

Using OSTA UDF on CD-R and CD-RW media

Version 0.10

December 4, 1996

Table of Contents

1. Overview.....	4
1.1 History.....	4
2. References.....	5
3. Definitions.....	6
4. Terms.....	8
5. Description.....	9
5.1 General.....	9
5.2 Information Control Block.....	9
5.3 Sequential File System Specific.....	10
5.3.1 Virtual Allocation Table.....	10
5.3.2 Virtual Partition Map.....	12
5.3.3 End of session data.....	12
5.4 Random Access File System Specific.....	13
5.4.1 Formatting.....	13
5.4.2 Sparable Partition Map.....	15
5.4.3 Host Based Defect Management.....	15
5.4.4 Read Modify Write Operation.....	15
5.4.5 Sparing Table.....	15
5.4.6 Unallocatable Areas.....	15
5.4.7 Levels of Compliance.....	15
6. Multisession and Mixed Mode.....	16
6.1 Volume Recognition Sequence.....	16
6.2 Anchor Volume Descriptor Pointer.....	16
6.3 UDF Bridge format.....	17
7. Sequential Access Implementation Strategies.....	18
7.1 Allocation of space.....	18
7.1.1 Strategy 1.....	18
7.1.2 Strategy 2.....	18
7.1.3 Strategy 3.....	18
7.2 Duplicate VAT.....	19
7.3 Duplicate ICB.....	19
8. Sequential File System Sample Sequence of Events.....	20
9. Sequential File System Example Disc Images (strategy 1).....	22

List of Tables

Table 1 - End of session data.....	12
Table 2 - Sector Map for formatted "blank" disc.....	22
Table 3 - Sector Map for disc with added directory and file.....	23
Table 4 - Underrun while recording a file.....	24

List of Figures

Figure 1 - Multisession UDF disc.....	17
Figure 2 - CD enhanced disc.....	17
Figure 3 - ISO 9660 converted to UDF.....	17
Figure 4 - Foreign format converted to UDF.....	18

1. Overview

This document describes implementing the UDF file system for use with CD recorders (either CD-R or CD-RW) as a "logical device" on a computer system. A logical device is one that is used in the same manner as a magnetic disk.

The technical details of the file system are described in OSTA UDF (Universal Disk Format) version 1.5 or higher and ECMA 167 / ISO 13346. OSTA UDF is a named domain of ISO 13346, and is thus entirely a subset of ISO 13346. Please see OSTA UDF for more information.

The implementation will use UDF as the interactive file system with an allowance for writing ISO 9660 structures when interchange is desired.

References enclosed in [] are references to ISO 13346. The references are in the form [x/a.b.c], where x is the section number and a.b.c is the paragraph or figure number.

1.1 History

OSTA UDF has been adopted as the file system for DVD players. Using this format instead of proprietary intermediate file systems will reduce confusion in the industry and provide a clear migration path to DVD.

UDF is the end result of several years of collaborative effort, including the effort put into developing ISO 13346, which in turn was based on ECMA 167. This proposal is a small incremental change to an existing standard. The added structures conform to the standard's conventions.

The CD-R file system is almost completely underrun proof; the CD-RW file system is completely underrun proof. The only parts that must be recorded without an underrun are the two volume descriptor sequences. As these structures are no more than 16 sectors each, this should not present a problem. The cost of an underrun during file recording is simply the link blocks (7); no data written to the disc needs to be re-recorded. No structures need to have pointers to unrecorded data which could become invalid if an underrun occurs.

No reserved tracks are required (though are allowed for certain performance enhancements), which eliminates the need for guessing sizes required for the file system relative to the data. It provides a deterministic method for finding all information; no searching for the most recent descriptors is necessary.

2. References

- ECMA 167 Volume and file structure of write-once and rewritable media using non-sequential recording for information interchange. This document was the basis for ISO 13346 and, in its second release, reflects the changes made to produce ISO 13346.
- ISO 9660:1988 Information Processing - Volume and File Structure of CD-ROM for Information Interchange
- IEC 908:1987 Compact disc digital audio system
- ISO/IEC 10149:1993 Information technology - Data Interchange on read-only 120mm optical data disks (CD-ROM based on the Philips/Sony "Yellow Book")
- Orange Book part-II Recordable Compact Disc System Part-II, N.V. Philips and Sony Corporation
- Orange Book part-III Recordable Compact Disc System Part-III, N.V. Philips and Sony Corporation
- ISO/IEC 13346:1995 Volume and file structure of write-once and rewritable media using non-sequential recording for information interchange. References enclosed in [] in this document are references to ISO 13346. The references are in the form [x/a.b.c], where x is the section number and a.b.c is the paragraph or figure number.
- OSTA UDF OSTA Universal Disk Format version 1.50. OSTA UDF is also called UDF in this document.

