Optical Storage Technology Association

UDF 2.50 approved errata

Base Document: OSTA Universal Disk Format Specification, revision 2.50
Document: UDF 2.50 approved errata
Date: September 15, 2003; last modified July 21, 2006

Purpose of this document:

This document contains the UDF Document Change Notices (DCNs) that were approved as UDF 2.50 errata by the OSTA UDF Committee.

Important notice: UDF 2.50 rules identical to UDF 2.60 for non-POW.

On issue of the UDF 2.60 specification, some minor modifications were made to the UDF 2.50 errata, see DCN-5122. The effect of these modifications is that for non-POW media, the rules are identical for UDF 2.50 plus 2.50 errata and UDF 2.60 plus 2.60 errata. POW stands for Pseudo OverWrite. A Major benefit from this is that an implementer can use the UDF 2.60 documents to implement both UDF 2.50 and 2.60, instead of going through the UDF 2.50 spec and all 2.50 errata DCNs for implementing UDF 2.50. The only thing that a UDF 2.50 implementer needs to know when implementing UDF 2.50 from the UDF 2.60 documents is that Pseudo OverWrite and pseudo-overwritable partitions are not allowed for UDF 2.50 and that the UDF Revision field in Domain Entity Identifiers and UDF Entity Identifiers must have a value #0250 instead of #0260 (see 2.1.5.3).

Note: This assumes that non-POW DCNs that are approved as UDF 2.60 errata are also approved as UDF 2.50 errata and that POW and non-POW issues are dealt with in separate DCNs.

History of this document:

30-04-2003: Release of the approved UDF revision 2.50 document.
22-01-2004: Added DCN-5101 till DCN-5109 incl. as approved on December 8, 2003.
25-06-2004: Replaced DCN-5101 by modified version as approved on June 14, 2004.
26-08-2004: Editorial addition in description of DCN-5107 about section numbering.
20-09-2004: Replaced DCN-5104 and DCN-5108 by modified versions as approved on September 14, 2004.
12-10-2004: Added DCN-5110, because its nature is errata for all UDF revisions, see section 6.3 last paragraph.
15-11-2004: Added DCN-5113 and DCN-5115 as approved on November 4, 2004. Replaced DCN-5107 by modified version approved on November 4, 2004.
13-12-2004: Added DCN-5121 as approved on December 6, 2004.
31-01-2005: Added DCN-5112 as approved by email on January 25, 2005.
28-02-2005: Added DCN-5122 and DCN-5119. Small modifications in DCN-5102 and DCN-5119 as indicated by DCN-5122, approved on February 28, 2005. See "Importance Notice" above.
01-03-2005: Release of the approved UDF revision 2.60 document.
16-12-2005: Added DCN 5152 as approved on September 27, 2005.
11-01-2006: Added DCN 5151, 5153 and 5156 as approved on December 05, 2005. There is an extra annex document to DCN 5156.
20-07-2006: Added DCN 5155, 5157, 5159, 5162 as approved on March 02, 2006.
21-07-2006: Added DCN 5154, 5160, 5161 as approved on June 12, 2006.

OSTA Optical Storage Technology Association

UDF 2.50 approved errata

Contents:

Document: OSTA Universal Disk Format **DCN-5101**
Subject: *Virtual, metadata and read-only partitions on one volume*
Date: September 15, 2003; Modified December 8, 2003 and June 8, 2004

Description:

Outcome of UDF Committee meeting September 15, 2003:

The intent of UDF 2.50 was to exclude the combination of a virtual and a metadata partition. It is assumed that a virtual partition is on a write once partition only, but that is not really defined in the UDF spec, so an explicit exclusion is needed.
Further "single partition" must be interpreted as "single Partition Descriptor", because a volume with a Metadata Partition always has already 2 partition maps.

Added on December 8, 2003 and June 8, 2004:
A Metadata Partition must be used for the overwritable partition in the special case of an overwritable partition and a read-only partition on one volume.

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Change:

In the first paragraph of 2.2.10 replace:

This partition map *shall* be recorded for volumes which contain a single partition having an access type of 1 (read only) or 4 (overwritable). It *shall not* be recorded in all other cases.

by:

A Metadata Partition Map *shall* be recorded for volumes that contain a single Partition Descriptor having an access type of 1 (read only) or 4 (overwritable) and do not have a Virtual Partition Map recorded in the LVD.
For the special case of two non-overlapping Partitions on one volume, one with an access type of read-only and one with an access type overwritable, there shall be a Metadata Partition Map associated with the overwritable partition.
A Metadata Partition Map *shall not* be recorded in all other cases.

In 2.2.13, third paragraph of page 40:

replace:

File Entries describing any other type of file data (including streams) shall use either “immediate” allocation, or LONG_ADs which shall reference the physical or sparable partition referenced by the metadata partition, to describe the data space of the file.

by:

File Entries describing any other type of file data (including streams) shall use either “immediate” allocation, or LONG_ADs which shall reference the physical or sparable partition referenced by the metadata partition map, to describe the data space of the file. In the special two partitions case mentioned in 2.2.10, with a read-only partition and an overwritable partition on one volume, the data space of the file or stream may also be located in the read-only partition.

Document: OSTA Universal Disk Format **DCN-5102**
Subject: *No Metadata Bitmap File required for read-only partition*
Date: September 17, 2003; last modified February 28, 2005

Description:

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.
Outcome of UDF Committee meetings of September 15, 2003 and February 28, 2005:
A Metadata Bitmap File and Space Bitmap shall not be recorded for a read-only partition.

Change:

In 2.2.10 2nd paragraph replace:

- Metadata Bitmap File Location = the address of of block containing the File Entry for the metadata bitmap file. This address shall be interpreted as a logical block number within the physical or sparable partition associated with this partition map (see above "Partition Number" field).

by:

- Metadata Bitmap File Location = the address of the block containing the File Entry for the Metadata Bitmap File. This address shall be interpreted as a logical block number within the physical or sparable partition associated with this partition map (see above "Partition Number" field). If the value of the Metadata Bitmap File Location field is equal to #FFFFFFFF, no File Entry for the Metadata Bitmap File is defined.

In 2.2.13, replace:

When a Type 2 Metadata Partition map is recorded, the Metadata File, Metadata Mirror File and Metadata Bitmap File shall also be recorded and maintained.

by:

When a Type 2 Metadata Partition map is recorded, the Metadata File, Metadata Mirror File and Metadata Bitmap File shall also be recorded and maintained. The sole exception is that a Metadata Bitmap File shall not be recorded for a read-only partition.

In 2.2.13 at the bottom of page 40 replace:

Logical blocks allocated to the Metadata or Metadata Mirror Files shall be marked as allocated in the partition unallocated space bitmap, therefore a mechanism to determine available blocks within the metadata partition is needed. This is accomplished through the Metadata Bitmap file.

by:

Logical blocks allocated to the Metadata or Metadata Mirror Files shall be marked as allocated in the partition unallocated space bitmap, therefore a mechanism to determine available blocks within the metadata partition is needed. This is accomplished through the Metadata Bitmap File. A Metadata Bitmap File shall not be recorded for a read-only partition.

On page 42 replace:

- An unused block marked free in the Metadata Bitmap File.

by:

- An unused block that is available for use.

At the end of 2.3.3 Partition Header Descriptor
add a note:

NOTE 2: A Space Table or Space Bitmap shall not be recorded for a read-only partition or for a file system using a VAT.

Document: OSTA Universal Disk Format **DCN-5103**
Subject: *Equivalence for Metadata File and Metadata Mirror File*
Date: September 16, 2003

Description:

Outcome of UDF Committee meeting September 15, 2003:
This DCN describes the clarification about equivalence between Metadata File and it's Mirror File. Unused logical blocks in Metadata File and it's Mirror File do not need to be identical.

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Change:

In **2.2.13 Metadata Partition**

Replace at the first sentence in 6th paragraph on page 40:

If the *Duplicate Metadata Flag* is set in the Metadata Partition Map *Flags* field, the Metadata Mirror File shall be maintained dynamically so that it contains identical data to the Metadata File at all times.

by:

If the *Duplicate Metadata Flag* is set in the Metadata Partition Map *Flags* field, the Metadata Mirror File shall be maintained dynamically so that it contains identical contents to the Metadata File at all times. Unused logical blocks in the Metadata File and Metadata Mirror File may not be identical.

Document: OSTA Universal Disk Format **DCN-5104**
Subject: *Next extent for Metadata File and Metadata Mirror File*
Date: September 16, 2003; modified September 14, 2004.

Description:

Outcome of UDF Committee meeting September 15, 2003 and September 14, 2004:
The requirements for the Allocation Descriptors in File Entry of Metadata File and Metadata Mirror File are amended, see 2.2.13.1.

- For 2nd bullet, type "next extent of allocation descriptors" is also allowed.
- For 3rd bullet, the extent length is described in bytes instead of logical blocks and the rule is not valid for extent type "next extent of allocation descriptors".
- For 4th bullet, the rule only makes sense for extent type "recorded and allocated".

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Change:

In **2.2.13.1 Metadata File (and Metadata Mirror File)**
replace last 3 bullets of first bullet list:

- …
- Either be of type "allocated and recorded" or type "not allocated".
- Have an extent length that is an integer multiple of the *Allocation Unit Size* specified in the Metadata Partition Map.
- Have a starting logical block number which is an integer multiple of the *Alignment Unit Size* specified in the Metadata Partition Map.

by:

- …
- Not specify an extent of type "not recorded but allocated".
- Extents of type "recorded and allocated" or "not allocated" shall have an extent length that is an integer multiple of (*Allocation Unit Size* multiplied by logical block size). The *Allocation Unit Size* is specified in the Metadata Partition Map.
- Extents of type "recorded and allocated" shall have a starting logical block number that is an integer multiple of the *Alignment Unit Size* specified in the Metadata Partition Map.

Document: OSTA Universal Disk Format **DCN-5105**
Subject: *Terminating Descriptor in Metadata Partition*
Date: September 16, 2003

Description:

Outcome of UDF Committee meeting September 15, 2003:
Terminating Descriptor may be recorded within Metadata Partition as a terminator in File Set Descriptor Sequence. Further the last bullet in the 2nd bullet list of 2.2.13.1 should not refer to the Metadata Bitmap File, because it may not be present for a read-only partition (see DCN-5102).

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Change:

In **2.2.13.1 Metadata File (and Metadata Mirror File)**

Replace:

The Allocation Descriptors for this file shall describe only logical blocks which contain one of the below data types. No user data or other metadata may be referenced.

- FSD
- ICB
- Extent of Allocation Descriptors (see 2.3.11).
- Directory or stream directory data (i.e. FIDs)
- An unused block marked free in the Metadata Bitmap File.

By to add Terminating Descriptor:

The Allocation Descriptors for this file shall describe only logical blocks which contain one of the below data types. No user data or other metadata may be referenced.

- FSD
- Terminating Descriptor
- ICB
- Extent of Allocation Descriptors (see 2.3.11).
- Directory or stream directory data (i.e. FIDs)
- An unused block that is available for use.

