



1

2 **NAND Software Working Group**

3 **Invitation for Collaboration on Software Interface** 4 **Standardization**

5 **Version 1.0 – 04 13 June 2006**

6

7

Copyright © 2004 MIPI Alliance, Inc. All rights reserved.

8

9

10

The material contained herein is not a license, either expressly or impliedly, to any IPR owned or controlled by any of the authors or developers of this material or MIPI. The material contained herein is provided on an "AS IS" basis and to the maximum extent permitted by applicable law, this material is provided AS IS AND WITH ALL FAULTS, and the authors and developers of this material and MIPI hereby disclaim all other warranties and conditions, either express, implied or statutory, including, but not limited to, any (if any) implied warranties, duties or conditions of merchantability, of fitness for a particular purpose, of accuracy or completeness of responses, of results, of workmanlike effort, of lack of viruses, and of lack of negligence.

11

12

13

14

15

16

ALSO, THERE IS NO WARRANTY OR CONDITION OF TITLE, QUIET ENJOYMENT, QUIET POSSESSION, CORRESPONDENCE TO DESCRIPTION OR NON-INFRINGEMENT WITH REGARD TO THIS MATERIAL. IN NO EVENT WILL ANY AUTHOR OR DEVELOPER OF THIS MATERIAL OR MIPI BE LIABLE TO ANY OTHER PARTY FOR THE COST OF PROCURING SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA, OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, OR SPECIAL DAMAGES WHETHER UNDER CONTRACT, TORT, WARRANTY, OR OTHERWISE, ARISING IN ANY WAY OUT OF THIS OR ANY OTHER AGREEMENT RELATING TO THIS MATERIAL, WHETHER OR NOT SUCH PARTY HAD ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.

17

18

19

20

21

22

23

24

25

Questions pertaining to this document, or the terms or conditions of its provision, should be addressed to:

T10/06-379r0

NAND Software Working Group

Invitation for Collaboration

26
27 MIPI Alliance, Inc.
28 c/o IEEE-ISTO
29 445 Hoes Lane
30 Piscataway,
31 Attn: Board Secretary

NJ

08854

32 **Introduction**

33

34 **Contributors List**

35

COMPANY	NAME
M-Systems Flash Disk Pioneers	Avraham Shimor
Micron Technology, Inc.	Wanmo Wong
Samsung Electronics	Yejin Moon
STMicroelectronics N.V.	Antonio Furno
STMicroelectronics N.V.	Thierry Vuillaume
Texas Instruments Incorporated	Leonardo Estevez
Toshiba Corporation	Shuji Takano
Toshiba Corporation	Patrick Elzer

T10/06-379r0

NAND Software Working Group

Invitation for Collaboration

36 Contents

37 Draft Version 1.0 – 00 07 March 2006i

38 Introduction iii

39 1 Overview7

40 1.1 Scope7

41 1.2 Purpose8

42 2 Terminology8

43 2.1 Definitions8

44 2.2 Abbreviations9

45 2.3 Acronyms9

46 3 References9

47 4 Problems of the current Hard Disk or block device emulation11

48 4.1 Power Loss Robustness12

49 4.2 Longevity.....12

50 4.3 Flash Media Reliability13

51 5 Possible new application domain requirements.....13

52 5.1 Digital Rights Management (DRM)13

53 5.2 Compression14

54 5.3 Encryption/security.....15

55 5.4 Data portability15

56 6 Performance issues16

57 6.1 Audio/video streaming requirement **Error! Bookmark not defined.**

58 7 Review of non Block Device oriented alternatives16

59 7.1 Journaling Flash File System (JFFS2 and JFFS3).....17

60 7.2 SCSI Object-based Storage Device model (OSD).....17

61 7.3 FTP Storage19

T10/06-379r0

NAND Software Working Group

Invitation for Collaboration

62 8 Attributes of the “ideal solution”19

63 Invitation for Collaboration on Software Interface 64 Standardization

65 1 Overview

66 1.1 Scope

67 What is MIPI ?

68 The Mobile Industry Processor Interface (MIPI) Alliance (<https://www.mipi.org/>) is an open membership
69 organization that includes over 90 leading companies in the mobile industry that share the objective of
70 defining and promoting open specifications for interfaces to mobile application processors. Through these
71 specifications, the MIPI Alliance intends to speed deployment of new services to mobile users by
72 establishing specifications for standard hardware and software interfaces to mobile application processors
73 and encouraging the adoption of those standards throughout the industry value chain.