3. Definitions

Audio session	Audio session contains one or more audio tracks, and no data track.
Audio track	Audio tracks are tracks that are designated to contain audio sectors specified in the ISO/IEC 908.
CD-RW	CD-ReWritable. An overwritable CD defined in Orange Book, part-III.
CD-R	CD-Recordable. A write once CD defined in Orange Book, part-II.
Clean File System	The file system on the media conforms to this standard.
Data track	Data tracks are tracks that are designated to contain data sectors specified in the ISO/IEC 10149.
Dirty File System	A file system that is not a clean file system.
Fixed Packet	An incremental recording method in which all packets in a given track are of a length specified in the Track Descriptor Block. Addresses presented to a CD drive are translated by the drive according to the Method 2 addressing specified in Orange Book parts-II and -III.
ICB	A control node in ISO 13346. See [4/14.6], [4/14.7], [4/14.8], and [4/14.9]
Logical Block Address	An address relative to the beginning of a partition, as defined in ISO 13346.
Media Block Address	The address of a sector as it appears on the medium, before any mapping performed by the device.
Physical Address	An address used when accessing the medium, as it would appear at the interface to the device.
Random Access File System	A file system for randomly writable media, either write once or rewritable
Sequential File System	A file system for sequentially written media (e.g. CD-R)
Session	The tracks of a volume shall be organized into one or more sessions as specified by the Orange Book part-II. A session shall be a sequence of one or more tracks, the track numbers of which form a contiguous ascending sequence.
Track	The sectors of a volume shall be organized into one or more tracks. A track shall be a sequence of sectors, the sector numbers of which form a contiguous ascending sequence. No sector shall belong to more than one track. <i>Note: There may be gaps between tracks; that is, the last sector of a track need not be adjacent to the first sector of the next track.</i>
UDF	OSTA Universal Disk Format 1.01
Variable Packet	An incremental recording method in which each packet in a given track is of a host determined length. Addresses presented to a CD drive are as specified in Method 1 addressing in Orange Book parts II and III.

VAT ICB	A File Entry ICB that describes a file containing a Virtual Allocation Table.
Virtual Address	An address described by a Virtual Allocation Table entry.
VAT	The Virtual Allocation Table (VAT) provides a Logical Block Address for each Virtual Address. The Virtual Allocation Table is used with sequential write once media.

4. Terms

May	Indicates an action or feature that is optional.
Optional	Describes a feature that may or may not be implemented. If implemented, the feature shall be implemented as described.
Shall	Indicates an action or feature that is mandatory and must be implemented to claim compliance to this standard.
Should	Indicates an action or feature that is optional, but its implementation is strongly recommended.
Reserved	A reserved field is reserved for future use and shall be set to zero. A reserved value is reserved for future use and shall not be used.

5. Description

5.1 General

Two file system strategies are defined in this document. The first is a sequential file system model. This model is designed for devices that are not randomly writable; these devices require that, in general, data be appended to previously written data. The second strategy is defined for rewritable, randomly addressable media and is the random access file system model. The unit of writing may be larger than a sector. The random access file system model also contains a description of host based defect management.

CD-R shall use the sequential file system model. CD-RW should use the random access file system model, but may use the sequential file system model. The two strategies shall not be used within one medium.

ISO 13346 requires an Anchor Volume Descriptor Pointer (AVDP [3/10.2]) at sector 256 and either n or $(n - 256)$, where n is the last recorded Physical Address on the media. A random access file system session is always contained in closed session; a sequential access file system may be in an intermediate state before the session is closed yet still be interchangeable; see 5.3.3, End of session data.

Recording shall be performed in compliance with UDF and ISO 13346. The multisession rules below shall apply for finding the volume recognition sequence and the Anchor Volume Descriptor Pointer.

ISO 9660 requires a Primary Volume Descriptor (PVD) at sector 16. If an ISO 9660 file system is desired, it may contain references to the same files as those referenced by ISO 13346 structures, or reference a different set of files, or a combination of the two.

It is assumed that early implementations will record some ISO 9660 structures but that as implementations of UDF become available, the need for ISO 9660 structures will decrease.

5.2 Information Control Block

One of the fundamental structures of ISO 13346 and UDF is the ICB (Information Control Block) [4/8.10, 4/14.6, 4/14.9]. All files and directories are described by an ICB. The ICB contains file attributes, permissions, association with an extended attribute, and the file location(s). A file can be identified by either a list of extents that contain the file or the file's data may be included in the ICB directly. To include a file in the ICB, the sum of the file size and the ICB size must be less than one sector.

ICBs are recorded with various strategies. Most of these strategies are designed for random access, write once applications. As CD-R drives are not randomly accessible for write, these strategies do not work well. Strategy type 4 is used on CD media as the VAT (see section 5.3.1) provides an apparently rewritable space. The ICB contains no indirections to ICBs to come, and is recorded as a single sector. Strategies are defined in [4/14.6.2].

5.3 Sequential File System Specific

UDF allows an intermediate state on CD-R media in which only one AVDP is recorded; this single AVDP shall be at sector 256 or sector 512 according to the multisession rules below.

Sequential file system writing shall be performed with variable packet writing. This allows maximum space efficiency for both large and small updates. Variable packet writing is more compatible with CD-ROM drives as current models do not support method 2 addressing required by fixed packets.

