operating system concept. Each unit contains its own Files-11 structure; however, files on the various units in a volume set may be referenced with a relative volume number, which uniquely determines which unit in the set the file is located on. Other sections in this specification will make occasional reference to volume sets and relative volume numbers where hooks for their implementation exist. Since volume sets have not been implemented as yet, however, no complete specification is provided here.

3.0 Files

Any data in a volume or volume set that is of any interest (i.e., all blocks not available for allocation) is contained in a file. A file is an ordered set of virtual blocks, where a virtual block is an array of 512 8 bit bytes. The virtual blocks of a file are consecutively numbered from 1 to n, where n blocks have been allocated to the file. The number assigned to a virtual block is called (obviously) its virtual block number, or VBN. Each virtual block is mapped to a unique logical block in the volume set by Files-11. Virtual blocks may be processed in the same manner as logical blocks. Any array of bytes less than 65k in length may be read or written, provided that the transfer starts on a virtual block boundary and that its length is a multiple of four.

3.1 File ID

Each file in a volume set is uniquely identified by a File ID. A File ID is a binary value consisting of 48 bits (3 PDP-11 words). It is supplied by the file system when the file is created, and must be supplied by the user whenever he wishes to reference a particular file.

The three words of the File ID are used as follows:

Word 1 File Number

Locates the file within a particular unit of the volume set. File numbers must lie in the range 1 through 65535. The set of file numbers on a unit is moderately (but not totally) dense; at any instant in time a file number uniquely identifies one file within that unit.

Word 2 File Sequence Number

Identifies the current use of an individual file number on a unit. File numbers are re-used; when a file is deleted its file number becomes available for future use for some other file. Each time a file number is re-used, a different file sequence number is assigned to distinguish the uses of that file number. The file sequence number is essential since it is perfectly legal for users to remember and attempt to use a File ID long after that file has been deleted.

Word 3 Relative Volume Number

Identifies which unit of a volume set the file is located on. Volume sets are at present not implemented; the only legal value for the

relative volume number in any context is zero.

3.2 File Header

Each file on a Files-11 volume is described by a file header. The file header is a block that contains all the information necessary to access the file. It is not part of the file; rather, it is contained in the volume's index file. (The index file is described in section 5.1). The header block is organized into four areas, of which the first three are variable in size.

3.2.1 Header Area

The information in the header area permits the file system to verify that this block is in fact a file header and, in particular, is the header being sought by the user. It contains the file number and file sequence number of the file, as well as its ownership and protection codes. This area also contains offsets to the other areas of the file header, thus defining their size. Finally, the header area contains a user attribute area, which may be used by the user to store a limited amount of data describing the file.

3.2.2 Ident Area

The ident area of a file header contains identification and accounting data about the file. Stored here are the primary name of the file, its creation date and time, revision count, date, and time, and expiration date.

3.2.3 Map Area

The map area describes the mapping of virtual blocks of the file to the logical blocks of the volume. The mapping data consists of a list of retrieval pointers. Each retrieval pointer describes one logically contiguous segment of the file. The map area also contains the linkage to the next extension header of the file, if such exists.

3.2.4 End Checksum

The last two bytes of the file header contain a 16 bit additive checksum of the remaining 255 words of the file header. The checksum is used to help verify that the block is in fact a file header.

3.3 Extension Headers

Since the file header is of fixed size, it is inevitable that for some files the mapping information will not fit in the allocated space. A file with a large amount of mapping data is therefore represented with a chain of file headers. Each header maps a consecutive set of virtual blocks; the extension linkage in the map area links the headers together in order of ascending virtual block numbers.

Multiple headers are also needed for files that span units in a volume set. A header may only map logical blocks located on its unit; therefore a multi-volume file is represented by headers on all units that contain portions of that file.

3.4 File Header - Detailed Description

This section describes in detail the items contained in the file header. Each item is identified by a symbol which represents the offset address of that item within its area in the file header. Any item may be located in the file header by locating the area to which it belongs and then adding the value of its offset address. Users who concern themselves with the contents of file headers are strongly urged to use the offset symbols. The symbols may be defined in assembly language programs by calling and invoking the macro FHDOPS, which may be found in the macro library of any system that supports Files-11. Alternatively, one may find the macro in the file F11MAC.MAC, which may be obtained from the author.

3.4.1 Header Area Description

The header area of the file header always starts at byte 0. It contains the basic information needed for checking the validity of accesses to the file.

3.4.1.1 H.IDOF 1 Byte Ident Area Offset

This byte contains the number of 16 bit words between the start of the file header and the start of the ident area. It defines the location of the ident area and the size of the header area.

3.4.1.2 H.MPOP 1 Byte Map Area Offset

This byte contains the number of 16 bit words between the start of the file header and the start of the map area. It defines the location of the map area and, together with H.IDOF, the size of the ident area.