Document: OSTA Universal Disk Format **DCN-5106**
Subject: *Metadata Mirror File FEs and AEDs always far apart*
Date: September 26, 2003

Description:

Outcome of UDF Committee meeting September 15, 2003:
Not only the data for the Metadata File and its mirror must be located far apart, but also their File Entries and possible Allocation Extent Descriptors. The latter does NOT depend on the value of the Duplicate Metadata Flag.
This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Change:

In the NOTE after the second dotted list of 2.2.13.1 replace:

NOTE: In the case where the *Duplicate Metadata Flag* in the Metadata Partition Map is set, the allocations for the Metadata File and Metadata Mirror File should be as far apart (physically) as possible. Typically this is achieved by maximizing the difference between the start LBNs of the extents belonging to the file and its mirror. Likewise the file entries for these two files should be recorded as far apart as possible. Some drive/media …

by:

NOTE: The File Entry and possible Allocation Extent Descriptors of the Metadata File should be recorded as far apart (physically) as possible from those of the Metadata Mirror File. The same counts for the allocated extents of these two files in the case that the *Duplicate Metadata Flag* in the Metadata Partition Map is set. Typically, recording far apart is achieved by maximizing the difference between the start LBNs of the descriptors and extents belonging to the file and its mirror. Some drive/media …

Document: OSTA Universal Disk Format **DCN-5107**
Subject: *Clarify overlapping of Sparing Table with a partition*
Date: September 25, 2003; modified November 4, 2004

Description:

This DCN is for the next UDF revision after 2.50 **and for the errata of UDF 1.50 thru 2.50**.
Sparing Area may overlap with a partition, but it is not clearly defined in UDF 1.50 thru UDF 2.50 whether instances of the Sparing Table may overlap with a partition or not.
History of this DCN:
In the UDF committee meeting, September 15, 2003 it was decided not to allow overlap of the Sparing Table with a partition and a DCN-5107 with that content and title *"Sparing Table itself shall not overlap with a partition"* was approved on December 8, 2003. However, it appeared that existing implementations were broken in this way. Therefore, it was decided in the meeting, November 4, 2004 to reopen DCN-5107 and reverse the previous decision and clearly define that overlap is allowed and how shall be acted in that case with respect to the Non-Allocated Space Stream. The DCN title was changed accordingly.
For UDF 1.50, Non-Allocated Space List must be read instead of Non-Allocatable Space Stream. For UDF 1.50 thru UDF 2.01, section 2.2.11 must be read instead of 2.2.12.

Change:

In 2.2.12, in the 3rd paragraph of page 38 replace:

Available sparing areas may be anywhere on the media, either inside or outside of a partition. If located inside a partition, sparable space shall be marked as allocated and shall be included in the Non-Allocatable Space Stream.

by:

Available sparing areas and instances of the Sparing Table may be anywhere on the media, either inside or outside of a partition. If overlapping with a partition, the overlapping part shall be marked as allocated and shall be included in the Non-Allocatable Space Stream.

Document: **OSTA Universal Disk Format** **DCN-5108**
Subject: *Descriptor CRC Length Uint16 overflow rules*
Date: October 04, 2003; modified September 14, 2004

Description:

This DCN is for the next UDF revision after 2.50 **and for the errata of the UDF revisions 1.02 thru 2.50**.
This issue was discussed in the UDF Committee meeting September 15, 2003.
Mind that section numbers may differ for previous revisions.

Since the definition of UDF, until now nobody seems to have been aware of the fact that the value of the Descriptor CRC Length can easily overflow for some descriptors. The default value for the CRC Length is (length of descriptor – 16), so if a descriptor length grows beyond 65551 (= 16 + MAX_UINT16) bytes, the CRC length does no longer fit in the Uint16 Descriptor CRC Length Field. Candidates for this are the descriptors that have a 'no max' length in the table of UDF section 5.1. However, the Sparing Table is missing in this table, and also has a "no max" length. This omission will also be repaired in this DCN. The 'no max' candidates are:
Logical Volume Descriptor (LVD), Logical Volume Integrity Descriptor (LVID), Unallocated Space Descriptor(USD), Space Bitmap Descriptor (SBD) and Sparing Table (ST).

The descriptors that practically can grow above a length of 65551 are the SBD and ST because:

- For the LVD, 1017 Type 2 partition maps can be recorded before the LVD length grows beyond 65551 bytes. This could only occur for a multiple volume Volume Set. Such a Volume Set has the restriction of one partition map for each volume.
- For the LVID, in the case of 8000 partition maps, still 1471 bytes are left for the Implementation Use field. UDF itself only needs 46 bytes for Implementation Use.
- For the USD, there is room for 8190 extent_ad's, which should be enough for the registration of unallocated space outside the partitions.

The ECMA intent of the CRC is to detect descriptor damage. This is a valuable mechanism. For the Space Bitmap Descriptor (SBD), a CRC length of (descriptor length – 16) or 0 is allowed, because the bitmap can be updated very frequently, and it would be cumbersome to calculate the CRC over the whole descriptor for each bitmap change. For the Sparing Table (ST), only (descriptor length – 16) is allowed for the CRC length. The ST however is not updated frequently and a CRC calculation over the Map Entries is very valuable, e.g. for distinguishing between damaged an undamaged instances of the Sparing Table, in the case that a block (not the first one) of a ST instance is accidentally overwritten. Further the most important Map Entries of the ST are in the beginning,

because the Map Entries are sorted on the 'Original Location' field, moving the unused and defective Map Entries to the end of the descriptor.
Because of these considerations and in order to keep the rules as simple as possible, no extra exceptions for the CRC Length rule of the ST are introduced, but the default rule specified in 2.2.1.2 and 2.3.1.2 are adapted in order to cope with the 16 bits overflow case. An extra benefit is that also the theoretical cases (LVD, LVID and USD) are covered by this rule.

Changes:

*In the table **5.1 Descriptor Lengths** add a row with:*

Sparing Table	no max

*In **2. Basic Restrictions & Requirements** replace:*

Descriptor CRCs CRCs shall be supported and calculated for all Descriptors except for the Space Bitmap Descriptor. There is a CRC length special case for the Allocation Extent Descriptor.

by:

Descriptor CRCs CRCs shall be supported and calculated for all Descriptors. There are exception rules for the Descriptor CRC Length of the Space Bitmap Descriptor and the Allocation Extent Descriptor.

*In **2. Basic Restrictions & Requirements** remove row:*

Space Bitmap Descriptor CRC not required.

*Replace the complete text of section **2.2.1.2** by:*

CRCs shall be supported and calculated for each descriptor. Unless otherwise specified, the value of the Descriptor CRC Length field shall be set to the minimum of the following two values: ((Size of the Descriptor) - (Length of Descriptor Tag)); 65535. When reading a descriptor, the Descriptor CRC should be validated.

NOTE 1: The *DescriptorCRCLength* field must not be used to determine the actual length of the descriptor or the number of bytes to be read. These lengths do not match in all cases because of possible *DescriptorCRCLength* truncation to 65535 and other *DescriptorCRCLength* exceptions as specified in this standard.

*Replace the complete text of section **2.3.1.2** by:*

The same applies as for volume structure *DescriptorCRCLength* values, see 2.2.1.2.

In 2.2.13.2 on page 43, first bullet, replace:

- The descriptor tag fields *DescriptorCRC* and *DescriptorCRCLength* for this SBD shall be set to zero.

by:

- The descriptor tag *DescriptorCRCLength* field for this SBD shall be set to zero or 8. The value of 8 is recommended.

and in 2.3.8.1 replace:

2.3.8.1 struct Tag DescriptorTag
The calculation and maintenance of the *DescriptorCRC* field of the Descriptor Tag for the *SpaceBitmap* descriptor is optional. If the CRC is not maintained then both the *DescriptorCRC* and *DescriptorCRCLength* fields shall be ZERO.

by:

2.3.8.1 struct Tag DescriptorTag
There are exception rules for the SBD *DescriptorCRCLength.* If the default value for the *DescriptorCRCLength* as defined by 2.3.1.2 is not used, then *DescriptorCRCLength* shall be either 8 or zero. The value of 8 is recommended.

Document: **OSTA Universal Disk Format** **DCN-5109**
Subject: *Clarification of NOTE on page 41*
Date: October 02, 2003

Description:

This issue was discussed in the UDF Committee meeting of September 15, 2003.
This DCN is meant for the next UDF revision after 2.50 **and for the errata of UDF 2.50**.

The NOTE of 2.2.13 below the figure on page 41 can easily be misunderstood.

Changes:

*In **2.2.13 Metadata Partition** below the figure on page 41 replace:*

NOTE: the LBN values used in the diagram above are for illustrative purposes only and are not fixed. The partition references used are fixed as a consequence of the Metadata Partition implementation.

by:

NOTE: The LBN values used in the diagram above are for illustrative purposes only and are not fixed. The partition reference numbers used are determined by the order of the partition maps in the LVD.

Document: OSTA Universal Disk Format **DCN-5110**
Subject: *Appoint OS Identifier for UNIX - NetBSD*
Date: August 16, 2004

Description:

Request from NetBSD community. This DCN is meant for the next UDF revision after 2.50 **and for the errata of the UDF revisions 1.02 thru 2.50**.

Change:

In 2nd table of 6.3 after:

4	7	UNIX - FreeBSD

add a row:

4	8	UNIX - NetBSD

Document:	**OSTA Universal Disk Format**	**DCN-5112**
Subject:	BD non-POW media *recommendations for UDF 2.50*	
Date:	November 4, 2004; modified December 15, 2004	

Description:

This DCN is for the next UDF revision after 2.50 (UDF 2.60) **and for the UDF 2.50 errata**.
The Blu-ray Disc (BD) Format for consumer appliances uses UDF 2.50 as the file system for the Blu-ray Read-Only disc (BD-ROM) and Blu-ray Rewritable disc (BD-RE). The Blu-ray Recordable disc (BD-R) using SRM without Logical OverWrite (non-POW), also uses UDF 2.50. BD-R using SRM with Logical OverWrite (POW) uses UDF 2.60, rather than UDF 2.50. POW stands for Pseudo OverWrite.
The purpose of this proposal is to provide enough information for the requirements and restrictions in the Blu-ray Disc Format, and to support good interchangeability between both computer systems and consumer appliances using Blu-ray Disc.
The text in this DCN, is meant for the UDF 2.50 errata and describes the requirements and recommendations for BD-R non-POW media, so all the BD media that use UDF 2.50.
The text in this DCN is also copied to the UDF 2.60 specification, except for the text in section 6.y.3. The text of 6.y.3 for the UDF 2.60 specification is in a separate DCN-5114.

Change:

Insert a new section 6.y to describe the recommendations for Blu-ray Disc:

6.y Recommendations for Blu-ray Disc media

This appendix defines the requirements and recommendation on volume and file structures for Blu-ray Disc (BD) media, to support data interchange among computer systems and consumer appliances. These requirements do not apply to the discs when the use of the discs is limited to computer systems and there is no necessity to provide interchangeability with consumer appliances. Specific requirements related to BDAV and BDMV application usage are described in section 6.y.4.

Blu-ray Disc has the following three types of media:

- Blu-ray Disc Read-Only Format (BD-ROM)
- Blu-ray Disc Rewritable Format (BD-RE)
- Blu-ray Disc Recordable Format (BD-R)

BD-R can use either SRM with LOW or SRM without LOW, for details see section 6.y.3. BD-ROM, BD-RE and BD-R using SRM without LOW, all use UDF revision 2.50. BD-R using SRM with LOW uses UDF revision 2.60, rather than 2.50.

Common characteristics and requirements for these three media types are:

1. Logical sector size is 2048 bytes.
2. ECC Block Size is 65536 bytes (64KB)
3. Sparable Partition Map and Sparing Table shall not be recorded.
4. Non-Allocatable Space Stream shall not be recorded.

6.y.1 Requirements for Blu-ray Disc Read-Only Format (BD-ROM)

A Blu-ray Read-Only disc (BD-ROM) is a Read-Only medium. The BD-ROM File System Format shall comply with UDF revision 2.50 and has the following additional requirements:

For Volume Structure:

1. The Partition Descriptor Access Type shall be 1 (read-only).
2. Three Anchor Volume Descriptor Pointers should be recorded.

For File Structure:

3. Unallocated Space Table and Unallocated Space Bitmap shall not be recorded.
4. Metadata Bitmap File shall not be recorded.

NOTE: Duplication of Metadata File data is optional. When robustness is required, it is recommended that duplication is used and that Metadata File and Metadata Mirror File data and descriptors are recorded at the physically inner radius area and outer radius area, respectively.