The Logical Volume Integrity Descriptor [3/10.10] is recorded and the volume is marked as open. This provides compatibility with versions of UDF implementations that are not aware of the extensions for CD. Logical volume integrity can be verified by finding the VAT ICB at the last recorded Physical Address. If the VAT ICB is present, the volume is clean; otherwise it is dirty.

The Partition Header descriptor [4/14.3] specifies no Unallocated Space Table [4/10.1], no Unallocated Space Bitmap [4/10.1], no Partition Integrity Table [4/11], no Freed Space Table [4/10.1], and no Freed Space Bitmap [4/10.1]. The drive is capable of reporting free space directly, eliminating the need for a separate descriptor.

ISO 13346 as specified requires a randomly addressable media, even for write once media. As this file system is written sequentially (within a track), some modification is needed for efficient application of ISO 13346.

This standard introduces no new ISO 13346 structures. The added information, which fits within ISO 13346 standard structures, describes:

1. a virtual partition or partitions
2. the contents of a special file, called the Virtual Allocation Table

The Virtual Partition does not have an immediate physical mapping. The Virtual Partition exists within the physical partition identified in the partition map. The VAT (which describes the Virtual Partition) is recorded in the physical partition identified in the partition map.

Each surface contains 0 or 1 write once partitions and 0 or 1 virtual partitions. CD media should contain 1 write once partition and 1 virtual partition. Multisurface implementations may not have virtual partitions on every surface.

The physical partition should be as large as the media, as the partition descriptor cannot be updated without starting a new session. Partitions are allowed to be not fully recorded.

5.3.1 Virtual Allocation Table

The Virtual Allocation Table (VAT) is a map that translates Virtual Addresses to logical. It is recorded as a file identified by a normal File Entry ICB (referred to hereafter as a VAT ICB) [4/14.9] which allows great flexibility in building the table. The VAT ICB is the last sector recorded in any transaction, which is easy to find on any CD-R mechanism. The VAT shall be used on CD-R media. The data of the VAT may be recorded at any location.

Each file and directory is described by a single direct ICB. The ICB is written after the associated file data to allow for data underruns during writing, which will cause logical gaps in the file data. The ICB can be written afterward and correctly identify all extents of the file data.

It is intended that file system data be written in virtual space and that file data be written in logical space. Any structure that is referenced in virtual space may be easily "overwritten" by

writing a replacement structure and updating the proper entry in the VAT to point to the replacement. At a minimum, ICBs that describe directories should be in virtual space. Placing file ICBs in virtual space allows files to be updated without updating the parent directory.

There are a couple of unique characteristics of the VAT. The file type is 0, indicated that the contents of the file are not specified by ISO 13346. The VAT ICB is found by finding the last recorded sector. The VAT is not identified by any File Identifier Descriptors, and thus its link count is zero. Error recovery schemes can find the last valid VAT by finding ICBs with file type 0 and examining the contents for the regid at the end of the table.

This file, when small, can be embedded in the ICB that describes it. If it is larger, it can be recorded in a sector or sectors preceding the ICB. The sectors do not have to be contiguous, which allows writing only new parts of the table if desired. This allows small incremental updates, even on discs with many directories. Each sector can hold entries that represent up to 512 directories.

When the VAT is small (a small number of directories on the disc), the VAT is updated by writing a new file ICB with the VAT embedded. When the VAT becomes too large to fit in the ICB, writing a single sector with the VAT and a second sector with the ICB is required. Beyond this point, more than one sector is required for the VAT. However, as multiple extents are supported, updating the VAT may consist of writing only the sector or sectors that need updating and writing the ICB with pointers to all of the pieces of the VAT.

The Virtual Allocation Table is used to redirect requests for certain information to the proper logical location. The indirection provided by this table provides the appearance of direct overwrite capability. For example, the sector describing the root directory could be referenced as virtual sector 1. Over the course of updating the disc, the root directory may change. When it changes, a new sector describing the root directory is written, and its Logical Block Address is recorded as the Logical Block Address corresponding to virtual sector 1. Nothing that references virtual sector 1 needs to change, as it still points to the most current virtual sector 1 that exists, even though it exists at a new Logical Block Address.

The use of virtual addressing allows any desired structure to become effectively rewritable. The structure is rewritable when every pointer that references it does so only by its Virtual Address. When a replacement structure is written, the virtual reference does not need to change. The proper entry in the VAT is changed to reflect the new Logical Block Address of the corresponding Virtual Address and all virtual references then point to the new structure. All structures that require updating, such as directory ICBs, shall be referenced by a Virtual Address. As each structure is updated, its corresponding entry in the VAT ICB shall be updated.

The VAT is a sequence of Uint32 entries in a file. Each entry shall be the offset, in sectors, into the physical partition in which the VAT is located (which is identified in the partition map). The first entry is for the virtual partition sector 0, the second entry for virtual partition sector 1, etc. The Uint32 entries are followed by a regid identifying the VAT and a Uint32 entry indicating the location of the previous VAT ICB.

The number of entries in the table can be determined from the VAT file size in the ICB:

$$\text{Number of entries (N)} = \frac{\text{FileSize} - 36}{4}$$

An entry of 0xffffffff indicates that the virtual sector is currently unused.