- 3.4.1.3 H.FNUM 2 Bytes File Number
This word contains the file number of the file.
- 3.4.1.4 H.FSEQ 2 Bytes File Sequence Number
This word contains the file sequence number of the file.
- 3.4.1.5 H.FLEV 2 Bytes File Structure Level
The file structure level is used to identify different versions of Files-11 as they affect the structure of the file header. This permits upwards compatibility of file structures as Files-11 evolves, in that the structure level word identifies the version of Files-11 that created this particular file. This document describes version 1 of Files-11; the only legal contents for H.FLEV is 401 octal.
- 3.4.1.6 H.FOWN 2 Bytes File Owner UIC
H.PROG = H.FOWN+0 Programmer (Member) Number
H.PROJ = H.FOWN+1 Project (Group) Number
This word contains the binary user identification code (UIC) of the owner of the file. The file owner is usually (but not necessarily) the creator of the file.
- 3.4.1.7 H.FPRO 2 Bytes File Protection Code
This word controls what access all users in the system may have to the file. Accessors of a file are categorized according to the relationship between the UIC of the accessor and the UIC of the owner of the file. Each category is controlled by a four bit field in the protection word. The category of the accessor is selected as follows:
- System Bits 0 - 3
The accessor is subject to system protection if the project number of the UIC under which he is running is 10 octal or less.
- Owner Bits 4 - 7
The accessor is subject to owner protection if the UIC under which he is running exactly matches the file owner UIC.

- Group Bits 8 - 11
The accessor is subject to group protection if the project number of his UIC matches the project number of the file owner UIC.
- World Bits 12 - 15
The accessor is subject to world protection if he does not fit into any of the above categories.
- Four types of access intents are defined in Files-11: read, write, extend, and delete. Each four bit field in the protection word is bit encoded to permit or deny any combination of the four types of access to that category of accessors. Setting a bit denies that type of access to that category. The bits are defined as follows (these values apply to a right-justified protection field):
- FP.RDV Deny read access
FP.WRV Deny write access
FP.EXT Deny extend access
FP.DEL Deny delete access
- When a user attempts to access a file, protection checks are performed in all the categories to which he is eligible, in the order system - owner - group - world. The user is granted access to the file if any of the categories to which he is eligible grants him access.
- 3.4.1.8 H.FCHA 2 Bytes File Characteristics
H.UCHA = H.FCHA+0 User Controlled Char.
H.SCHA = H.FCHA+1 System Controlled Char.
- The user controlled characteristics byte contains the following flag bits:
- 1 Bit, Reserved.
- UC.NID Set if incremental dump (backup) is to be disabled for this file.
- UC.WBC Set if the file is to be write-back cached; i.e., if a cache is used for the file data, data written by a user is only written back to the disk when is it removed from the cache. Clear for write-through cache operation.

UC.RCK Set if the file is to be read-checked. All read operations on the file, including reads of the file header(s), will be performed with a read, read-compare to assure data integrity.

UC.WCK Set if the file is to be write-checked. All write operations on the file, including modifications of the file header(s), will be performed with a write, read-compare to assure data integrity.

UC.CNB Set if the file is allocated contiguous best effort; i.e., as contiguous as possible.

UC.DLK Set if the file is deaccess-locked. This bit is used as a flag warning that the file was not properly closed and may contain inconsistent data. Access to the file is denied if this bit is set.

UC.CON Set if the file is logically contiguous; i.e., if for all virtual blocks in the file, virtual block i maps to logical block $k+i$ on one unit for some constant k . This bit may be implicitly set or cleared by file system operations that allocate space to the file; the user may only clear it explicitly.

The system controlled characteristics byte contains the following flag bits:

- 3 Bits, Reserved.
- Reserved (Access Control List).

SC.SPL Set if the file is an intermediate file for spooling.

SC.DIR Set if the file is a directory.

SC.BAD Set if there is a bad data block in the file. This bit is as yet unimplemented. It is intended for dynamic bad block handling.

SC.MDL Set if the file is marked for delete. If this bit is set, further accesses to the file are denied, and the file will be physically deleted when no users are accessing it.

3.4.1.9 H.UFAT 32 Bytes User Attribute Area

This area is intended for the storage of a limited quantity of "user file attributes", i.e., any data the user deems useful for processing the file that is not part of the file itself. An example of the use of the user attribute area is presented in section 6.1 (FCS File Format).

3.4.1.10 S.HDHD 46 Bytes Size of Header Area

This symbol represents the total size of the header area containing all of the above entries.

3.4.2 Ident Area Description

The ident area of the file header begins at the word indicated by H.IDOF. It contains identification and accounting data about the file.

3.4.2.1 I.FNAM 6 Bytes File Name

These three words contain the name of the file, packed three Radix-50 characters to the word. This name usually, but not necessarily, corresponds to the name of the file's primary directory entry.

3.4.2.2 I.FTYP 2 Bytes File Type

This word contains the type of the file in the form of three Radix-50 characters.

3.4.2.3 I.FVER 2 Bytes Version Number

This word contains the version number of the file in binary form.

3.4.2.4 I.RVNO 2 Bytes Revision Number

This word contains the revision count of the file. The revision count is the number of times the file has been accessed for write.

3.4.2.5 I.RVDT 7 Bytes Revision Date

The revision date is the date on which the file was last deaccessed after being accessed for write. It is stored in ASCII in the form "DDMMYY", where DD is two digits representing the day of the month, MMM is three characters representing the month, and YY is the last two digits of the year.