6.y.2 Requirements for Blu-ray Disc Rewritable Format (BD-RE)

A Blu-ray Rewritable disc (BD-RE) is a non-sequential recording medium. A BD-RE drive performs read modify write operations when needed. Defect free logical space is provided by a BD-RE drive which performs defect management using the linear replacement algorithm.
The BD-RE File System Format shall comply with UDF revision 2.50 and has the following additional requirements:

For Volume Structure:

1. The Partition Descriptor Access Type shall be 4 (overwritable).

For File Structure:

2. An Unallocated Space Bitmap shall be recorded, no Unallocated Space Table.

NOTE: Duplication of Metadata File data is optional. When the user requires robustness rather than write performance, it is recommended that duplication is used and that Metadata File and Metadata Mirror File data and descriptors are recorded at the physically inner radius area and outer radius area, respectively.

Requirements for Defect Management:
Spare Area shall be assigned on a Blu-ray Rewritable disc, as the UDF file system requires Drive Defect Management by the drive system. In general, Spare Areas with the default size are assigned at format time.

NOTE: When the available clusters in Spare Area are exhausted, additional Spare Area can be allocated after all data is backed up to the other media. On the other hand, if a special utility tool can move some file data and volume structure on the disc in order to shorten the volume space, the Spare Area can be expanded preserving the file data on the disc.

6.y.3 Requirements for Blu-ray Disc Recordable Format (BD-R)

{note: this section is for UDF 2.50 errata only, for UDF 2.60 text, see DCN-5114}

A Blu-ray Recordable disc (BD-R) is a Write-Once medium that can use Sequential Recording Mode (SRM) either with or without Logical OverWrite (LOW). Drive based defect management using the linear replacement algorithm is supported.

For the BD-R File System Format for media using SRM with LOW (Pseudo OverWrite), UDF 2.60 is used. For further details, see the UDF 2.60 specification.

The BD-R File System Format for media using SRM without LOW shall comply with UDF revision 2.50 and has the following additional requirements:

For Volume Structure:

1. The Partition Descriptor Access Type shall be 1 (read-only) or 2 (write-once).

For File Structure:

2. Unallocated Space Table and Unallocated Space Bitmap shall not be recorded.
3. Only ICB Strategy Type 4 shall be used.

6.y.4 Information about AV Applications

The Blu-ray Disc Format has two types of AV Application Formats that are called “BDAV Application” and “BDMV Application”.

Information about BDAV Application Use
The ”BDAV Application” is a Video Recording Format for BD-RE discs and BD-R discs, including AV Stream and database for playback the AV Stream.
The “BDAV”, “BDAV1”, “BDAV2”, “BDAV3”, and “BDAV4” directories immediately under the root directory are reserved for the BDAV application.

Information about BDMV Application Use
The ”BDMV Application” is a Video Application Format for BD-ROM discs, including AV Stream and database for playback the AV Stream.
The “BDMV” directory immediately under the root directory is reserved for the BDMV application.

6.y.4.1 Requirements for BDAV and BDMV Application usage

The following additional requirements are applied for BDAV and BDMV Application usage:

1. A volume set shall consist of only one volume.
2. Only one prevailing Partition Descriptor shall be recorded in the Volume Descriptor Sequence.
3. A Metadata Partition Map shall be recorded.
4. Symbolic Links shall not be used for all files and directories (the value of the File Type field in the ICB shall not be 12).
5. Hard Link shall not be used for all files and directories.
6. Multisession and VAT recording shall not be used.

Document: **OSTA Universal Disk Format** **DCN-5113**
Subject: *Main and Reserve VDS far apart*
Date: October 12, 2004; November 4, 2004

Description:

Several measures have been taken to increase UDF robustness. Therefore, it is strange that there are no stricter rules for the position of the Main and Reserve Volume Descriptor Sequences. This DCN also avoids mentioning these rules in the recommendations for each new media type.
This DCN is meant for the next UDF revision after 2.50 and for the UDF 2.50 errata.

Change:

After section 2.2.3.2 add a note:

NOTE: The Main VDS extent and the Reserve VDS extent shall be recorded in different ECC blocks. The locations of these extents on the volume should be as far apart as physically possible. Typically this is achieved by maximizing the difference between the start LSNs of the extents. Care should be taken in case of special LSN address schemes, e.g. for multiple layer media.

Document: **OSTA Universal Disk Format** **DCN-5115**
Subject: *Enable UDF 2.50 POW read compatibility*
Date: October 14, 2004; modified: November 12 (editorial)

Description:

This DCN is meant for the next UDF revision after UDF 2.50 **and as a UDF 2.50 errata**. This DCN enables read compatibility for partitions with an Access Type that does not match with the media type and for partitions with an unknown Access Type value. The main reason to define this now and as a UDF 2.50 errata is to enable read-only access by a UDF 2.50 implementation for media with a UDF 2.60 File System using a Pseudo-Overwrite partition an a Minimum UDF Read Revision value of 2.50 (e.g. BD-R using POW).

Changes:

At the end of 2.2.14.2, add:

If the value of Access Type is not equal to any of the defined Access Type values or if the combination of the medium and drive is not capable of performing the write action denoted by the Access Type value, the partition shall be handled as a read-only partition (e.g. an overwritable partition on a write-once medium or in a read-only drive).

NOTE: The above rule is important in order to enable read-only access by a UDF 2.50 implementation for media with a higher UDF revision (e.g. UDF 2.60) using a Pseudo-Overwrite partition and a Minimum UDF Read Revision value of 2.50.

Document: OSTA Universal Disk Format **DCN-5119**
Subject: *Zero Information Length for Non-Allocatable Space Stream*
Date: October 26, 2004; last update February 28, 2005

Description:

This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**. There is no requirement for the Non-Allocatable Space Stream to have an Information Length value of zero as for the Non-Allocatable Space List in UDF 1.50 and lower. It is absurd to assume that there is any relevant data in this stream.

Change:

In 3.3.7.2 at the top of page 90 replace:

… The allocation descriptors shall have allocation type 1 (allocated but not recorded). …

by:

… The allocation descriptors shall have allocation type 1 (allocated but not recorded). The Information Length in the File Entry of this stream shall be zero; so all allocation descriptors are in the file tail. …

Document:	**OSTA Universal Disk Format**	**DCN-5121**
Subject:	*Clarification of Directory bit in parent FID*	
Date:	December 03, 2004; Modified December 06, 2004	

Description:

This DCN is meant for the next UDF revision after 2.50 **and for the errata of UDF revisions 2.00, 2.01 and 2.50**.
In ECMA 4/14.4.3, below the caption of figure 13 it is stated that:

"If the Parent bit is set to ONE, then the Directory bit shall be set to ONE."

This was true for ECMA 2nd edition, were there were no streams. However for ECMA 3rd edition a parent FID in a stream directory can point to a file or directory to which the stream directory is attached or to itself in the case of the System Stream Directory. In the case that a stream directory is attached to a file, the directory bit in the parent FID of the stream directory shall not be set. This is current practice but not described in the UDF spec.

Change:

*At the begin of current section **2.3.4.2 Uint8 FileCharacteristics** insert:*

2.3.4.2.1 Deleted bit

*At the end of current section **2.3.4.2 Uint8 FileCharacteristics** add:*

2.3.4.2.2 Parent bit and Directory bit

There is a flaw in the following statement in ECMA 4/14.4.3, below figure 13:

"If the Parent bit is set to ONE, then the Directory bit shall be set to ONE."

In spite of this statement, the Directory bit in a parent FID shall only be set to ONE if the FID identifies a directory or the System Stream Directory. If the parent FID identifies a file, the Directory bit shall be set to ZERO. The latter is the case for a parent FID in a Stream Directory that is attached to a file.

Document: **OSTA Universal Disk Format** **DCN-5122**
Subject: *Make UDF2.50 identical to UDF 2.60 for non-POW*
Date: November 12, 2004; modified February 22, 2005;

Description:

This DCN has consequences for the **UDF 2.50 errata only**, so no change for UDF 2.60. It only proposes a few small modifications in existing DCNs. These modifications only have an effect for the UDF 2.50 errata.

The goal is that UDF 2.50 together with the UDF 2.50 errata becomes 'guaranteed identical' to UDF 2.60 for the non-Pseudo OverWrite (non-POW) case.
This can be achieved by a small modification in DCN-5102 and addition of DCN-5119 to the UDF 2.50 errata. These modifications are explained here. After these modifications are excecuted, this DCN-5122 can effectively disappear, but it is maintained in the set of DCNs and UDF 2.50 errata as clarification.
None of these modifications have real implementation consequences for an existing UDF implementation.

A Major benefit from this is that an implementer can use a single specification document to implement both UDF 2.50 and 2.60, namely the UDF 2.60 spec, instead of going through the 2.50 spec and all errata DCNs for implementing UDF 2.50. So UDF 2.50 and 2.60 can be implemented in one action.
The only thing that a UDF 2.50 implementer needs to know is that Pseudo OverWrite and pseudo-overwritable partitions are not allowed in UDF 2.50 and that the UDF Revision field in Domain Entity Identifiers and UDF Entity Identifiers must have a value #0250 instead of #0260 (see 2.1.5.3). This will be explained in the heading of the UDF 2.50 errata document, so a UDF 2.50 implementer can decide for himself whether to use the UDF 2.50 + errata or the UDF 2.60 document.

What is needed to make (UDF 2.50 + errata) 'guaranteed identical' to (UDF 2.60 for non-POW media)?

If we would add all DCNs to the UDF 2.50 errata, we would have proved that 2.50 plus 2.50 errata is identical to 2.60.
DCNs that cannot be added to the 2.50 errata are the POW related ones, so 5111, 5114 and 5116. There is only one non-POW related issue that is in DCN-5116. It is dealt with below. Further, DCN-5120 is valid for UDF 2.60 only, so it does not make sense as UDF 2.50 errata.

Of the remaining DCNs, only 5100, 5117, 5118 and 5119 are not yet part of the UDF 2.50 errata.

5100: editorial DCN
5117: Common aspects of recording for different media (2.01-2.60 in fact, editorial)
5118: Clarify location of Partition Header Descriptor (1.02-2.60 in fact)
5119: Zero Information Length for Non-Allocatable Space Stream (2.00-2.60 in fact)

The DCNs 5100, 5117 and 5118 are editorial or clarification only and it does not make sense to add them to the UDF 2.50 errata document.

So the only DCN that has to be added to the UDF 2.50 errata is DCN-5119.

Changes: {in DCNs only}

In DCN-5119
replace: This DCN is meant for the next UDF revision after 2.50.
by: This DCN is meant for the next UDF revision after 2.50 **and for the UDF 2.50 errata**.

Remains the non-POW related issue in DCN-5116 that affects non-POW media.
It is the bitmap not required (2.50) versus bitmap shall not be recorded (UDF 2.60) case for read-only partitions. This also has no implementation consequences and can be added to the UDF 2.50 errata by three small changes to DCN-5102:

in DCN-5102 replace: … no Metadata Bitmap File is required for a read-only partition
by: … a Metadata Bitmap File shall not be recorded for a read-only partition

in DCN-5102 replace: A Metadata Bitmap File is not required for a read-only partition.
by: A Metadata Bitmap File shall not be recorded for a read-only partition.

add to DCN-5102:

At the end of 2.3.3 Partition Header Descriptor
add a note:

NOTE 2: A Space Table or Space Bitmap shall not be recorded for a read-only partition or for a file system using a VAT.

Document: **OSTA Universal Disk Format** **DCN-5151**
Subject: *Recommendations DVD-R DL LJR*
Date: June 15, 2005; last modified: December 05, 2005
Status: Approved December 05, 2005

Description:

This DCN is for the next UDF revision after 2.60 **and as errata for all previous UDF revisions**.
DVD-R DL LJR introduces a new method of recording named Layer Jump Recording (LJR) as described in the MMC and Mt Fuji specifications. Although similar to incremental recording, this new recording is slightly different. Reserved R-Zones and LJBs (Layer Jump Block) of DVD-R DL LJR do not match the definition of a single UDF track, but two logical tracks. Consequently, Border does not match the UDF session definition. LJR also introduces the possibility to remap anchor point sectors.
UDF multi-session is not straightforward on DVD-R DL LJR, so this DCN describes how to perform multi-Border / multi-session recording on DVD-R DL LJR

Change:

Add:

6.xx Recommendations for DVD-R DL LJR (Multi-Border recording)

This appendix defines the recommendations on volume and file structures for DVD-R DL LJR, to support the interchange of information between users of computer systems.