5.3.2 Virtual Partition Map

The Logical Volume Descriptor contains a list of partitions that make up a given volume. As the virtual partition cannot be described in the same manner as a physical partition, a Type 2 partition map is used.

The Logical Volume Descriptor should contain at least two partition maps. Normally, the first map is recorded as an ordinary Type 1 partition map [3/10.7.2]. The second is normally recorded as a Type 2 partition map [3/10.7.3]. Both the type 1 and type 2 maps identify the length of the individual map, the type of partition, the number of the physical partition being referenced, and the number of the volume (side) that contains the partition. The virtual map contains an identifier that can be used to determine the exact partition map type (type 2 maps are generic and may be used for other things such as a sparable partition). The presence of a virtual map also implies the presence of a VAT in the partition being referenced.

5.3.3 End of session data

A session is closed to enable reading by CD-ROM drives. The last complete session on a disc must conform to ISO 13346 which means that two AVDPs are recorded. This can be accomplished by writing data according to Table 1. Although not shown in the following example, the data may be written in multiple packets. There are only two critical pieces: the AVDP and the VAT ICB. The VAT must be the last recorded data sector (N) and the AVDP must fall at N - 256. The part in the middle can be null, contain multiple copies of the VAT, or anything the implementation sees fit.

Table 1 - End of session data

Count	Description
1	Anchor Volume Descriptor Pointer
255	Implementation specific. May contain user data, file system structures, and/or link areas.
1	VAT ICB.

The implementation specific data may contain repeated copies of the VAT and VAT ICB. Compatibility with drives that do not accurately report the location of the last sector will be enhanced.

Implementations shall ensure that enough space is available to record the end of session data.

5.4 Random Access File System Specific

CD-RW media needs to be formatted prior to use with UDF. CD-RW media is not randomly addressable at the packet level when the disc is blank; it must be pre-recorded with packets which can then be overwritten. To allow practical random addressability, fixed packets are used. UDF specifies a packet size of 32 user data blocks (64 kB).

UDF requires randomly addressable sectors. Random reads are supported by the devices. Random sector writes need to be handled by the host, usually in a device driver. This device driver (either below or part of the file system) needs to perform read-modify-write operations. The sequence of events would be to read a 32 sector packet, fill in the appropriate sectors with new data from the host, and write the 32 sectors back to the media.

The "read" cycle performs double duty. In addition to filling in data for sectors not written by the host, it can be used to detect when media is deteriorating. CD-RW devices can report a recovered error that indicates the data was read but that excessive ECC was needed to do so. This can act as an indicator to the host that the packet needs to be moved or "spared."

CD physical format does not provide for sparing or bad block mapping. Sparing is described in section 5.4.5. The host maintains a list of defects on the disc. The defect list is contained in a stream or extended attribute of the root directory (see section 5.4.6).

5.4.1 Formatting

Formatting consists of writing a lead-in, user data area, and lead-out. These areas may be written in any order. The host may direct the entire operation or use a format command to direct the drive to perform these operations. This physical format may be followed by a verification pass. Defects found during the verification pass are enumerated in the Unallocatable Areas Table (see 5.4.6). Finally, file system root structures are recorded. These mandatory file system and root structures include the Volume Recognition Sequence[2/8.3.1], Anchor Volume Descriptor Pointers[3/10.2], a Volume Descriptor Sequence[3/8.4.2], a File Set Descriptor[4/14.1], a Space Bitmap Descriptor[4/14.12], and a Root Directory [4/14.1.15].

The Anchor Volume Descriptor Pointers shall be recorded at sectors 256 and N - 256, where N is the Physical Address of the last addressable sector.

Allocation for sparing shall occur during the format process. The sparing allocation may be zero in length. Space for the sparing may be outside of the partition (in volume space) or within the partition (logical space). In either case, the spare areas are addressed in volume space. However, if the spare areas overlap the logical space, the space must be marked as used in the Space Bitmap Descriptor and listed in the unallocatable space table.

The free space descriptors shall be recorded and shall reflect space allocated to defective areas and sector sparing areas.

The format may include all available space on the medium. However, if requested by the user, a subset may be formatted to save formatting time. That smaller format may be later "grown" to the full available space.

5.4.1.1 Growing the Format

If the medium is partially formatted, it may be later grown to a larger size. This operation consists of:

1. Optionally erase the lead-in of the last session.
2. Optionally erase the lead-out of the last session.
3. Write packets beginning immediately after the last previously recorded packet.
4. Update the sparing table to reflect any new spare areas
5. Adjust the partition map as appropriate
6. Update the free space map to show new available area
7. Move the last AVDP to the new N - 256
8. Write the lead-in (which reflects the new track size)
9. Write the lead-out

5.4.2 Sparable Partition Map

CD-RW systems do not perform defect management. To provide an apparent defect-free space, a partition identified by a type 2 map is used. The partition map defines the partition number, packet size, and size and locations of the sparing tables. A partition descriptor is recorded as it normally is for a type 1 partition map.

5.4.3 Host Based Defect Management

The host shall perform defect management operations. The CD format was defined without any defect management; to be compatible with existing technology and components, the host must manage defects. There are two levels of defect management: Marking bad sectors at format time (see 5.4.6) and on-line sparing (see 5.4.2 and 5.4.5).