3.4.2.6 I.RVTI 6 Bytes Revision Time

The revision time is the time of day on which the file was last deaccessed after being accessed for write. It is stored in ASCII in the format "HHMMSS", where HH is the hour, MM is the minute, and SS is the second.

3.4.2.7 I.CRDT 7 Bytes Creation Date

These seven bytes contain the date on which the file was created. The format is the same as that of the revision date above.

3.4.2.8 I.CRTI 6 Bytes Creation Time

These six bytes contain the time of day at which the file was created. The format is the same as that of the revision time above.

3.4.2.9 I.EXDT 7 Bytes Expiration Date

These seven bytes contain the date on which the file becomes eligible to be deleted. The format is the same as that of the revision and creation dates above.

3.4.2.10 - 1 Byte (unused)

This unused byte is present to round up the size of the ident area to a word boundary.

3.4.2.11 S.IDHD 46 Bytes Size of Ident Area

This symbol represents the size of the ident area containing all of the above entries.

3.4.3 Map Area Description

The map area of the file header starts at the word indicated by H.MPOF. It contains the information necessary to map the virtual blocks of the file to the logical blocks of the volume.

3.4.3.1 M.ESON 1 Byte Extension Segment Number

This byte contains the value n, where this header is the n+1th header of the file; i.e., headers of a file are numbered sequentially starting with 0.

3.4.3.2 M.FRVN 1 Byte Extension Relative Volume No.

This byte contains the relative volume number of the unit in the volume set that contains the next

sequential extension header for this file. If there is no extension header, or if the extension header is located on the same unit as this header, this byte contains 0.

3.4.3.3 M.EFNU 2 Bytes Extension File Number

This word contains the file number of the next sequential extension header for this file. If there is no extension header, this word contains 0.

3.4.3.4 M.EFSQ 2 Bytes Extension File Sequence Number

This word contains the file sequence number of the next sequential extension header for this file. If there is no extension header, this word contains 0.

3.4.3.5 M.CTSZ 1 Byte Block Count Field Size

This byte contains a count of the number of bytes used to represent the count field in the retrieval pointers in the map area. The retrieval pointer format is described in section 3.4.3.9 below.

3.4.3.6 M.LBSZ 1 Byte LBN Field Size

This byte contains a count of the number of bytes used to represent the logical block number field in the retrieval pointers in the map area. The contents of M.CTSZ and M.LBSZ must add up to an even number.

3.4.3.7 M.USE 1 Byte Map Words In Use

This byte contains a count of the number of words in the map area that are presently occupied by retrieval pointers.

3.4.3.8 M.MAX 1 Byte Map Words Available

This byte contains the total number of words available for retrieval pointers in the map area.

3.4.3.9 M.RTRV variable Retrieval Pointers


This area contains the retrieval pointers that actually map the virtual blocks of the file to the logical blocks of the volume. Each retrieval pointer describes a consecutively numbered group of logical blocks which is part of the file. The count field contains the binary value n to represent a group of n+1 logical blocks. The logical block number field contains the logical

block number of the first logical block in the group. Thus each retrieval pointer maps virtual blocks j through $j+n$ into logical blocks k through $k+n$, respectively, where j is the total number plus one of virtual blocks represented by all preceding retrieval pointers in this and all preceding headers of the file, n is the value contained in the count field, and k is the value contained in the logical block number field.

Although the data in the map area provides for arbitrarily extensible retrieval pointer formats, Files-11 has defined only three. Of these, only the first is currently implemented; the other two are presented out of historical interest; they will never be supported.


Format 1: M.CTSZ = 1
M.LBSZ = 3

The total retrieval pointer length is four bytes. Byte 1 contains the high order bits of the 24 bit LBN. Byte 2 contains the count field, and bytes 3 and 4 contain the low 16 bits of the LBN.



Format 2: M.CTSZ = 2
M.LBSZ = 2


The total retrieval pointer length is four bytes. The first word contains a 16 bit count field and the second word contains a 16 bit LBN field.



Format 3: M.CTSZ = 2
M.LBSZ = 4

The total retrieval pointer length is six bytes. The first word contains a 16 bit count field and the second and third

words contain a 32 bit LBN field.



3.4.3.10 S.MPHD 10 Bytes Size of Map Area

This symbol represents the size of the map area, not including the space used for the retrieval pointers.

3.4.4 End Checksum Description

The header check sum occupies the last two bytes of the file header. It is verified every time a header is read, and is recomputed every time a header is written.