1. The volume and file structure should comply with UDF 2.00.
2. The Minimum UDF Read Revision and Minimum UDF Write Revision should be 2.00.
3. The length of logical sector and logical block shall be 2048 bytes.

Additionally, the following recommendations are made for DVD-R DL LJR:

The DVD-R DL LJR does not follow the usual session rules. On DVD-R DL LJR the start of each Border corresponds with the start of a new session, as usual. However, the end of each session is always the end of the disc. This results in overlapping sessions, which is not strictly according to the session definition in 1.3.2.
DVD-R DL LJR is a fixed size medium. Each R-Zone contains one or more LJB (Layer Jump Block). For each R-Zone, READ TRACK INFORMATION returns 1 (physical) track but UDF implementations need to consider it as two logical tracks per LJB: one on layer 0, one on layer 1. Boundary of the logical track containing the current NWA for the R-Zone is indicated by the Next Layer Jump Address.

The formula to calculate the start address of the second logical track (on L1) can be found in the Mt. Fuji specification.
Files may start in a track of layer 0, respectively 1, and continue in a track of layer 1, respectively 0, so UDF implementations should take care to write corresponding file extents.

For DVD-ROM drive compatibility, UDF implementation should close the Border.

6.xx.0 DVD-R DL LJR Differences
DVD-R DL LJR with remapping slightly differs from recommendations in 6.11

Differences with 6.11.3 Multi-session Usage:

- After the first session, at least 2 of the AVDPs at the logical sector numbers 256, *N*-256 and *N* and at least the AVDPs remapped in the previous session, are *remapped* from the last session. The remapping requires writing in the last session AVDPs with location tags of 256, *N*-256 and/or *N*, then instruct the drive to remap with the Remapping Address (Format Code = 24h) of SEND DISC STRUCTURE command, using Anchor Point Number 2, 3 and/or 4 for respectively 256, *N*-256 and/or *N*.
- After the first session, at least 2 of the AVDPs at logical sector numbers S+256, C-256 and C are written, where C is the LRA of the last session.

Document: **OSTA Universal Disk Format** **DCN-5152**
Subject: *Stream bit ZERO for main data stream*
Date: September 8, 2005; last modified September 27, 2005
Status: Approved September 27, 2005

Description:

This DCN is meant for the next UDF revision after 2.60 **and as errata for UDF 2.00 till UDF 2.60 inclusive**.
There is confusion whether the ICBTag Flags Stream Bit of an Extended File Entry must be set to ONE if a Stream Directory is attached, because this EFE is referenced by the Parent FID in the Stream Directory. The confusion is raised by the unfortunate text of Note 24 in ECMA 4/14.6.8 bit 13. The Stream Bit is meant to distinguish between the main data stream and named data streams as defined by ECMA 4/8.8.3. E.g. if a repair utility finds a File Entry or Extended File Entry with the Stream Bit set, it knows that it must search for a reference in a Stream Directory instead of a normal directory. Note 24 of ECMA 4/14.6.8 in fact aims at a different situation, i.e. a hard link between a named data stream of a file and the main data stream of another file. This type of hard link is not allowed by any UDF revision.

Change:

In 2.3.5.4 replace: **NOTE:**

by: **NOTE 1:**

and at the end of 2.3.5.4 add:

Bit 13 (*Stream*):

- Shall indicate (ONE) whether a File Entry or Extended File Entry defines a Named Stream or the main data stream of a file or directory, see ECMA 4/8.8.3 and UDF 3.3.5.
- Shall be set to ONE for a FE or EFE defining a Named Stream. It shall be set to ZERO in all other cases.

NOTE 2: The Stream bit shall be set to ZERO for the FE or EFE of the main data stream of a file or directory and for the FE or EFE of the System Stream Directory. This is so in spite of the fact that such a FE or EFE may be referenced by the Parent FID in a Stream Directory, thus excluding the parent FID case from Note 24 in ECMA 4/14.6.8.

Document: **OSTA Universal Disk Format** **DCN-5153**
Subject: *Relaxation of file timestamps relation rule*
Date: November 04, 2005
Status: Approved December 05, 2005

Description:

This DCN is for the next UDF revision after 2.60 and as errata for all previous UDF revisions 1.02 thru 2.60.
Because of a different definition of the file creation time in different Operating Systems, it is difficult for UDF implementations to always ensure that the Modification, Access and Attribute Date and Times "shall not be earlier than the File Creation Date and Time", as required by ECMA. Therefore these rules will be changed from mandatory to a recommendation as decided in the September 27, 2005 UDF committee meeting.
Editorial: "Time" replaced by "Date and Time" to be consistent with ECMA. This will be changed for the whole UDF spec, see the editorial DCN-5150.

Change:

In 2.3.6 replace:

struct timestamp AccessTime;
struct timestamp ModificationTime;
struct timestamp AttributeTime;

by:

struct timestamp **AccessDateAndTime;**
struct timestamp **ModificationDateAndTime;**
struct timestamp **AttributeDateAndTime;**

after section 2.3.6.8 add:

2.3.6.9 Access, Modification, Creation and Attribute Timestamps

ECMA sections 4/14.9.12-14 state that the Access, Modification and Attribute Date and Time "shall not be earlier than the File Creation Date and Time …". Because some Operation Systems have a different notion of "Creation Time", UDF changes this ECMA rule from mandatory into a recommendation by reading "*should* not be earlier" instead of "*shall* not be earlier" in ECMA 4/14.9.12-14.

NOTE: ECMA 4/14.9.12-14 only refers to the File Creation Date and Time in a File Times Extended Attribute. However, the File Times EA File Creation Date and Time shall not be recorded for an Extended File Entry. An EFE has its own Creation Date and Time field that shall be used, see 3.3.4.3.1 and ECMA 4/14.17.13-16.

Document:	**OSTA Universal Disk Format**	**DCN-5154**
Subject:	*Requirements for HD DVD Disc*	
Date:	December 5, 2005; last modified June 09, 2006	
Status:	Approved June 12, 2006	

Description:

This DCN is for the next UDF revision after 2.60 **and for the UDF 2.50 and UDF 2.60 errata**.

The High Density DVD (HD DVD) Format for consumer appliances uses UDF 2.50 as the file system for the High Density Read-Only disc (HD DVD-ROM), High Density Rewritable disc (HD DVD-RAM) and The High Density Recordable disc (HD DVD-R for SL/DL).

The purpose of this proposal is to provide enough information for the requirements in the HD DVD Format, and to support good interchangeability between both computer systems and consumer appliances using HD DVD.

The text in this DCN describes the requirements for HD DVD media, so all the HD DVD media that use UDF 2.50.

Change:

Insert a new section 6.z to describe the requirements for HD DVD Disc:

6.z Requirements for HD DVD Disc

This appendix defines the requirements and restrictions on volume and file structures for HD DVD media, including but not limited to HD DVD-ROM discs (6.z.1), HD DVD-RAM discs (6.z.2) and HD DVD-R for SL/DL discs (6.z.3), to support the interchange of information between users of both computer systems and consumer appliances. These requirements do not apply to the discs that are used in a computer system environment only and have no interchangeability with consumer appliances. The common requirements for these HD DVD discs are summarized as follows:

1. The volume and file structure shall comply with UDF 2.50.
2. The length of logical sector and logical block shall be 2048 bytes.
3. ECC block size is 32 sectors (64 KB).
4. A Main Volume Descriptor Sequence and a Reserve Volume Descriptor Sequence shall be recorded.
5. A HD DVD disc shall have a single volume with a single Partition Descriptor per side.
   Therefore, the volume sequence number shall be 1, the maximum volume sequence number shall be 1 and the Primary Volume Descriptor Interchange Level shall be 2.
6. Only ICB Strategy type 4 shall be used.

6.z.1 Requirements for HD DVD-ROM

The volume and file structure is simplified as for Read-Only discs.

For Volume Structure:

1. A partition on a HD DVD-ROM disc shall be a read-only partition specified as access type 1.
2. One of the Anchor Volume Descriptor Pointers should be recorded in the logical sector 256.
3. The Terminating Descriptor shall be recorded to terminate an extent of a Volume Descriptor Sequence.
4. The Unallocated Space Table and the Unallocated Space Bitmap shall not be recorded.
5. Sparable Partition Map and Sparing Table shall not be recorded.
6. Virtual Partition Map shall not be recorded.
7. Metadata Partition Map, Metadata File and Metadata Mirror File shall be recorded. Metadata Bitmap File shall not be recorded.

For File Structure:

Common requirements for HD DVD shall be applied.

6.z.2 Requirements for HD DVD-RAM

The volume and file structure is simplified as for Overwritable discs using non-sequential recording.

For Volume Structure:

1. A partition on a HD DVD-RAM disc shall be an overwritable partition specified as access type 4.
2. Sparable Partition Map and Sparing Table shall not be recorded.
3. Virtual Partition Map shall not be recorded.
4. Metadata Partition Map, Metadata File and Metadata Bitmap File shall be recorded.

For File Structure:

5. Non-Allocatable Space Stream shall not be recorded.

6.z.3 Requirements for HD DVD-R for SL/DL

The requirements for HD DVD-R for SL/DL discs are under Data updatable structure (VAT) or under HD DVD-ROM compatible structure (read-only partition). The volume and file structure is simplified as for Write-Once discs using sequential recording.
In the case of HD DVD-ROM compatible structure, the requirements are the same as for HD DVD-ROM, refer to 6.z.1. HD DVD-R DL only supports single session.

In the case of Data updatable structure (VAT), following restriction shall be applied.

For Volume Structure:

1. A partition shall be a write-once partition specified as access type 2.
2. The Unallocated Space Table and the Unallocated Space Bitmap shall not be recorded.
3. Sparable Partition Map and Sparing Table shall not be recorded.
4. Only one Open Logical Volume Integrity Descriptor shall be recorded in the Logical Volume Integrity Sequence.
5. Virtual Partition Map shall be recorded. Therefore Metadata Partition Map shall not be recorded.

For File Structure:

6. Only one File Set Descriptor shall be recorded.
7. Non-Allocatable Space Stream shall not be recorded.
8. Virtual Allocation Table and VAT ICB shall be recorded.
9. Metadata File, Metadata Mirror File and Metadata Bitmap File shall not be recorded.

Add a new entry for this DCN to the UDF history table in section 6.17:

Description	DCN number	Updated in UDF Revision	Minimum UDF Read Revision	Minimum UDF Write Revision
Requirements for HD DVD Disc	5154	x.yz	2.50	2.50

(x.yz is the next UDF revision after 2.60)

Document: **OSTA Universal Disk Format** **DCN-5155**
Subject: *Add recommendations for DVD+R DL and DVD+RW DL*
Date: November 29, 2005; last modified March 02, 2006;
June 12, 2006: Editorial change referring to DCN-5163.
Status: Approved March 02, 2006

Description:

This DCN is for the next UDF revision after 2.60 **and as errata for all previous UDF revisions**.
The recommendations for DVD+R and DVD+RW must be adapted to the Dual Layer versions that are available now. Further, the current text is improved. For choosing the Sparing Packet Length for UDF revisions 1.50 and 2.00, see also errata DCN-5163.

Change:

Replace the complete appendix 6.13 by:

6.13 Recommendations for DVD+R and DVD+RW Media

DVD+R and DVD+RW single layer and dual layer media require special consideration due to their nature. The following information and guidelines are established to ensure interchange.

- Logical Sector Size is 2048 Bytes
- 2048 Bytes user data transfer for read and write
- ECC Block Size is 32768 bytes (32KB) and the first sector number of an ECC block shall be an integral multiple of 16

Single layer DVD+R and DVD+RW media have a maximum capacity of 4.7 Gbytes. Dual layer DVD+R DL and DVD+RW DL media have a maximum capacity of 8.5 Gbytes. For more detailed information, see the SCSI-3 MMC command set specification and DVD+R and DVD+RW Basic Format Specification documents. For Mount Rainier formatted DVD+MRW media, see appendix 6.14.