5.4.4 Read Modify Write Operation

CD-RW media requires large writable units as each unit incurs a 14KB overhead. The file system requires a 2KB writable unit. The difference in write sizes is handled by a read-modify-write operation by the host. An entire packet is read, the appropriate portions are modified, and the entire packet written to the CD.

Packets may not be aligned to 32 sector boundaries. Please see section 5.4.7.

5.4.5 Sparing Table

Sparing Tables point to space allocated for sparing and contains a list of mappings of defective sectors to their replacements. Separate copies of the sparing tables shall be recorded in separate packets. All sparing tables shall be kept up to date.

The sparing tables contain map entries that consist of an "original location" which is a logical address (within a partition) and a "mapped location" which is a physical address (in volume space). The volume space is a superset of the partition space. The scheme described in UDF allows spare areas to be outside of any partition if desired.

5.4.6 Unallocatable Areas

Unallocatable areas include defects and space reserved for other use such as sparing. The defect list shall be generated at format time. All space indicated by the defect list shall also be marked as allocated in the free space map. The defect list shall be recorded as a file stream (if UDF 2.0 is used) or an extended attribute (if UDF 1.50 is used) of the root directory.

5.4.7 Levels of Compliance

Levels of compliance were introduced to allow simple implementations for widespread use while providing for the flexibility required in some situations. Level 1 is the easiest to implement and contains all of the features necessary for a rich file system. Levels 2 and 3 add features that have a more limited application.

6. Multisession and Mixed Mode

The Volume Recognition Sequence [2/8.3.1] and Anchor Volume Descriptor Pointer locations are specified by ISO 13346 to be at a location relative to the beginning of the disc. The beginning of a disc shall be determined from a base address 's' for the purposes of finding the VRS and AVDP. The scheme is identical to that used for finding the PVD of an ISO 9660 multisession disc.

's' is the Physical Address of the first data sector in the first recorded data track in the last existent session of the volume. 's' is the same value currently used in multisession ISO 9660 recording. The first track in the session shall be a data track.

'n' is the physical sector number of the last recorded data sector on a disc.

If random write mode is used, the media may be formatted with zero or one audio sessions followed by exactly one writable data session containing one track. Other session configurations are possible but not described here. There shall be no more than one writable partition or session at one time, and this session shall be the last session on the disc.

6.1 Volume Recognition Sequence

The volume recognition area [2/8.3] of a UDF Bridge format shall be the part of the volume space [2/8.2] starting at sector $s + 16$.

The volume recognition space shall end in the track in which it begins. As a result of this definition, the volume recognition area [2/8.3] always exists in the last session of a disc.

6.2 Anchor Volume Descriptor Pointer

Anchor Volume Descriptor Pointers is recorded at the following logical sector numbers: $s + 256$ and $n - 256$. The AVDP at sector $n - 256$ shall be recorded before closing a session; it may not be recorded while a session is open. If the AVDP at sector $s + 256$ is not recorded, an AVDP shall be recorded at sector $s + 512$. The AVDP at 512 is to allow reserving of the first track for an ISO 9660 structure to be written later. When this track is recorded, the AVDP will be recorded in its proper location of $s + 256$.