3.4.4.1 H.CKSM 2 Bytes Block Checksum

This word is a simple additive checksum of all other words in the block. It is computed by the following PDP-11 routine or its equivalent:

```

MOV      Header-address,R0
CLR      R1
MOV      #255.,R2
10$:    ADD      (R0)+,R1
        SOB     R2,10$
        MOV     R1,(R0)

```

3.4.A File Header Layout

Header Area

H.MPOF	Map Area Offset	Ident Area Offset	H.IDOF
	File Number		H.FNUM
	File Sequence Number		H.FSEQ
	File Structure Level		H.FLEV
H.PROJ	File Owner UIC		H.FOWN
	File Protection		H.FPRO
H.SCHA	System Char.	User Char.	H.FCHA
	User Attribute Area		H.UCHA
			H.UFAT
			S.HDHD

Ident Area

	File Name	I.FNAM
	File Type	I.FTYP
	Version Number	I.FVER
	Revision Number	I.RVNO
	Revision Date	I.RVDT
I.RVTI	Revision Time	

I.CRDT	Creation Date	I.CRTI
	Creation Time	I.EXDT
	Expiration Date	S.IDHD
	(not used)	

Map Area

M.ERVN	Extension RVN	Ext. Seg. Num.	M.ESQN
	Extension File Number		M.EFNU
	Extension File Seq. Num.		M.EFSQ
M.LBSZ	LBN Field Size	Count Field Size	M.CTSZ
M.MAX	Map Words Avail.	Map Words in Use	M.USE
	Retrieval Pointers		S.MPHD
	File Header Checksum		M.RTRV
			H.CKSM

4.0 Directories

Files-11 provides directories to allow the organization of files in a meaningful way. While the File ID is sufficient to locate a file uniquely on a volume set, it is hardly mnemonic. Directories are files whose sole function is to associate file name strings with File ID's.

4.1 Directory Heirarchies

Since directories are files with no special attributes, directories may list files that are in turn directories. Thus the user may construct directory heirarchies of arbitrary depth and complexity to structure his files as he pleases.

4.1.1 User File Directories

Current implementations of Files-11 all support a two level directory heirarchy which is tied in with the user identification mechanism of the operating system. Each UIC is associated with a user file directory (UFD). References to files that do not specify a directory are generally defaulted to the UFD associated with the user's UIC. All UFD's are listed in the volume's MFD under a file name constructed from the UIC. A UIC of [n,m] associates with a directory name of "nnnmmm.DIR;1", where nnn and mmm are n and m padded out to three digits each with leading zeroes. Note that all number conversions are done in octal.

Two points should be noted here. The UFD structure described here is not intrinsically part of the Files-11 on-disk structure; rather, it is a convenient cataloging system applied by various operating systems. Also, there is no hard and fast relationship between the owner UIC of a file and the UFD in which it is listed. Generally, they will correspond, but not necessarily.

4.2 Directory Structure

A directory is a file consisting of 16 byte records. It is structured as an FCS fixed length record file, with no carriage control attributes (see section 6 for a description of FCS files). Each record is a directory entry. The entries are not required to be ordered, or densely packed, nor do they have any other relationship to each other, except that no two entries in one directory may contain the same name, type, and version. Each entry contains the following:

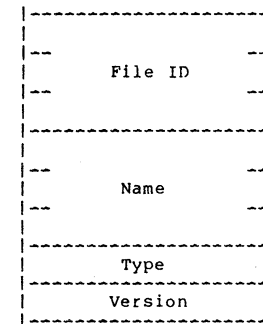
File ID The three word binary File ID of the file that this directory entry represents. If the file number portion of the File ID field is zero, then this record is empty and may be used for a new

directory entry.

Name The name of the file may be up to 9 characters. It is stored as three words, each containing three Radix-50 packed characters.

Type The type of the file (also historically referred to as the extension) may be up to three characters. It is stored as one word of Radix-50 packed characters.

Version The version number of the file is stored in binary in one word.



4.3 Directory Protection

Since directories are files with no special characteristics, they may be accessed like all other files, and are subject to the same protection mechanism. However, implementations of Files-11 support three special functions for the management of directories, namely FIND, REMOVE, and ENTER. A user performing such a directory operation must have the following privileges to be allowed the various functions:

Find: READ
 Remove: READ, WRITE
 Enter: READ, WRITE

Note that the same privilege is required for both enter and remove. The recovery for an operation that involves a remove at the beginning of the sequence is an enter.

5.0 Known Files

Clearly any file system must maintain some data structure on the medium which is used to control the file organization. In Files-11 this data is kept in five files. These files are created when a new volume is initialized. They are unique in that their File ID's are known constants. These five files have the following uses:

File ID 1,1,0 is the index file. The index file is the root of the entire Files-11 structure. It contains the volume's bootstrap block and the home block, which is used to identify the volume and locate the rest of the file structure. The index file also contains all of the file headers for the volume, and a bitmap to control the allocation of file headers.

File ID 2,2,0 is the storage bitmap file. It is used to control the allocation of logical blocks on the volume.

File ID 3,3,0 is the bad block file. It is a file containing all of the known bad blocks on the volume.

File ID 4,4,0 is the volume master file directory (or MFD). It forms the root of the volume's directory structure. The MFD lists the five known files, all first level user directories, and whatever other files the user chooses to enter.

File ID 5,5,0 is the system core image file. Its use is operating system dependent; its basic purpose is to provide a file of known File ID for the use of the operating system.

5.1 Index File

The index file is File ID 1,1,0. It is listed in the MFD as INDEXF.SYS;1. The index file is the root of the Files-11 structure in that it provides the means for identification and initial access to a Files-11 volume, and contains the access data for all files on the volume (including itself).