Special care must be taken when UDF structures should be recorded ‘physically far apart’, see 2.2.3.2 and 2.2.13.1. For dual layer media, physically far apart is not the same as logically far apart.

6.13.1 Use of UDF on DVD+R media

DVD+R single layer and dual layer media can be used for disc at once, session at once and incremental recording. Multisession is supported. The general rules of appendix 6.11 apply.

6.13.2 Use of UDF on DVD+RW media

DVD+RW single layer and dual layer media are random readable and writable. Where needed, the DVD+RW drive performs Read-Modify-Write cycles to accomplish this. For DVD+RW media, the drive does not perform defect management. DVD+RW media provide the following features:

- Random read and write access
- Background physical formatting
- The Media Type is Overwritable (partition Access Type 4, overwritable)

Multisession is not supported for DVD+RW media.

6.13.2.1 Requirements

- Sparing shall be managed by the host via the Sparable Partition and a Sparing Table
- The sparing Packet Length shall be 16 logical blocks (32 KB, one ECC block). For UDF revisions 1.50 and 2.00, the sparing Packet Length may be 32 logical blocks, see errata DCN-5163.
- Defective packets known at logical format time shall be allocated by the Non-Allocatable Space Stream, see 3.3.7.2

Preparing a blank DVD+RW medium for UDF usage is done by physical formatting and logical formatting. Physical formatting is writing basic physical structures and writing data to all sectors once (de-icing for Read-Only device compatibility). Logical formatting is writing the mandatory basic UDF file system structures, see 6.13.2.3. Physical formatting can be done prior to logical formatting. This is called physical pre-formatting. However this will take much time. DVD+RW offers the possibility of background physical formatting, see 6.13.2.2. Logical formatting, writing of user data and eject and re-insert of the disc can be performed while background physical formatting is in progress. Physical formatting may be followed by a verification pass.

6.13.2.2 Background physical formatting

When background physical formatting is started, some minimal amount of formatting will be performed and then the de-icing operation will continue in background. From that moment, logical formatting and writing of user data can be performed. The disc can be ejected before background formatting has finished. For such an early eject, the background formatting process must be suspended, using a so-called compatibility stop or a quick stop. After a compatibility stop, the medium is compatible with Read-Only devices. For a compatibility stop, the drive must format all non-recorded areas in between recorded areas at the inner side of the disc. This could cause a significant delay in the early eject process. While background formatting is not yet complete, the delay for a compatibility stop can be reduced greatly by temporarily adapting the file system allocation strategy, see 6.13.2.4. When writing is resumed to a medium where background formatting was suspended, the background physical formatting process must be resumed too. Background physical formatting starts at the inner side of the disc. After a compatibility stop, an uninterrupted part at the inner side of the disc is recorded on layer L0 and for the dual layer disc also an equal part at the inner side of the disc on layer L1. These parts on L0 and L1 correspond to two equal portions, one at the beginning and one at the end of the UDF volume space respectively.

6.13.2.3 Logical formatting

Logical formatting is writing the mandatory basic UDF file system structures, such as VRS, AVDP, VDS, FSD, Root Directory and Sparing Tables.
An AVDP shall be recorded at two of the following locations: 256, N-256 and N, where N is the last valid address in the volume space. However, it is recommended to record an AVDP at all three locations, especially for the DVD+RW DL disc, where the regions for recording of the AVDPs are on both layers at the inner side of the disc, so physically not far apart. Allocation for sparing shall occur during the logical formatting process. The sparing allocation may be zero in length. Defective packets known at logical format time shall not be spared using the Sparing Table but added to the Non-Allocatable Space Stream. Not spared defective packets and packets used for a Sparing Table instance or as sparing area shall always be marked as allocated. Inside the UDF partition space, these packets shall be added to the Non-Allocatable Space Stream and consequently be marked as allocated in the partition Space Set, see 2.2.12 and 3.3.7.2. Outside the UDF partition space, these packets shall be marked as allocated by the Unallocated Space Descriptor. The background physical formatting status shall not influence recording of the LVID. At early eject, the LVID shall be recorded in the same way as it will be recorded on Overwritable media that do not support background physical formatting.

6.13.2.4 Restrictions for DVD-ROM compatibility during background formatting

The restrictions mentioned here are only recommended if DVD-ROM compatibility is required at an early eject while background physical formatting is not yet complete. When background physical formatting is complete, DVD-ROM compatibility is a fact and no restrictions are needed. The restrictions all aim at reduction of compatibility stop delay at an early-eject.
The restrictions during background physical formatting are:

- For single layer DVD+RW, only the first AVDP at 256 must be recorded after background physical formatting has been started. The second and third AVDP must be written after background formatting is complete. As long as there is only one AVDP recorded, the file system is in an intermediate state and is not strictly in compliance with ECMA 167. The dual layer DVD+RW DL does not have this restriction, because all AVDPs can be recorded immediately after background formatting has been started. This is possible because of physical formatting on both layers as described above in 6.13.2.2.
- In order to reduce delay in case of a compatibility stop at early eject, the allocation strategy must be restricted as long as background formatting is not yet complete. The lowest logical sector addresses at the beginning of the volume space and for dual layer DVD+RW DL also the highest logical sector addresses at the end of the volume space should be allocated and recorded first in order to reduce compatibility stop delay.

Document: **OSTA Universal Disk Format** **DCN-5156**
Subject: *Macintosh OS X additions*
Date: November 30, 2005; last modified December 05, 2005
Status: Approved December 05, 2005
Editorial: Small C-code correction, January 11, 2006

Description:

This DCN is meant for the next UDF revision after 2.60, and as errata for previous UDF revisions.
The changes defined in this DCN refer to the UDF 2.60 specification text. However, most of these changes are also relevant for the appropriate sections in previous UDF specifications starting with UDF 1.02. In UDF 2.50, an OS Class 3 with OS Identifier value 1 was introduced for Macintosh OS X, see 6.3.2. However, all references to "Macintosh" in the text of the UDF specifications 1.02 till 2.60 inclusive are in fact for "Macintosh OS 9 and older" and there are no specific rules for Macintosh OS X yet. This DCN wants to distinguish clearly between Macintosh OS X and Macintosh OS 9/older rules and will add Macintosh OS X specific rules where needed.

Changes:

In section 2.2.6.4 remove 2 occurrences of:

This information is needed by the Macintosh OS.

In the title of 3.3.1.1.1 replace: Macintosh
by: Macintosh OS 9/older, Macintosh OS X

Add at the end of section 3.3.1.1.1:

In Macintosh OS X, additional rules regarding the hidden bit are in section 3.3.4.5.4.2.

In the title of 3.3.2.1.2 replace: Macintosh
by: Macintosh OS 9/older

In the title of 3.3.2.1.3 replace: UNIX
by: UNIX, Macintosh OS X

In section 3.3.3.3, in the title of the "Default Permission Values table" replace: Mac OS

by: Mac OS 9/older

replace: UNIX & OS/400
by: UNIX, OS/400, Mac OS X

add at the end of NOTE 1:

Under Macintosh OS X, the *attribute* permission shall either be treated in the same way as UNIX, or be specified by the user.

add at the end of NOTE 2:

Under Macintosh OS X, the *delete* permission shall either be treated in the same way as UNIX, or be specified by the user.

In the title of the "Processing Permissions table"
replace: Mac OS
by: Mac OS 9/older

Add a column at the Processing Permissions table with the following values:

Mac OS X
E
E
E
E
E
E
Note 4
Note 4
Note 4
Note 4

In the paragraph before NOTE 3
replace 2 occurrences of: Macintosh
by: Macintosh OS 9/older

add at the end of section 3.3.3.3:

NOTE 4: When processing permissions under Macintosh OS X, if an implementation chooses to treat the *attribute* permission the same way as UNIX, this permission shall be ignored; if an implementation allows the user to set the *attribute* permission, this permission shall be enforced. Similarly, if an implementation chooses to treat the *delete* permission the same way as UNIX, this permission shall be ignored; if an implementation allows the user to set the *delete* permission, this permission shall be enforced.

At the end of section 3.3.4.5.4.2 change: **NOTE:**
by: **NOTE 1:**
and add a second note:

NOTE 2: Macintosh OS X handles the invisible flag of the Finder Info specially. The invisible flag of the Finder Info is the 15th bit of the FdFlags of UDFFInfo for a file, or the 15th bit of the FrFlags of UDFDInfo for a directory.

- After the value of the Finder Info of a file or a directory is read, the value of the hidden bit in the File Characteristics of this file or directory's File Identifier Descriptor (FID) shall be copied to the invisible flag of the Finder Info returned to the application. If this file or directory does not have a Finder Info and the hidden bit in the FID is set, an all-zero Finder Info with only the invisible flag set shall be returned to the application. If more than one FID points to this file, the invisible flag of the Finder Info returned to the application shall be set to the same value as the value of the hidden bit of the last FID that was used to find this file. The on-disk data shall not be modified when reading.

- When updating the Finder Info on the media, the invisible flag of the Finder Info shall be copied to the hidden bit of the FID that points to this file or directory. If more than one FID points to the file, the hidden bit of at least one FID shall be updated according to the invisible flag of the Finder Info. Which FID is updated is up to the implementation.

This rule improves the interoperability of hidden files between Windows and Macintosh OS X so that hidden files on Windows are hidden on Macintosh OS X and vice versa. For files with hard links, the behavior of hidden files is undefined.

In the title and text of section 4.2.2.1.3 replace: Macintosh
by: Macintosh OS 9/older

Add section 4.2.2.1.7:

4.2.2.1.7 Macintosh OS X

Due to the restrictions imposed by the Mac OS X operating system environment, on the *FileIdentifier* associated with a file or a directory the following methodology shall be employed to handle *FileIdentifier(s)* under the above-mentioned operating system environment:

1. FileIdentifier Lookup: Upon request for a “lookup” of a *FileIdentifier*, a case-sensitive comparison shall be performed. If the case-sensitive comparison fails, a case-insensitive comparison may be performed.

2. Validate FileIdentifier: If the *FileIdentifier* is a valid Mac OS X file identifier for the current system interface, then do not apply the following steps.

3. Invalid Characters: A *FileIdentifier* that contains characters considered invalid within a Mac OS X file name shall have them translated into "_" (#005F). Multiple sequential invalid characters shall be translated into a single "_" (#005F) character. Reference appendix 6.7.2 on invalid characters for a complete list.

4. Long FileIdentifier: In the event that the name does not fit into the limit of the current system interface, the new *FileIdentifier* will consist of the first *N* characters of the *FileIdentifier* at this step in the process, where *N* is the largest possible value such that the first *N* characters of the *FileIdentifier* plus 5 characters (the size of the CRC) is valid in the current system interface.

5. FileIdentifier CRC: Since through the above step 3 and/or 4 process character information from the original *FileIdentifier* is lost, the chance of creating a duplicate *FileIdentifier* in the same directory increases. To greatly reduce the chance of having a duplicate *FileIdentifier* the *file name* shall be modified to contain a CRC of the original *FileIdentifier*.

The CRC has 5 characters. It starts with the separator '#', and followed by a 4 digit CS0 Hex representation of the 16-bit CRC of the original CS0 *FileIdentifier*.

If there is a file extension, the new *FileIdentifier* shall be transformed from the following:

(first *N* characters of the *FileIdentifier* obtained after step 3 without the file extension and the '.' before the file extension) '#' (four characters of CRC) '.' (file extension)

Otherwise if there is no file extension, the new *FileIdentifier* shall be transformed from the following:

(first *N* characters of the *FileIdentifier* obtained after step 3) '#' (four characters of CRC)

In both cases, *N* is the largest possible value such that the transformed *FileIdentifier* is valid in the current system interface.

In section 6.3.2
replace: Macintosh OS X and later releases.
by: Macintosh OS X - generic (includes Cheetah, Puma, Jaguar, Panther, Tiger, and later releases based on the same code base).