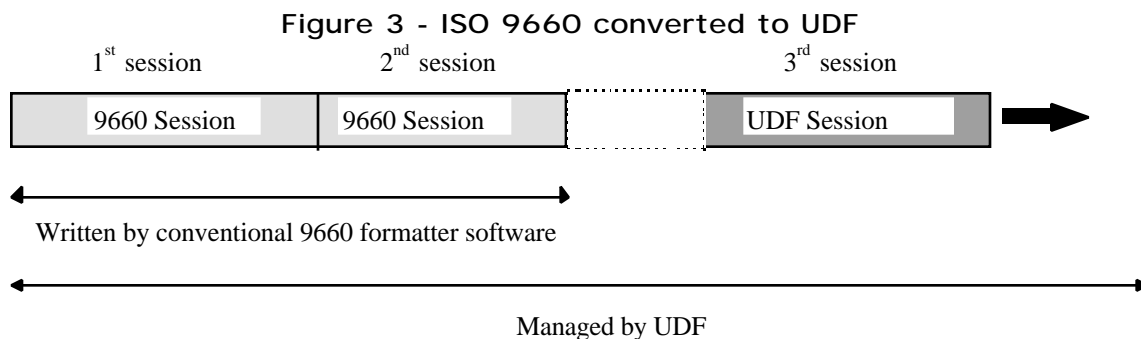

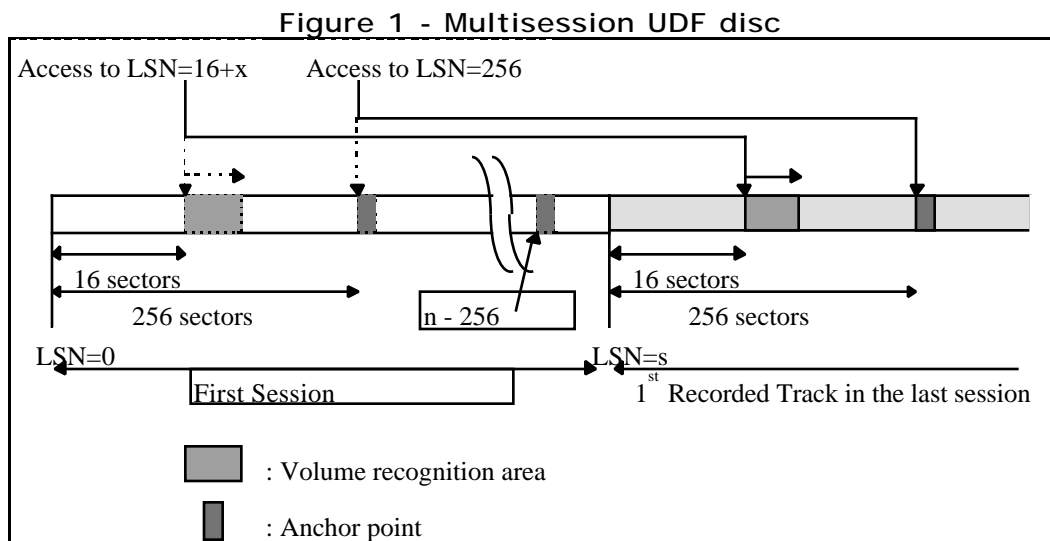
6.3 UDF Bridge format

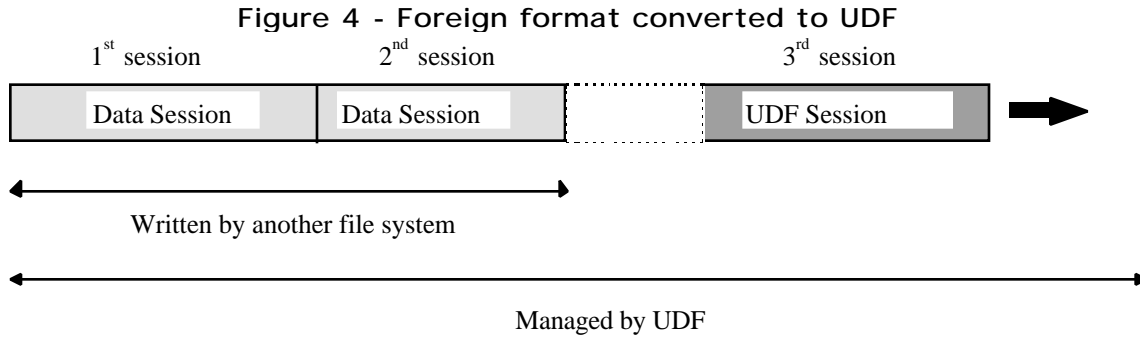
The UDF Bridge format allows UDF to be added to a disc that may contain another file system. A UDF Bridge disc shall contain a UDF file system in its last session. The last session shall follow the rules described in "Multisession and Mixed Mode" section above.

The disc may contain sessions that are based on ISO 9660, audio, vendor unique, or a combination of file systems. The UDF Bridge format allows CD enhanced discs to be created.

The UDF session may contain pointers to data in other sessions, pointers to data only within the UDF session, or a combination of both.

Some examples of UDF bridge discs are shown in Figures 1 - 4.





7. Sequential Access Implementation Strategies

7.1 Allocation of space

The UDF file system allows at least three strategies for arranging information on the CD:

1. All information is written to a single track (the invisible track)
2. A track is reserved at the beginning of the free space. All file and directory ICBs are written in the reserved track and all file data is written to the data (invisible) track.
3. Two tracks are reserved. The first track will eventually hold the ISO 9660 file system, the second track is used for file and directory ICBs, and the invisible track holds the file data.

7.1.1 Strategy 1

Strategy 1 is the most space efficient when used as the only file system. It is also the simplest. However, closing to ISO 9660 requires closing the current session and writing an additional session with ISO 9660 in it.

The initialization of the disc shall consist of writing the ISO 9660 PVD (if the 9660 system is recorded at all), the ISO 13346 AVDP and the required supporting structures. A complete ISO 9660 structure may be written. The file set shall contain a text file identifying the disc as a UDF disc. It may also contain drivers to help interchange with other systems. All writing occurs in the unreserved last track. The 9660 filesystem would not overlap the UDF file system; the primary use of 9660 would be identification of the media and possible distribution of drivers to allow the recipient to read the media.

7.1.2 Strategy 2

Strategy 2 is similar to strategy 1 except that performance gains will be seen due to the clustering of the file system information in a small physical area.

If the file system area becomes full, the file system track shall be closed. The invisible track shall be closed. A new file system track shall be reserved. Writing may then continue into the new reserved file system track and file data into the invisible track.

7.1.3 Strategy 3

Strategy 3 provides a higher degree of compatibility with ISO 9660. The first track is reserved for track at once recording, which is filled in when the session is closed. However, due to minimum track lengths, the ISO 13346 Anchor Volume Descriptor Pointer cannot be initially

recorded at sector 256. It shall be recorded at sector 512. The AVDP shall also be recorded in sector 256 when the file system is closed.