74 Founded by ARM, Nokia, STMicroelectronics and Texas Instruments, the MIPI Alliance is intended to
75 complement existing standards bodies such as the Open Mobile Alliance and 3GPP, with a focus on
76 microprocessors, peripherals and software interfaces.

77 What is the MIPI NAND SW Working Group ?

78 The NAND SW Working Group has been established by the MIPI Alliance in the end of 2005 to
79 investigate on SW interfaces standardization for a better and faster integration of NAND products into SW
80 platforms. The group is composed of leading NAND flash industry players and in particular mains NAND
81 manufacturers.

82 The Working Group will focus on:

- 83 ○ Identifying a state-of-the-art SW architecture comprising the embedded *NAND-like Flash* storage
84 management and anticipating coming NAND flash usage.
- 85 ○ Standardizing the SW interfaces (functions, parameters, return codes...) between the identified
86 architecture components that would be under responsibility of NAND flash manufacturer, and
87 publishing corresponding specification(s).
- 88 ○ Developing additional documentation to help to use and to advertise this interface (white paper,
89 examples...).

90 As a necessary factor for success, the architecture and interfaces defined by the NAND SW Working
91 Group will aim to stay compatible with existing and future Operating Systems architecture.

92 93 What are the problems ? 94

95 Traditionally, as the de-facto standard industry practice, NAND Flash devices in embedded and mobile
96 platforms (much like in PC platforms) are used to emulate Hard Disks or block devices. There are two
97 types of problems associated with this contemporary approach:

- 98 • The current Hard Disk emulation paradigm is problematic in the sense that it creates performance
99 and reliability problems originating from mismatch between this block device emulation and the
100 unique characteristics of the NAND Flash technology;
- 101 • But maybe even more significantly, there is a growing awareness in the Flash industry that certain
102 functional topics like security, DRM, compression, etc., are very difficult to implement over the
103 block-device, sector oriented operation of the Hard Disk emulation.

104

105 The NAND SW WG is seeking solutions to improve the usage of NAND flash on file systems, to fasten
106 integration of NAND into platforms and to integrate emerging mass storage technologies.

107

108 **Scope and Document overview**

109 This paper provides a summary/review of the topics related the interface between NAND Flash storage
110 devices and the embedded/mobile computing platforms and the corresponding Operating Systems:

- 111 ○ Elaboration of the problems in the current block-device emulation (FAT/FTL) approach.
- 112 ○ Overview of the new functional topics of security, DRM, compression and data portability.
- 113 ○ Description of associated Performance Issues.
- 114 ○ Review of alternative (non block-device or non FAT oriented) storage management methods,
115 focusing on their interfaces with the Operating Systems.
- 116 ○ Outline of a “whish list” requirements specification for a new “ideal” solution.

117 .

118 **1.2 Purpose**

119 The purpose and the intended use of this document is to serve as a **background material**, to facilitate the
120 interaction between leading Operating System vendors and mobile/embedded computing platform vendors
121 with the MIPI NAND Software Workgroup. The objective of the interaction is to **jointly investigate** the
122 **need and possibility** to extend functionality of the NAND Flash management and interface software
123 beyond the current de-facto industry standard practice of block-device disk emulation. The result of this
124 investigation will be used to align future NAND SW WG standardization efforts with OS and platform
125 vendor’s roadmap.