In the Appendix 6.18 Developers Registration Form,
replace the entry for: *Macintosh*
by two entries for: *Mac OS 9 and Mac OS X* respectively

editorial note: This change may be overruled by a complete review of the Developer Registration Form as described in the separate DCN-5162.

In appendix 6.7.2 apply the following diff to the name conversion algorithm:

editorial note: The complete C code of appendix 6.7.2 as a result of this change is in a separate document named dcn-5156-annex.

```
*** nameconv-org.c  Thu Nov 17 13:59:25 2005
--- nameconv.c      Fri Dec  9 10:53:31 2005
***************
*** 21,30 ****
    *   Define WIN_NT
    *   Define MAXLEN = 255
    *
!   * Macintosh:
!   *   Define MAC.
    *   Define MAXLEN = 31.
    *
    * UNIX
    *   Define UNIX.
    *   Define MAXLEN as specified by unix version.
--- 21,34 ----
    *   Define WIN_NT
    *   Define MAXLEN = 255
    *
!   * Macintosh OS 9/older:
!   *   Define MAC9.
    *   Define MAXLEN = 31.
    *
+   * Macintosh OS X:
+   *   Define MACOSX
+   *   Define MAXLEN = 255
+   *
    * UNIX
    *   Define UNIX.
    *   Define MAXLEN as specified by unix version.
***************
```

```
*** 43,49 ****
   * byte needs to be unsigned 8-bit, and unicode_t needs to
   * be unsigned 16-bit.
   */
! typedef unsigned int unicode_t;
  typedef unsigned char byte;

  /*** PROTOTYPES ***/
--- 47,53 ----
   * byte needs to be unsigned 8-bit, and unicode_t needs to
   * be unsigned 16-bit.
   */
! typedef unsigned short unicode_t;
  typedef unsigned char byte;

  /*** PROTOTYPES ***/
***************
*** 54,59 ****
--- 58,82 ----
   * printable under your implementation.
   */
  int UnicodeIsPrint(unicode_t);
+
+ #ifdef MACOSX
+ size_t GetMaxUnicodeLen(unicode_t *name, /* the unicode name  */
+ size_t charcnt, /* number of unicode characters */
+ size_t maxUtf8Len); /* maximum size of the utf-8 buffer in bytes */
+
+
+ /*****************************************************************
+  * this function returns the number of bytes required to encode
+  * bytecnt/2 unicode characters into the encoding required by the
+  * current system.
+  *
+  * For example, in Mac OS X 10.4 (Tiger), this is UTF-8 encoding
+  * normalized to NFD (decomposed) from.
+  *
+  * The implementation of this function is not included in this standard.
+  */
+ size_t UTF8EncodeLength(unicode_t *str, size_t bytecnt, int flag);
+ #endif

  /*****************************************************************
   * Translates a long file name to one using a MAXLEN and an illegal
***************
*** 67,79 ****
  int UDFTransName(
  unicode_t *newName,/*(Output)Translated name. Must be of length MAXLEN*/
  unicode_t *udfName, /* (Input)  Name from UDF volume.*/
! int udfLen,         /* (Input)  Length of UDF Name.  */
  {
     int index, newIndex = 0, needsCRC = FALSE;
!    int extIndex, newExtIndex = 0, hasExt = FALSE;
! #ifdef (OS2 | WIN_95 | WIN_NT)
     int trailIndex = 0;
  #endif
     unsigned short valueCRC;
     unicode_t current;
     const char hexChar[] = "0123456789ABCDEF";
--- 90,106 ----
```

```
  int UDFTransName(
  unicode_t *newName,/*(Output)Translated name. Must be of length MAXLEN*/
  unicode_t *udfName, /* (Input)  Name from UDF volume.*/
! int udfLen)         /* (Input)  Length of UDF Name.  */
  {
     int index, newIndex = 0, needsCRC = FALSE;
!    int extIndex = 0, newExtIndex = 0, hasExt = FALSE;
! #if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
     int trailIndex = 0;
  #endif
+ #ifdef MACOSX
+    int decomposedUtf8len = 0;
+ #endif
+
     unsigned short valueCRC;
     unicode_t current;
     const char hexChar[] = "0123456789ABCDEF";
***************
*** 111,117 ****
            }
        }

! #ifdef (OS2 | WIN_95 | WIN_NT)
        /* Record position of last char which is NOT period or space. */
        else if (current != PERIOD && current != SPACE)
        {
--- 138,144 ----
            }
        }

! #if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
        /* Record position of last char which is NOT period or space. */
        else if (current != PERIOD && current != SPACE)
        {
***************
*** 127,135 ****
        {
           needsCRC = TRUE;
        }
     }

! #ifdef (OS2 | WIN_95 | WIN_NT)
     /* For OS2, 95 & NT, truncate any trailing periods and\or spaces. */
     if (trailIndex != newIndex - 1)
     {
--- 154,168 ----
        {
           needsCRC = TRUE;
        }
+
+ #ifdef MACOSX
+       decomposedUtf8len += UTF8EncodeLength(&current, 2, UTF_DECOMPOSED);
+       if (decomposedUtf8len >= MAXLEN)
+          needsCRC = TRUE;
+ #endif
     }

! #if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
     /* For OS2, 95 & NT, truncate any trailing periods and\or spaces. */
     if (trailIndex != newIndex - 1)
```

```
     {
***************
*** 153,159 ****

              if (IsIllegal(current) || !UnicodeIsPrint(current))
              {
-                needsCRC = 1;
                 /* Replace Illegal and non-displayable chars
                  * with underscore.
                  */
--- 186,191 ----
***************
*** 172,177 ****
--- 204,214 ----
           }

           /* Truncate filename to leave room for extension and CRC. */
+ #ifdef MACOSX
+          maxFilenameLen = (MAXLEN - 5) -
+              UTF8EncodeLength(ext, localExtIndex*2, UTF_DECOMPOSED) - 1;
+          newIndex = GetMaxUnicodeLen(newName, newExtIndex, maxFilenameLen);
+ #else
           maxFilenameLen = ((MAXLEN - 5) - localExtIndex - 1);
           if (newIndex > maxFilenameLen)
           {
***************
*** 181,196 ****
           {
              newIndex = newExtIndex;
           }
        }
        else if (newIndex > MAXLEN - 5)
        {
           /*If no extension, make sure to leave room for CRC. */
           newIndex = MAXLEN - 5;
        }
        newName[newIndex++] = CRC_MARK; /* Add mark for CRC. */

        /*Calculate CRC from original filename from FileIdentifier. */
!       valueCRC = unicode_cksum(udfName, udfLen);
        /* Convert 16-bits of CRC to hex characters. */
        newName[newIndex++] = hexChar[(valueCRC & 0xf000) >> 12];
        newName[newIndex++] = hexChar[(valueCRC & 0x0f00) >> 8];
--- 218,238 ----
           {
              newIndex = newExtIndex;
           }
+ #endif
        }
        else if (newIndex > MAXLEN - 5)
        {
           /*If no extension, make sure to leave room for CRC. */
+ #ifdef MACOSX
+          newIndex = GetMaxUnicodeLen(newName, newIndex, MAXLEN - 5);
+ #else
           newIndex = MAXLEN - 5;
+ #endif
        }
        newName[newIndex++] = CRC_MARK; /* Add mark for CRC. */
```

```
        /*Calculate CRC from original filename from FileIdentifier. */
!       valueCRC = unicode_cksum((unsigned short *)udfName, udfLen);
        /* Convert 16-bits of CRC to hex characters. */
        newName[newIndex++] = hexChar[(valueCRC & 0xf000) >> 12];
        newName[newIndex++] = hexChar[(valueCRC & 0x0f00) >> 8];
***************
*** 210,216 ****
     return(newIndex);
  }

! #ifdef (OS2 | WIN_95 | WIN_NT)
  /*************************************************************************
   * Decides if a Unicode character matches one of a list
   * of ASCII characters.
--- 252,258 ----
     return(newIndex);
  }

! #if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
  /*************************************************************************
   * Decides if a Unicode character matches one of a list
   * of ASCII characters.
***************
*** 238,244 ****
     }
     return(found);
  }
! #endif /* OS2 */

  /*************************************************************************
   * Decides whether the given character is illegal for a given OS.
--- 280,286 ----
     }
     return(found);
  }
! #endif /* if defined(OS2) || defined(WIN_95) || defined(WIN_NT) */

  /*************************************************************************
   * Decides whether the given character is illegal for a given OS.
***************
*** 249,256 ****
   */
  int IsIllegal(unicode_t ch)
  {
! #ifdef MAC
!    /* Only illegal character on the MAC is the colon. */
     if (ch == 0x003A)
     {
        return(1);
--- 291,298 ----
   */
  int IsIllegal(unicode_t ch)
  {
! #ifdef MAC9
!    /* Only illegal character on the Mac OS 9/older is the colon. */
     if (ch == 0x003A)
     {
        return(1);
***************
*** 259,266 ****
```

```
     {
        return(0);
     }
! #elif defined UNIX
!    /* Illegal UNIX characters are NULL and slash. */
     if (ch == 0x0000 || ch == 0x002F)
     {
        return(1);
--- 301,308 ----
     {
        return(0);
     }
! #elif defined(UNIX) || defined(MACOSX)
!    /* Illegal UNIX and Mac OS X characters are NULL and slash. */
     if (ch == 0x0000 || ch == 0x002F)
     {
        return(1);
***************
*** 269,275 ****
     {
        return(0);
     }
! #elif defined (OS2 | WIN_95 | WIN_NT)
     /* Illegal char's for OS/2 according to WARP toolkit. */
     if (ch < 0x0020 || UnicodeInString("\\/:*?\"<>|", ch))
     {
--- 311,317 ----
     {
        return(0);
     }
! #elif defined(OS2) || defined(WIN_95) || defined(WIN_NT)
     /* Illegal char's for OS/2 according to WARP toolkit. */
     if (ch < 0x0020 || UnicodeInString("\\/:*?\"<>|", ch))
     {
***************
*** 279,283 ****
     {
        return(0);
     }
! #endif
  }
--- 321,356 ----
     {
        return(0);
     }
! #endif
!    return 1; // should never reach here
  }
+
+ #ifdef MACOSX
+
+ /*****************************************************************
+  * given the maximum size of the utf8 buffer, return the number of
+  * unicode characters that can fit in the utf8 buffer
+  */
+ size_t GetMaxUnicodeLen(
+ unicode_t *name, /* the unicode name  */
+ size_t charcnt, /* number of unicode characters */
+ size_t maxUtf8Len) /* maximum size of the utf-8 buffer in bytes */
+ {
```

```
+    size_t len, i;
+
+    len = 0;
+    for (i=0; i<charcnt; i++)
+    {
+       len += UTF8EncodeLength(name++, 2, UTF_DECOMPOSED);
+       if (len > maxUtf8Len)
+          break;
+    }
+    return i;
+ }
+
+ int UnicodeIsPrint(unicode_t ch)
+ {
+    return 1;
+ }
+
+ #endif
```

Document: OSTA Universal Disk Format Annex to DCN-5156
Subject: *Annex to DCN-5156: Resulting C code of 6.7.2*
Date: December 05, 2005
Editorial: Small C-code correction, January 11, 2006

Description:

This document is an annex to DCN-5156.