7.2 Duplicate VAT

The VAT is always recorded at the end of the invisible track. However, a VAT ICB may also be recorded at the end of a separate file system track. This second VAT ICB shall identify the VAT ICB in the data area as its parent. VAT ICBs recorded in the data area shall identify the previous VAT ICB recorded in the data area as its parent. The duplicate VAT ICB may be used to check file system integrity or aid in error recovery.

A duplicate VAT shall be recorded if strategies one or two are implemented.

7.3 Duplicate ICB

Each set of file data shall have at least one ICB that identifies it. The logical and physical location of an ICB is not specified. The ICB may be recorded in the data area, the file system area, or both. If recorded in both areas, the ICB in the file system area shall be identified by the VAT and the ICB in the data area shall not appear in the file system hierarchy. If recorded in both areas, the data area ICB shall be used solely for error recovery procedures.

8. Sequential File System Sample Sequence of Events

The following is an example of a normal sequence of events:

1. Format the disc:
 - a) [strategies 2 and 3] Reserve track(s)
 - a) [strategy 1] Write ISO 13346 volume recognition structures (which may include ISO 9660 structures)
 - b) Record the AVDP [3/10.2], PVD [3/10.1], PD [3/10.5], LVD [3/10.6], etc.
 - c) Record a root directory (which may be in the root directory ICB) and a root directory ICB.
 - d) Record a VAT (which may be contained in the VAT ICB) and a VAT ICB.
2. Write file data. This data consists of any new information, either for a completely new file or new data within an existing file.
3. Write the file ICB [4/14.9]. This ICB describes the file, including a list of all extents of the file. If the extent list is sufficiently large to overflow the sector, the continued Allocation Extent Descriptor shall be recorded before this file ICB. This ICB shall be recorded in the file system track if strategies 2 or 3 are used.
4. Write the directory data (if not embedded in the directory ICB). The directory must be updated to point to the file's new ICB (and the file version number may also be incremented). Directory data shall be recorded in the file system track if strategies 2 or 3 are used.
5. Write the directory ICB [4/14.9]. The directory ICB shall be recorded in the file system track if strategies 2 or 3 are used.
6. Repeat from step 2 if large updates are desired.
7. Write the VAT data (if not embedded in the VAT ICB). The VAT needs to reflect the new Logical Block Addresses of the new or updated directories.
8. Write the VAT ICB to the data area. If strategies 2 or 3 are used, record a VAT ICB in the file system track. The disc is now in a clean state.
9. Repeat from step 2 as necessary.

To "close" the disc to ISO 9660 compatibility (Strategies 1 and 2)

1. Close the session.
2. Generate ISO 9660 image for existing directories.
3. Write a session containing the ISO 9660 image and UDF end of session data.
4. Close this new session.

To "close" the disc to ISO 9660 compatibility (Strategy 3)

1. Generate ISO 9660 image for existing directories.
2. Write the ISO 9660 image to the first reserved track.
3. Write UDF end of session data.
4. Close the session.

For disc insertion:

1. Read sector $s + 256$. This sector shall contain the Anchor Volume Descriptor Pointer, which identifies the Volume Descriptor Sequence. If sector $s + 256$ is not recorded, read sector 512. See "Multisession and Mixed Mode" section above.

2. Read the Volume Descriptor Sequence. It will contain at least the Logical Volume Descriptor and at least one Partition Descriptor.
Cache the partition information and the File Set Descriptor location.
3. Obtain the "Next Writable Address" for the last (invisible) track.
Subtract 8 from the NWA to get the Physical Address of the last written sector.
4. Read the last written sector.
Verify that the sector contains an ICB for file type 0. Verify that the ICB describes a VAT by checking the identifier at the end of the VAT. If not a valid VAT, go to "Error Recovery."
5. Extract and cache the VAT. The VAT location will be identified by the VAT ICB.

Error Recovery:

1. Read sectors from (NWA - 8) to 257 until a VAT ICB (file type 0) is found. Verify that the ICB describes a file that has a VAT Entity Identifier at the end. This operation is more efficient if large blocks are read from the disc at one time.
2. Record a new VAT ICB at the end of the track. Advanced recovery techniques would examine all data following the found ICB for other types of ICB and recover those structures also.

9. Sequential File System Example Disc Images (strategy 1)

Table 2 - Sector Map for formatted "blank" disc

VSN	LSN	PSN	Description	General
-	-	0-15	Implementation specific, i.e. system boot area.	ISO 9660 Format Disc Init
-	-	16	ISO 9660 Primary Volume Descriptor	
-	-	17	ISO 9660 Volume Descriptor Set Terminator	
-	-	18	ISO 13346 Beginning Extended Area Descriptor	
-	-	19	ISO 13346 Volume Structure Descriptor "NSR02" [3/9.1]	
-	-	20	ISO 13346 Terminating Extended Area Descriptor	
-	-	21	Blank (may be data but must be non-descriptor)	
-	-	22-28	Link blocks (not needed if streaming can be guaranteed)	
-	-	29	ISO 9660 type L path table - describes root directory	
-	-	30	ISO 9660 type M path table - describes root directory	
-	-	31	ISO 9660 root directory - describes README.TXT and, optionally, other files, e.g. drivers	
-	-	32	Data for file "README.TXT"	
-	-	33-56	Data for file "CD_UDF.VXD"	
-	-	57-63	Link blocks (not needed if streaming can be guaranteed)	
-	-	64-95	Data for file "CD_UDF.IFS"	
-	-	96-102	Link blocks	
-	-	103	ISO 13346 Primary Volume Descriptor	ISO 13346 Format Disc Init
-	-	104	ISO 13346 Logical Volume Descriptor. This structure points to the FSD in partition 1 (the virtual partition), sector 0.	