5.1.1 Bootstrap Block

Virtual block 1 of the index file is the volume's boot block. It is always mapped to logical block 0 of the volume. If the volume is the system device of an operating system, the boot block contains an operating system dependent program which reads the operating system into memory when the boot block is read and executed by a machine's hardware bootstrap. If the volume is not a system device, the boot block contains a small program that outputs

a message on the system console to inform the operator to that effect.

5.1.2 Home Block

Virtual block 2 of the index file is the volume's home block. The logical block containing the home block is the first good block on the volume out of the sequence 1, 256, 512, 768, 1024, 1280, 256n. The purpose of the home block is to identify the volume as Files-11, establish the specific identity of the volume, and serve as the ground zero entry point into the volume's file structure. The home block is recognized as a home block by the presence of checksums in known places and by the presence of predictable values in certain locations.

Items contained in the home block are identified by symbolic offsets in the same manner as items in the file header. The symbols may be defined in assembly language programs by calling and invoking the macro HMBOFS, which may be found in the macro library of any system that supports Files-11. Alternatively, one may find the macro in the file FILEMAC.MAC, which is available from the author.

5.1.2.1 H.IBSZ 2 Bytes Index File Bitmap Size

This 16 bit word contains the number of blocks that make up the index file bitmap. (The index file bitmap is discussed in section 5.1.3.) This value must be non-zero for a valid home block.

5.1.2.2 H.IBLB 4 Bytes Index File Bitmap LBN

This double word contains the starting logical block address of the index file bitmap. Once the home block of a volume has been found, it is this value that provides access to the rest of the index file and to the volume. The LBN is stored with the high order in the first 16 bits, followed by the low order portion. This value must be non-zero for a valid home block.

5.1.2.3 H.FMAX 2 Bytes Maximum Number of Files

This word contains the maximum number of files that may be present on the volume at any time. This value must be non-zero for a valid home block.

5.1.2.4 H.SBCL 2 Bytes Storage Bitmap Cluster Factor

This word contains the cluster factor used in the storage bitmap file. The cluster factor is the number of blocks represented by each bit in the

storage bitmap. Volume clustering can not be implemented in ODS-1; the only legal value for this item is 1.

5.1.2.5 H.DVTY 2 Bytes Disk Device Type

This word is an index identifying the type of disk that contains this volume. It is currently not used and always contains 0.

5.1.2.6 H.VLEV 2 Bytes Volume Structure Level

This word identifies the volume's structure level. Like the file structure level, this word identifies the version of Files-11 which created this volume and permits upwards compatibility of media as Files-11 evolves. The volume structure level is affected by all portions of the Files-11 structure except the contents of the file header. This document describes Files-11 version 1; the only legal values for the structure level are 401 and 402 octal. The former (401) is the standard value for most volumes. The latter (402) is an advisory that the volume contains a multiheader index file. (A multiheader index file is required to support more than about 26,000 files. The index file may in fact be multiheader without the volume having a structure level of 402).

5.1.2.7 H.VNAM 12 Bytes Volume Name

This area contains the volume label as an ASCII string. It is padded out to 12 bytes with nulls. The volume label is used to identify individual Files-11 volumes.

5.1.2.8 - 4 Bytes Not Used

5.1.2.9 H.VOWN 2 Bytes Volume Owner UIC

This word contains the binary UIC of the owner of the volume. The format is the same as that of the file owner UIC stored in the file header.

5.1.2.10 H.VPRO 2 Bytes Volume Protection Code

This word contains the protection code for the entire volume. Its contents are coded in the same manner as the file protection code stored in the file header, and it is interpreted in the same way in conjunction with the volume owner UIC. All operations on all files on the volume must pass both the volume and the file protection check to be permitted. (Refer to the discussion on file protection in section 3.4.1.7).

5.1.2.11 H.VCHA 2 Bytes Volume Characteristics

This word contains bits which provide additional control over access to the volume. The following bits are defined:

CH.NDC Obsolete, used by RSX-11D/IAS. Set if device control functions are not permitted on this volume. Device control functions are those which can threaten the integrity of the volume, such as direct reading and writing of logical blocks, etc.

CH.NAT Obsolete, used by RSX-11D/IAS. Set if the volume may not be attached, i.e., reserved for the sole use by one task.

CH.SDI Set if the volume contains only a single directory. If this bit is set, no directories should be created on the volume other than the MPD. The access methods should also be informed of this situation, e.g. by setting the DV.SDI bit in the device characteristics word.

5.1.2.12 H.DFPR 2 Bytes Default File Protection

This word contains the file protection that will be assigned to all files created on this volume if no file protection is specified by the user.

5.1.2.13 - 6 Bytes Not Used

5.1.2.14 H.WISZ 1 Byte Default Window Size

This byte contains the number of retrieval pointers that will be used for the "window" (in core file access data) when files are accessed on the volume, if not otherwise specified by the accessor.

5.1.2.15 H.FIEX 1 Byte Default File Extend

This byte contains the number of blocks that will be allocated to a file when a user extends the file and asks for the system default value for allocation.

5.1.2.16 H.LRUC 1 Byte Directory Pre-access Limit

This byte contains a count of the number of directories to be stored in the file system's directory access cache. More generally, it is an estimate of the number of concurrent users of the

volume and its use may be generalized in the future.