Resulting C code of appendix 6.7.2 after applying DCN-5156:

```
/*****************************************************************
 * OSTA UDF compliant file name translation routine for OS/2,
 * Windows 95, Windows NT, Macintosh and UNIX.
 * Copyright 1995 Micro Design International, Inc.
 * Written by Jason M. Rinn.
 * Micro Design International gives permission for the free use of the
 * following source code.
 *****************************************************************
 *
 * To use these routines with different operating systems.
 *
 * OS/2
 *   Define OS2
 *   Define MAXLEN = 254
 *
 * Windows 95
 *   Define WIN_95
 *   Define MAXLEN = 255
 *
 * Windows NT
 *   Define WIN_NT
 *   Define MAXLEN = 255
 *
 * Macintosh OS 9/older:
 *   Define MAC9.
 *   Define MAXLEN = 31.
 *
 * Macintosh OS X:
 *   Define MACOSX
 *   Define MAXLEN = 255
 *
 * UNIX
 *   Define UNIX.
 *   Define MAXLEN as specified by unix version.
 */

#define ILLEGAL_CHAR_MARK 0x005F
#define CRC_MARK          0x0023
#define EXT_SIZE           5
#define TRUE               1
#define FALSE              0
#define PERIOD            0x002E
#define SPACE             0x0020

/*****************************************************************
 * The following two typedef's are to remove compiler dependancies.
 * byte needs to be unsigned 8-bit, and unicode_t needs to
 * be unsigned 16-bit.
 */
```

```
typedef unsigned short unicode_t;
typedef unsigned char byte;

/*** PROTOTYPES ***/
int IsIllegal(unicode_t ch);
unsigned short unicode_cksum(register unsigned short *s, register int n);

/* Define a function or macro which determines if a Unicode character is
 * printable under your implementation.
 */
int UnicodeIsPrint(unicode_t);

#ifdef MACOSX
size_t GetMaxUnicodeLen(unicode_t *name, /* the unicode name  */
size_t charcnt, /* number of unicode characters */
size_t maxUtf8Len); /* maximum size of the utf-8 buffer in bytes */


/*****************************************************************
 * this function returns the number of bytes required to encode
 * bytecnt/2 unicode characters into the encoding required by the
 * current system.
 *
 * For example, in Mac OS X 10.4 (Tiger), this is UTF-8 encoding
 * normalized to NFD (decomposed) from.
 *
 * The implementation of this function is not included in this standard.
 */
size_t UTF8EncodeLength(unicode_t *str, size_t bytecnt, int flag);
#endif

/*****************************************************************
 * Translates a long file name to one using a MAXLEN and an illegal
 * char set in accord with the OSTA requirements.  Assumes the name has
 * already been translated to Unicode.
 *
 * RETURN VALUE
 *
 *    Number of unicode characters in translated name.
 */
int UDFTransName(
unicode_t *newName,/*(Output)Translated name. Must be of length MAXLEN*/
unicode_t *udfName, /* (Input)  Name from UDF volume.*/
int udfLen)         /* (Input)  Length of UDF Name.  */
{
   int index, newIndex = 0, needsCRC = FALSE;
   int extIndex = 0, newExtIndex = 0, hasExt = FALSE;
#if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
   int trailIndex = 0;
#endif
#ifdef MACOSX
   int decomposedUtf8len = 0;
#endif

   unsigned short valueCRC;
   unicode_t current;
   const char hexChar[] = "0123456789ABCDEF";

   for (index = 0; index < udfLen; index++)
   {
      current = udfName[index];

      if (IsIllegal(current) || !UnicodeIsPrint(current))
      {
         needsCRC = TRUE;
        /* Replace Illegal and non-displayable chars with underscore. */
         current = ILLEGAL_CHAR_MARK;
         /* Skip any other illegal or non-displayable characters. */
```

```
         while(index+1 < udfLen && (IsIllegal(udfName[index+1])
                     || !UnicodeIsPrint(udfName[index+1])))
         {
            index++;
         }
      }

      /* Record position of extension, if one is found. */
      if (current == PERIOD && (udfLen - index -1) <= EXT_SIZE)
      {
         if (udfLen == index + 1)
         {
            /* A trailing period is NOT an extension. */
            hasExt = FALSE;
         }
         else
         {
            hasExt = TRUE;
            extIndex = index;
            newExtIndex = newIndex;
         }
      }

#if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
      /* Record position of last char which is NOT period or space. */
      else if (current != PERIOD && current != SPACE)
      {
         trailIndex = newIndex;
      }
#endif

      if (newIndex < MAXLEN)
      {
         newName[newIndex++] = current;
      }
      else
      {
         needsCRC = TRUE;
      }

#ifdef MACOSX
      decomposedUtf8len += UTF8EncodeLength(&current, 2, UTF_DECOMPOSED);
      if (decomposedUtf8len >= MAXLEN)
         needsCRC = TRUE;
#endif
   }

#if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
   /* For OS2, 95 & NT, truncate any trailing periods and\or spaces. */
   if (trailIndex != newIndex - 1)
   {
      newIndex = trailIndex + 1;
      needsCRC = TRUE;
      hasExt = FALSE; /* Trailing period does not make an extension. */
   }
#endif

   if (needsCRC)
   {
      unicode_t ext[EXT_SIZE];
      int localExtIndex = 0;
      if (hasExt)
      {
         int maxFilenameLen;
         /* Translate extension, and store it in ext. */
         for(index = 0; index<EXT_SIZE && extIndex + index +1 < udfLen;
              index++ )
         {
```

```
            current = udfName[extIndex + index + 1];

            if (IsIllegal(current) || !UnicodeIsPrint(current))
            {
               /* Replace Illegal and non-displayable chars
                * with underscore.
                */
               current = ILLEGAL_CHAR_MARK;
               /* Skip any other illegal or non-displayable
                * characters.
                */
               while(index + 1 < EXT_SIZE
                           && (IsIllegal(udfName[extIndex + index + 2])
                           || !UnicodeIsPrint(udfName[extIndex + index + 2])))
               {
                  index++;
               }
            }
            ext[localExtIndex++] = current;
         }

         /* Truncate filename to leave room for extension and CRC. */
#ifdef MACOSX
         maxFilenameLen = (MAXLEN - 5) -
             UTF8EncodeLength(ext, localExtIndex*2, UTF_DECOMPOSED) - 1;
         newIndex = GetMaxUnicodeLen(newName, newExtIndex, maxFilenameLen);
#else
         maxFilenameLen = ((MAXLEN - 5) - localExtIndex - 1);
         if (newIndex > maxFilenameLen)
         {
            newIndex = maxFilenameLen;
         }
         else
         {
            newIndex = newExtIndex;
         }
#endif
      }
      else if (newIndex > MAXLEN - 5)
      {
         /*If no extension, make sure to leave room for CRC. */
#ifdef MACOSX
         newIndex = GetMaxUnicodeLen(newName, newIndex, MAXLEN - 5);
#else
         newIndex = MAXLEN - 5;
#endif
      }
      newName[newIndex++] = CRC_MARK; /* Add mark for CRC. */

      /*Calculate CRC from original filename from FileIdentifier. */
      valueCRC = unicode_cksum((unsigned short *)udfName, udfLen);
      /* Convert 16-bits of CRC to hex characters. */
      newName[newIndex++] = hexChar[(valueCRC & 0xf000) >> 12];
      newName[newIndex++] = hexChar[(valueCRC & 0x0f00) >> 8];
      newName[newIndex++] = hexChar[(valueCRC & 0x00f0) >> 4];
      newName[newIndex++] = hexChar[(valueCRC & 0x000f)];

      /* Place a translated extension at end, if found. */
      if (hasExt)
      {
         newName[newIndex++] = PERIOD;
         for (index = 0;index < localExtIndex ;index++ )
         {
            newName[newIndex++] = ext[index];
         }
      }
   }
   return(newIndex);
```

```
}

#if defined(OS2) || defined(WIN_95) || defined(WIN_NT)
/************************************************************************
 * Decides if a Unicode character matches one of a list
 * of ASCII characters.
 * Used by OS2 version of IsIllegal for readability, since all of the
 * illegal characters above 0x0020 are in the ASCII subset of Unicode.
 * Works very similarly to the standard C function strchr().
 *
 * RETURN VALUE
 *
 *    Non-zero if the Unicode character is in the given ASCII string.
 */
int UnicodeInString(
unsigned char *string,  /* (Input) String to search through.   */
unicode_t ch)  /* (Input) Unicode char to search for. */
{
   int found = FALSE;
   while (*string != '\0' && found == FALSE)
   {
      /* These types should compare, since both are unsigned numbers. */
      if (*string == ch)
      {
         found = TRUE;
      }
      string++;
   }
   return(found);
}
#endif /* if defined(OS2) || defined(WIN_95) || defined(WIN_NT) */

/************************************************************************
 * Decides whether the given character is illegal for a given OS.
 *
 * RETURN VALUE
 *
 *    Non-zero if char is illegal.
 */
int IsIllegal(unicode_t ch)
{
#ifdef MAC9
   /* Only illegal character on the Mac OS 9/older is the colon. */
   if (ch == 0x003A)
   {
      return(1);
   }
   else
   {
      return(0);
   }
#elif defined(UNIX) || defined(MACOSX)
   /* Illegal UNIX and Mac OS X characters are NULL and slash. */
   if (ch == 0x0000 || ch == 0x002F)
   {
      return(1);
   }
   else
   {
      return(0);
   }
#elif defined(OS2) || defined(WIN_95) || defined(WIN_NT)
   /* Illegal char's for OS/2 according to WARP toolkit. */
   if (ch < 0x0020 || UnicodeInString("\\/:*?\"<>|", ch))
   {
      return(1);
   }
   else
```

```
   {
      return(0);
   }
#endif
   return 1; // should never reach here
}

#ifdef MACOSX

/*****************************************************************
 * given the maximum size of the utf8 buffer, return the number of
 * unicode characters that can fit in the utf8 buffer
 */
size_t GetMaxUnicodeLen(
unicode_t *name, /* the unicode name  */
size_t charcnt, /* number of unicode characters */
size_t maxUtf8Len) /* maximum size of the utf-8 buffer in bytes */
{
   size_t len, i;

   len = 0;
   for (i=0; i<charcnt; i++)
   {
      len += UTF8EncodeLength(name++, 2, UTF_DECOMPOSED);
      if (len > maxUtf8Len)
         break;
   }
   return i;
}

int UnicodeIsPrint(unicode_t ch)
{
   return 1;
}

#endif
```

Document: **OSTA Universal Disk Format** **DCN-5157**
Subject: *Unicode Version and Unicode Normalization Form*
Date: December 6, 2005; last modified December 14, 2005
Status: Approved March 02, 2006

Description:

This DCN is meant for the next UDF revision after 2.60, and as errata for all UDF revisions 1.02 till 2.60 included.
This DCN enables the use of d-characters from newer Unicode Standard versions than strictly defined in UDF section 2.1.1.
Further, Unicode Normalization Form C (NFC), as used by Windows is recommended for recording of d-character identifiers on all UDF media. This also avoids e.g. file identifiers that are 'optically identical' but are not identical for UDF because they are represented in a different normalization form on the medium.
The changes proposed in this DCN are with respect to the current UDF 2.60 text. Note that UDF revisions 1.02 and 1.50 are currently referring to Unicode Standard Version 1.1 opposed to Unicode Standard Version 2.0 as currently for UDF 2.00 and higher revisions. It is now proposed to let all UDF revisions refer to The Unicode Standard 4.0 as a *reference version*.

Change:

In section 2.1.1

Replace:

OSTA CS0 shall consist of the d-characters specified in The Unicode Standard, Version 2.0 (ISBN 0-201-48345-9 from Addison-Wesley Publishing Company http://www.awl.com/, see also http://www.unicode.org/), excluding #FEFF and #FFFE, stored in the *OSTA Compressed Unicode* format which is defined as follows:

by:

OSTA CS0 shall consist of the d-characters specified in The Unicode Standard, excluding the characters #FEFF and #FFFE. The Unicode Standard reference version is Version 4.0 (ISBN 0-321-18578-1 from Addison-Wesley Publishing Company http://www.awl.com/, see also http://www.unicode.org/). Because of the stability policy defined in the Unicode Standard (http://www.unicode.org/standard/stability_policy.html), also older or newer Unicode versions can be used without expecting backward or forward compatibility problems.

To improve interoperability among different platforms, the Unicode d-character identifiers stored on UDF media should be normalized to Normalization Form C (NFC), see Unicode Standard Annex #15 (http://www.unicode.org/unicode/reports/tr15).