-	-	105	ISO 13346 Partition Descriptor (partition 0). Partition spans physical sectors 257-333,000.
-	-	106	ISO 13346 Terminating Descriptor
-	-	107-113	Link blocks
-	-	114-248	Unused or more drivers
-	-	249-255	Link blocks
-	-	256	ISO 13346 Anchor Volume Descriptor Pointer (AVDP) Identifies the Volume Descriptor Sequence at sectors 103-106. This allows rewriting the VDS if an underrun occurs.
0	0	257	File Set Descriptor. Points to the root directory at partition 1 (the virtual partition), sector 1.
1	1	258	Root Directory ICB with embedded directory
-	2	259	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0 Virtual Sector 1 is identified as logical sector 1

1. VSN is the virtual Sector Number. This number is an index into the Virtual Allocation Table.
2. LSN is the logical Sector Number. This number is relative to the start of partition 0.
3. PSN is the physical Sector Number. This number is relative to the start of the disc (per Yellow/Orange Books). This is equivalent to the ISO 9660 Logical Block Number.

Links may occur at points other than indicated or may not appear at all. The volume identifier sequence and the data for each file shall be recorded contiguously.

Maximum compatibility is achieved if there are no links in sectors 0-256.

After the formatting, files can be added to the disc. To add the file FOO.TXT in the directory BAR, the disc image appears as follows. FOO.TXT is assumed to be 4 sectors long. The following example assumes no underruns occurred.

Table 3 - Sector Map for disc with added directory and file

VSN	LSN	PSN	Description	General
-----	-----	-----	-------------	---------

-	-	0-256	No change from previous table.	Old Stuff
0	0	257	File Set Descriptor. Points to the root directory at partition 1 (the virtual partition), sector 1. (no change)	
-	1	258	Root Directory ICB with embedded directory. <i>Note change: this is no longer virtual sector 1!</i> This ICB has been replaced by the one at sector 273.	
-	2	259	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0. However, as this is not the last VAT, it is ignored.	
-	3-9	260-266	Link Blocks	
-	10-13	267-270	Contents of file FOO.TXT	Add a file and directory
-	14	271	File ICB describing FOO.TXT in sectors 267-270	
2	15	272	Directory ICB for directory BAR which contains an entry for FOO.TXT. The directory information is embedded in the ICB.	
1	16	273	Directory ICB for the root directory. It contains an entry for the BAR subdirectory.	
-	17	274	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0 Virtual Sector 1 is identified as logical sector 16 Virtual Sector 2 is identified as logical sector 15	

Table 4 - Underrun while recording a file

Next, here's an example of a write in which an underrun occurs. File HP.DAT is being written, and it is 20M in size. The file is being written into the root directory.

VSN	LSN	PSN	Description	General
-	-	0-256	No change from previous table.	Old Stuff
0	0	257	File Set Descriptor. Points to the root	

			directory at partition 1 (the virtual partition) sector 1. (no change)	
-	1	258	Root Directory ICB with embedded directory.	
-	2	259	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0. However as this is not the last VAT it is ignored.	
-	3-9	260-266	Link Blocks	
-	10-13	267-270	Contents of file FOO.TXT	File and dir added last time
-	14	271	File ICB describing FOO.TXT in sectors 267-270	
2	15	272	Directory ICB for directory BAR which contains an entry for FOO.TXT. The directory information is embedded in the ICB.	
-	16	273	Directory ICB for the root directory. It contains an entry for the BAR subdirectory.	
-	17	274	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0 Virtual Sector 1 is identified as logical sector 16 Virtual Sector 2 is identified as logical sector 15	
-	18-24	275-281	Link Blocks	
-	25-8216	282-8473	Data for HP.DAT. An underrun occurs after writing the first 16M, so this is only the first part.	New file Added
-	8217-8223	8474-8480	Link Blocks	
-	8224-10271	8481-10528	Remaining 4M of HP.DAT	
-	10272	10529	The root directory. It contains an entry for the BAR subdirectory and an entry for HP.DAT. The entry for HP.DAT has two extents: 8192 sectors at logical sector 25 and 2048 sectors at logical sector 8224.	
1	10273	10530	Directory ICB for the root directory. It points to the root directory at sector	

			10272. The directory is not embedded in the ICB, which will happen when the directory gets too large to fit.
-	10274	10531	VAT ICB with embedded map. Virtual Sector 0 is identified as logical sector 0 Virtual Sector 1 is identified as logical sector 10273 Virtual Sector 2 is identified as logical sector 15