5.1.2.17 H.REVD 7 Bytes Date of Last Home Block Revision

This field ill defined field is in the standard ASCII date format and reflects the date of the last modifications to fields in the home block.

5.1.2.17 H.REVC 2 Bytes Count of Home Block Revisions

This field reflects the number of above mentioned modifications.

5.1.2.17 - 2 Bytes Not Used

5.1.2.18 H.CHK1 2 Bytes First Checksum

This word is an additive checksum of all entries preceding in the home block (i.e., all those listed above). It is computed by the same sort of algorithm as the file header checksum (see section 3.4.4.1).

5.1.2.19 H.VDAT 14 Bytes Volume Creation Date

This area contains the date and time that the volume was initialized. It is in the format "DDMMYYHHMSS", followed followed by a single null. (The same format is used in the ident area of the file header, section 3.4.2).

5.1.2.20 - 382 Bytes Not Used

This area is reserved for the relative volume table for volume sets. This field will not be used, although some versions of DSC referenced this area.

5.1.2.21 H.PKSR 4 Bytes Pack Serial Number

This area contains the manufacturer supplied serial number for the physical volume. For last track devices, the pack serial number is contained on the volume in the manufacturer data. For other devices the user must supply this information manually. The serial number is contained in the home block for convenience and consistency.

5.1.2.22 - 12 Bytes Not Used

5.1.2.23 H.INDN 12 Bytes Volume Name

This area contains another copy of the ASCII volume label. It is padded out to 12 bytes with spaces. It is placed here in accordance with the volume identification standard (STD 167).

5.1.2.24 H.INDO 12 Bytes Volume Owner

This area contains an ASCII expansion of the volume owner UIC in the form "[proj,prog]". Both numbers are expressed in decimal and are padded to three digits with leading zeroes. The area is padded out to 12 bytes with trailing spaces. It is placed here in accordance with the volume identification standard (STD 167).

5.1.2.25 H.INDF 12 Bytes Format Type

This field contains the ASCII string "DECFILE11A" padded out to 12 bytes with spaces. It identifies the volume as being of Files-11 format. It is placed here in accordance with the volume identification standard (STD 167).

5.1.2.26 - 2 Bytes Not Used


5.1.2.27 H.CHK2 2 Bytes Second Checksum

This word is the last word of the home block. It contains an additive checksum of the preceding 255 words of the home block, computed according to the algorithm listed in section 3.4.4.1.

5.1.2.A Home Block Layout

	Index File Bitmap Size	H.IBSZ
	Index File	H.IBLR
	Bitmap LBN	
	Maximum Number of Files	H.FMAX
	Storage Bitmap Cluster Factor	H.SBCL
	Disk Device Type	H.DVTY
	Volume Structure Level	H.VLEV
		H.VNAM
	Volume Name	
	(not used)	
	Volume Owner UIC	H.VOWN
	Volume Protection	H.VPRO
	Volume Characteristics	H.VCHA
	Default File Protection	H.DFPR
	(not used)	
H.PIEX	Def. File Extend	H.WISZ
H.REVD	Def. Window Size	
	Directory Limit	H.LRUC
	Volume Modification Date	

	Volume Modification Count	H.REVC
	(not used)	
	First Checksum	H.CHK1
		H.VDAT
	Volume Creation Date	
	(not used)	
	Pack Serial Number	H.PKSR
	(not used)	
	Volume Name	H.INDN



5.1.3 Index File Bitmap

The index file bitmap is used to control the allocation of file numbers (and hence file headers). It is simply a bit string of length n , where n is the maximum number of files permitted on the volume (contained in offset H.FMAX in the home block). The bitmap spans over as many blocks as is necessary to hold it, i.e., max number of files divided by 4096 and rounded up. The number of blocks in the bitmap is contained in offset H.IBSZ of the home block.

The bits in the index file bitmap are numbered sequentially from 0 to $n-1$ in the obvious manner, i.e., from right to left in each byte, and in order of increasing byte address. Bit j is used to represent file number $j+1$: if the bit is 1, then that file number is in use; if the bit is 0, then that file number is not in use and may be assigned to a newly created file.

The index file bitmap starts at virtual block 3 of the index file and continues through VBN $2+m$, where m is the number of blocks in the bitmap. It is located at the logical block indicated by offset H.IBLB in the home block.

5.1.4 File Headers

The rest of the index file contains all the file headers for the volume. The first 16 file headers (for file numbers 1 to 16) are logically contiguous with the index file bitmap to facilitate their location; the rest may be allocated wherever the file system sees fit. Thus the first 16 file headers may be located from data in the home block (H.IBSZ and H.IBLB) while the rest must be located through the mapping data in the index file header. The file header for file number n is located at virtual block $2+m+n$ (where m is the number of blocks in the index file bitmap).

5.2 Storage Bitmap File

The storage bitmap file is File ID 2,2,0. It is listed in the MFD as BITMAP.SYS;1. The storage bitmap is used to control the available space on a unit. It consists of a storage control block which contains summary information about the unit, and the bitmap itself which lists the availability of individual blocks.