NOTE 1: Since Windows uses NFC form, most existing UDF media and UDF implementations on Windows (including those that are not aware of Unicode normalization) already follow this recommendation.
UDF implementations using a different Normalization Form should still write d-character identifiers in NFC form onto the UDF medium and perform conversion to or from that different Normalization Form when needed. An example of this is MAC OS using Normalization Form D (NFD).
Implementations must be aware that normalization conversions of d-character identifiers may increase or decrease the number of Unicode characters of the identifier.

Unicode characters are stored in the *OSTA Compressed Unicode* format, which is defined as follows:

replace (2 occurrences): Unicode 2.0
by: Unicode

replace: **NOTE:**
by: **NOTE 2:**

Document: OSTA Universal Disk Format **DCN-5159**
Subject: *Add additional recommendations for BD Read-only Disc*
Date: January 24, 2006
Status: Approved March 02, 2006

Description:

This DCN is for the next UDF revision after UDF 2.60 **and for the UDF 2.50 and UDF 2.60 errata**.

The purpose of this DCN is to provide additional information for the BD Read-Only Disc Format to support good interchangeability between both computer systems and consumer appliances using Blu-ray Read-Only Disc.

For BD Read-Only disc with "BDMV Application", there are two types of discs with an ECC Block Size of 64KB or 32KB. Also, "BDMV Application" has a new additional directory immediately under the root directory to certify interactive applications.

Change:

In the second paragraph of section 6.16

replace: • Blu-ray Disc Read-Only Format (BD-ROM)
by: • Blu-ray Disc Read-Only Format (BD-ROM), see note below

and replace: 2. ECC Block Size is 65536 bytes (64KB)
by: 2. ECC Block Size is 65536 bytes (64KB), see note below

Add a following note at the end of section 6.16:

NOTE: There is a Blu-ray Read Only Format with the "BDMV Application" specified on a disc with a capacity of 4.7 Gbytes or 8.5 Gbytes. Its ECC Block Size is 32768 bytes (32KB). All other requirements for this format are the same as for BD-ROM.

In the third paragraph of section 6.16.4 replace:

The "BDMV Application" is a Video Application Format for BD-ROM discs, including AV Stream and database for playback the AV Stream.
The "BDMV" directory immediately under the root directory is reserved for the BDMV application.

by:

The ”BDMV Application” is a Video Application Format for BD-ROM discs, including AV Stream and database for playback of the AV Stream. It also supports certification of interactive applications.
The “BDMV” and “CERTIFICATE” directories immediately under the root directory are reserved for the BDMV application.

Document: **OSTA Universal Disk Format** **DCN-5160**
Subject: *More prominent role for Extended File Entry*
Date: January 27, 2006; last modified March 16, 2006
Status: Approved June 12, 2006

Description:

This DCN is for the next UDF revision after 2.60 and as clarification for UDF 2.00 and higher revisions.
Since UDF 2.00, the Extended File Entry descriptor should be used instead of the File Entry descriptor, see 3.3.5. However, the sections 2.3.6 and 3.3.3 are only about FE, no trace of EFE and there are no specific EFE sections. The result is that in most cases FE is used by implementations.
Sections 2.3.6 and 3.3.3 are adapted in such a way that it covers both EFE and FE with a more prominent role for EFE. No rule changes in this DCN.
Editorial: Mind that the approved DCN-5153 updates the same sections.

Changes:

Add a new entry for this DCN to the UDF history table in section 6.17:

Description	DCN number	Updated in UDF Revision	Minimum UDF Read Revision	Minimum UDF Write Revision
More prominent role for Extended File Entry	5160	x.yz	2.00	2.00

(editorial: x.yz is the next UDF revision after 2.60)

In 3.3.5.1 replace:

File Entries and Extended File Entries may be freely mixed. In particular, compatibility with old reader implementations can be maintained for certain files.

by:

File Entries and Extended File Entries may be freely mixed. In particular, compatibility with old reader implementations can be maintained for certain files. However, the use of an Extended File Entry instead of a File Entry is recommended, see 3.3.5.

Replace section 2.3.6 by:

2.3.6 Extended File Entry and File Entry

```
struct ExtendedFileEntry {              /* ECMA 167 4/14.17 and 4/14.9 */
        struct tag          DescriptorTag;
        struct icbtag       ICBTag;
        Uint32              Uid;
        Uint32              Gid;
        Uint32              Permissions;
        Uint16              FileLinkCount;
        Uint8               RecordFormat;
        Uint8               RecordDisplayAttributes;
        Uint32              RecordLength;
        Uint64              InformationLength;
        Uint64              ObjectSize;                 /* EFE only */
        Uint64              LogicalBlocksRecorded;
        struct timestamp    AccessDateAndTime;
        struct timestamp    ModificationDateAndTime;
        struct timestamp    CreationDateAndTime;        /* EFE only */
        struct timestamp    AttributeDateAndTime;
        Uint32              Checkpoint;
        byte                Reserved[4];                /* EFE only */
        struct long_ad      ExtendedAttributeICB;
        struct long_ad      StreamDirectoryICB;         /* EFE only */
        struct EntityID     ImplementationIdentifier;
        Uint64              UniqueID,
        Uint32              LengthofExtendedAttributes;
        Uint32              LengthofAllocationDescriptors;
        byte                ExtendedAttributes[];
        byte                AllocationDescriptors[];
}
```

The total length of an *Extended File Entry* (*EFE)* or *File Entry (FE)* shall not exceed the size of one logical block. It is recommended to use an *EFE* instead of an *FE* for all cases.

An *EFE* is a superset of an *FE*, which means that an *EFE* has all fields of an *FE* with interleaved some extra fields that are marked in the structure above with "**EFE only**". Note that the offsets of identical fields may be different for *EFE* and *FE*. Generally, "*Extended File Entry*" can replace "*File Entry*" throughout the text of this specification.

If a Metadata Partition Map is recorded on a volume, then all (*Extended) File Entries*, Allocation Descriptor Extents and directory data shall be recorded in the Metadata Partition - i.e. in logical blocks allocated to the Metadata and/or Metadata Mirror File.
For details including exceptions see section 2.2.13.

Replace section 3.3.3 by:

3.3.3 Extended File Entry and File Entry

```
struct ExtendedFileEntry {              /* ECMA 167 4/14.17 and 4/14.9 */
        struct tag          DescriptorTag;
        struct icbtag       ICBTag;
        Uint32              Uid;
        Uint32              Gid;
        Uint32              Permissions;
        Uint16              FileLinkCount;
        Uint8               RecordFormat;
        Uint8               RecordDisplayAttributes;
        Uint32              RecordLength;
        Uint64              InformationLength;
        Uint64              ObjectSize;                 /* EFE only */
        Uint64              LogicalBlocksRecorded;
        struct timestamp    AccessDateAndTime;
        struct timestamp    ModificationDateAndTime;
        struct timestamp    CreationDateAndTime;        /* EFE only */
        struct timestamp    AttributeDateAndTime;
        Uint32              Checkpoint;
        byte                Reserved[4];                /* EFE only */
        struct long_ad      ExtendedAttributeICB;
        struct long_ad      StreamDirectoryICB;         /* EFE only */
        struct EntityID     ImplementationIdentifier;
        Uint64              UniqueID,
        Uint32              LengthofExtendedAttributes;
        Uint32              LengthofAllocationDescriptors;
        byte                ExtendedAttributes[];
        byte                AllocationDescriptors[];
}
```

See section 2.3.6 for rules and distinction between Extended File Entry (EFE) and File Entry (FE).

Document: **OSTA Universal Disk Format** **DCN-5161**
Subject: *Treat Fixed Packets in the same way as ECC Blocks*
Date: February 21, 2006; last modified March 27, 2006
Status: Approved June 12, 2006

Description:

This DCN is for the next UDF revision after 2.60 **and as errata for the UDF revisions 1.50 till 2.60 inclusive**.
UDF rules for ECC blocks, like alignment etc., must also apply for fixed packet media like CD-RW. The easiest way to accomplish this is to add a remark to the ECC Block and Fixed Packet definitions. It would e.g. be strange not to align Metadata Partition extents on fixed packet boundaries for CD-RW when there is no Sparable Partition. Further, it seems that it is not clearly defined that the logical sector address of the first sector of a fixed packet must be an integer multiple of the packet length.

Changes:

In 1.3.2 replace:

ECC Block Size (bytes)	This term refers to values defined in relevant device and/or media specifications. The reader should consult the appropriate document – for example, the "MMC" or "Mt. Fuji" specifications for CD/DVD class media. For media exposing no such concept externally (e.g. hard disc) this term shall be interpreted to mean the sector size of the media.
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 according to the Method 2 addressing specified in Orange Book parts-II and -III.

by:

ECC Block Size (bytes)	This term refers to values defined in relevant device and/or media specifications. The reader should consult the appropriate document – for example, the "MMC" or "Mt. Fuji" specifications for CD/DVD class media. For media exposing no such concept externally (e.g. hard disc) this term shall be interpreted to mean the sector size of the media. Although not strictly the same, media with fixed packets, like CD-RW, also have to apply to the ECC block rules in this specification, where a fixed packet is assumed to be equal to an ECC Block.
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 according to the Method 2 addressing specified in Orange Book parts-II and -III. On a fixed packet medium with a UDF file system, the packets shall be equal in size for all tracks of the medium. The logical sector address of the first sector of each packet shall be an integer multiple of the number of logical sectors per Fixed Packet. Fixed Packets media must also obey to ECC Block rules, see the ECC Block Size definition above.

In 6.10.2.5 replace:

Note that packets may not be aligned to 32 sector boundaries.

by:

Note that packets and tracks shall be aligned to 32 sector boundaries, see the Fixed Packet definition in 1.3.2.

Add a new entry for this DCN to the UDF history table in section 6.17:

Description	DCN number	Updated in UDF Revision	Minimum UDF Read Revision	Minimum UDF Write Revision
Treat Fixed Packets in the same way as ECC Blocks	5161	x.yz	1.50	1.50

(editorial: x.yz is the next UDF revision after UDF 2.60)

Document: **OSTA Universal Disk Format** **DCN-5162**
Subject: *Simplification of UDF Developer Registration Form*
Date: February 27, 2006
Status: Approved March 02, 2006

Description:

This DCN is for the next UDF revision after 2.60, all previous UDF revisions and for the "Registered UDF Developers" section of the OSTA UDF web page.
On the December 2005 UDF committee meeting it was decided to simplify the UDF Developer Registration Form on the OSTA UDF web page section "Registered UDF Developers" and in appendix 6.18 of the UDF specification. Instead of detailed choices about the support of UDF revisions, media types and Operating Systems, each developer company is offered a text box where this information can be described in short.

Changes:

Add a new entry for this DCN to the UDF history table in section 6.17:

Description	DCN number	Updated in UDF Revision	Minimum UDF Read Revision	Minimum UDF Write Revision
Simplification of UDF Developer Registration Form	5162	x.yz	1.02	1.02

(x.yz is the next UDF revision after 2.60)

Replace the Developer Registration Form on the OSTA UDF web page and in appendix 6.18 by:

Optical Storage Technology Association

OSTA Universal Disk Format Specification Developer Registration Form

Name: _____

Company: _____

Address: _____

_____

City: _____

State/Province: _____

Zip/Postal Code: _____

Country: _____

Phone: _____ FAX: _____

Email: _____

Please indicate what value you plan to use as EntityID "*Developer ID" to identify your implementation. Developer ID: "*_____"

The Developer ID should uniquely identify your company as well as your product, see note 2 of section 2.1.5.2 in the latest UDF specification. The Developer ID should not start with "*UDF". The registered developer id can be extended with a suffix containing e.g. version information, as long as the total Developer ID (including "*") does not exceed 23 characters.

Please indicate which UDF revisions you plan to support:

_____

_____

Please indicate which media types you plan to support:

_____

_____

Please indicate which Operating Systems you plan to support:

_____

_____

O Please add my email address to the OSTA UDF email reflector.

O Please send an OSTA Membership kit.

Send completed form to OSTA, see http://www.osta.org/osta/contact.htm