5.2.1 Storage Control Block

Virtual block 1 of the storage bitmap is the storage control block. It contains summary information intended to optimize allocation of space on the unit. The storage control block has the following format for disks with less than 4096126, (516,096 blocks):

- (3 bytes) Not used (zero)
- (1 byte) Number of storage bitmap blocks (less than 127)
- (2 bytes) Number of free blocks in 1st bitmap block
- (2 bytes) Free block pointer in 1st bitmap block
- .
- .
- (2 bytes) Number of free blocks in nth bitmap block
- (2 bytes) Free block pointer in nth bitmap block
- (4 bytes) Size of the unit in logical blocks

For larger disks the following format is used:

- (3 bytes) Not used (zero)
- (1 byte) Number of storage bitmap blocks (greater than 126)
- (4 bytes) Size of the unit in logical blocks
- (246 bytes) Not used (zero)

Note: Current implementations of Files-11 do not correctly initialize the word pairs containing number of free blocks and free block pointer for each bitmap block, nor are these values maintained as space is allocated and freed on the unit. They are therefore best looked upon as 2n garbage words and should not be used by future implementations of Files-11 until the disk structure is formally updated.

5.2.2 Storage Bitmap

Virtual blocks 2 through n+1 are the storage bitmap itself. It is best viewed as a bit string of length m, numbered from 0 to m-1, where m is the total number of logical blocks on the unit rounded up to the next multiple of 4096. The bits are addressed in the usual manner (packed right to left in sequentially numbered bytes). Since each virtual block holds 4096 bits, n blocks, where $n = m/4096$, are used to hold the bitmap. Bit j of the bitmap represents logical block j of the volume; if the bit is set, the block is free; if clear, the block is allocated. Clearly the last k bits of the bitmap are always clear, where k is the difference between the true size of the volume and m, the length of the bitmap.

The size of the bitmap is limited to 256 blocks. In fact, due to existing implementations on all RSX systems, the retrieval pointers must be in one of the following two forms:

1. A single retrieval pointer mapping the entire BITMAP.SYS file.
2. Two retrieval pointers, the first mapping the storage control block only, and the second mapping the entire bitmap proper.

This restriction limits ODS-1 to a volume of 4096255 blocks (1,044,480 blocks or about 500 megabytes).

5.3 Bad Block File


The bad block file is File ID 3,3,0. It is listed in the MFD as BADBLK.SYS;1. The bad block file is simply a file containing all of the known bad blocks on the volume.

5.3.1 Bad Block Descriptor

Virtual block 1 of the bad block file is the bad block descriptor for the volume. It is always located on the last good block of the volume. This block may contain a listing of the bad blocks on the volume produced by a bad block scan program or diagnostic. The format of the bad block data is

identical to the map area of a file header, except that the first four entries (M.ESQN, M.ERVN, M.EFNU, and M.EFSO) are not present. The last word of the block contains the usual additive checksum. (See section 3.4.3 for a description of the map area.) This block is included in the bad block file to save the data it contains for future re-initialization of the volume.

Bad Block Descriptor Layout



5.4 Master File Directory

The master file directory is File ID 4,4,0. It is listed in the MFD (itself) as 000000.DIR;1. The MFD is the root of the volume's directory structure. It lists the five known files, plus whatever the user chooses to enter. In the two level UFD structure described in section 4.1.1, the MFD contains entries for all user file directories.

5.5 Core Image File

The core image file is File ID 5,5,0. It is listed in the MFD as CORIMG.SYS;1. Its use is operating system dependent. In general, it provides a file of known File ID for the use of the operating system, for use as a swap area, for example, or as a monitor overlay area, etc.

6.0 FCS File Structure

File Control Services (FCS) is a user level interface to Files-11. Its principal feature is a record control facility that allows sequential processing of variable length records and sequential and random access to fixed length record files. FCS interfaces to the virtual block

facility provided by the basic Files-11 structure.

6.1 FCS File Attributes

FCS stores attribute information about the file in the file's user attribute area (H.UFAT - see section 3.4.1.9). It uses only the first 7 words; the rest are ignored by FCS, but are reserved by DEC. (RMS uses an additional 3 words, 10 words in all, for relative and indexed file attributes.) The following items are contained in the attribute area; they are identified by the usual symbolic offsets (relative to the start of the attribute area). The offsets may be defined in assembly language programs by calling and invoking the macro FDOFPS DEFSL. Flag values and bits may be defined by calling and invoking the macro FCSRTS. These macros are in the system macro library of any operating system that supports Files-11. Alternatively, all these values are defined in the system object library of any system that supports Files-11, and may be obtained at link time.

6.1.1 F.RTYP 1 Byte Record Type

This byte identifies which type of records are contained in this file. The following three values are legal:

R.FIX Fixed length records.
R.VAR Variable length records.
R.SEQ Sequenced Variable Length records

6.1.2 F.RATT 1 Byte Record Attributes

This byte contains record attribute bits that control the handling of records under various contexts. The following flag bits are defined:

FD.FTN Use Fortran carriage control if set. The first byte of each record is to be interpreted as a standard Fortran carriage control character when the record is copied to a carriage control device.

FD.CR Use implied carriage control if set. When the file is copied to a carriage control device, each record is to be preceded by a line feed and followed by a carriage return. Note that the FD.FTN and FD.CR bits are mutually exclusive.

FD.PRN Used to indicate that the two byte sequence number field for R.SEQ record format is to be interpreted as print

control information (see Section 6.2.3.1 for format of print information).

FD.BLK Records do not cross block boundaries if set. Generally, there will be dead space at the end of each block; how this is handled is explained in the description of record formats in section 6.2.

6.1.3 F.RSIZ 2 Bytes Record Size

In a fixed length record file, this word contains the size of the records in bytes. In a variable or sequenced variable length record file, this word contains the size in bytes of the longest record in the file.

6.1.4 F.HIBK 4 Bytes Highest VBN Allocated

This 32 bit number is a count of the number of virtual blocks allocated to the file. Since this value is maintained by FCS, it is usually correct, but it is not guaranteed since FCS is a user level package.

6.1.5 F.EFBK 4 Bytes End of File Block

This 32 bit number is the VBN in which the end of file is located. Both F.HIBK and F.EFBK are stored with the high order half in the first two bytes, followed by the low order half.

6.1.6 F.PFBY 2 Bytes First Free Byte

This word is a count of the number of bytes in use in the virtual block containing the end of file; i.e., it is the offset to the first byte of the file available for appending. Note that an end of file that falls on a block boundary may be represented in either of two ways. If the file contains precisely n blocks, F.EFBK may contain n and F.PFBY will contain 512, or F.EFBK may contain n+1 and F.PFBY will contain 0.

6.1.7 S.FATT 14 Bytes Size of Attribute Block

This symbol represents the total number of bytes in the FCS file attribute block.

6.1.A FCS File Attributes Layout

|-----|-----|

F.RATT	Record Attr.	Record Type	F.RTYP
	Record Size (Bytes)		F.RSIZ
	Highest VBN		F.HIBK
	Allocated		
	End of File		F.EPBK
	VBN		
	First Free Byte		F.FFBY S.FATT

6.2 Record Structure

This section describes how records are packed in the virtual blocks of a disk file. In general, FCS treats a disk file as a sequentially numbered array of bytes. Records are numbered consecutively starting with 1.

6.2.1 Fixed Length Records

In a file consisting of fixed length records, the records are simply packed end to end with no additional control information. If the record length is odd, each record is padded with a single byte. The content of the pad byte is undefined. For direct access, the address of a record is computed as follows:

Let: n = record number
 k = record size (in bytes)
 m = byte address of record in file
 q = number of records per block
 j = VBN containing the start of the record
 i = byte offset within VBN j

then $h = ((k+1)/2)2$ (rounded up record length)
 $m = (n-1)h$
 $j = m/512+1$ (truncated)
 $i = m \bmod 512$

The previous discussion assumes that records cross block boundaries (that is, FD.BLK is not set). If records do not cross block boundaries, they are limited to 512 bytes, and the following equations apply (the variables are defined as above):

$h = ((k+1)/2)2$ (rounded up record length)
 $q = 512/k$ (truncated)
 $j = (n-1)/q+1$ (truncated)
 $i = ((n-1) \bmod q)h$

6.2.2 Variable Length Records

In a file consisting of variable length records, records may be up to 32767 bytes in length. Each record is preceded by a two byte binary count of the bytes in the record (the count does not include itself). For example, a null record is represented by a single zero word. The byte count is always word aligned; i.e., if a record ends on an odd byte boundary, it is padded with a single byte. The content of the pad byte is undefined.

If records do not cross block boundaries (FD.BLK is set), they are limited to a size of 512 bytes. A byte count of -1 is used as a flag to signal that there are no more records in a particular block. The remainder of that block is then dead space and the next record in the file starts at the beginning of the next block.

6.2.3 Sequenced Variable Length Records

The format of a sequenced file is identical to a variable length record file except that a two byte sequence number field is located immediately after the byte count field of each record. This field contains a binary value which is usually interpreted as the line number of that record (see Section 6.1.2 FD.PRN and Section 6.2.3.1). The sequence number is not returned as part of the data when a record is read, but is available separately. Note that the record byte count field counts the sequence number field as well as the data of the record.

6.2.3.1 Format of Two Byte Print Control Field in R.SEO Records

If the FD.PRN bit is set in the record attribute then the two byte "sequence number" field is used to contain carriage control data for the record. Byte 0 is print control information to act upon before the record data is output to a unit record device; byte 1 is print control information to act upon after the record data has been output to a unit record device.

The format of each byte is as follows:

Bit 7	Bits 6-0	Meaning		
0	0	No carriage control		
0	count(1-127)	"count" new lines (CR/LF)		
Bit 7	Bit 6	Bit 5	Bits 4-0	Meaning

1	0	0	ASCII C0 set	ASCII char to output (CR,FF etc.)
1	0	1	ASCII C1 set	ASCII char (8 bit code) to output
1	1	0	CODE (0-63)	Device specific code
1	1	1	-	Reserved

NOTE

The print control field is not currently supported
by FCS or RMS-11.