126 **Clarification.** This document is highlighting some problems related to the current software approach for
127 NAND devices; however **this document doesn’t impose any solutions**, it’s a starting point for
128 **collaboration with leading OS vendors** and mobile/embedded computing platform vendors in definition
129 of NAND Flash software interface with the objective of maximizing generality and industry acceptance.

130 **2 Terminology**

131 **2.1 Definitions**

132 **Module:** Coherent group of functions, classes or other software definition handling a specific subject of
133 the software framework. Each Module is identified by a single name and acronym.

134

135 **2.2 Abbreviations**

136 e.g. For example

137 **2.3 Acronyms**

138 DRM Digital Right Management

139 INCITS InterNational Committee for Information Technology Standards

140 JFFS Journaling Flash File System

141 MIPI Mobile Industry Processor Interface

142 OSD Object-based Storage Device

143 SCSI Small Computer System Interface

144 FTL Flash Translation Layer

145 FAT File Allocation Table

146 LBA Logical Block Address

147 PBA Physical Block Address

148 SW Software

149 HW Hardware

150 API Application Programming Interface

151 LFN Long File Name

152 MMC Multi Media Card

153 SD Secure Digital [Card]

154

155 **3 References**

156

157 [JFFS] [The Journaling Flash File System](#), David Woodhouse, Red Hat, Inc.,
158 dwmw2@cambridge.redhat.com

159

160 [JFFS2] [JFFS : The Journalling Flash File System](#), David Woodhouse, Red Hat, Inc.,
161 dwmw2@cambridge.redhat.com

162

163 [JFFS3] [JFFS3 design issues](#), Artem B. Bityutskiy, dedekind@infradead.org

- 164 [OSD2] [Object-Based Storage Devices - 2 \(OSD-2\)](#) INCITS T10 Technical Committee,;
165 2004/10/04, Rev: 00, Status: Development, Project: 1729-D, File: osd2r00.pdf (1313078 bytes)
- 166 [OSD DEMO] SEAGATE, EMULEX AND IBM TEAM UP TO DEMONSTRATE
167 INDUSTRY'S FIRST STANDARDS-COMPLIANT OBJECT-BASED STORAGE DEVICES
- 168 [OSD PROJ1] Seagate Object Based Storage Project
- 169 [OSD PROJ2] IBM Haifa Labs – Object Store Project
- 170 [PTP] [ISO 15740:2005](#): Electronic still picture imaging. Picture transfer protocol (PTP) for digital
171 still photography devices.

172 **4 Problems of the current Hard Disk or block device emulation**

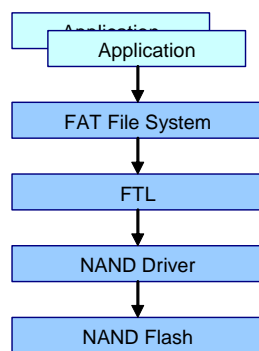
173

174 Traditionally, as the de-facto standard industry practice, NAND Flash devices in embedded and mobile
 175 platforms (much like in PC platforms) are used to emulate Hard Disks or block devices. The obvious
 176 advantage in this practice is that it is leveraging on existing Operating System software components: as far
 177 as either the Operating System or even the application software is concerned the actual technology
 178 implementation of the NAND Flash devices is totally transparent: for all practical purposes it behaves like
 179 yet another Hard Disk Drive. In the context of this paper, by *Hard Disk or block device emulation* we refer
 180 to the abstraction whereby a storage device, independent of the underlying physical storage technology
 181 employed (such as magnetic, optical, or flash) is presenting or exporting a simple, yet powerful, atomic
 182 interface of allowing the reading and writing a fixed sized information unit or sector, usually addressable
 183 by a linear index (so the storage device is logically represented as a linear array of sectors). On top of this
 184 simple scheme, the Operating Systems and its storage and file management components implement
 185 additional software based layers of abstractions like for example the FAT type file system which was
 186 introduced in the 1980s. (FAT is extant in small systems. Microsoft chose FAT as the file system in
 187 Windows CE, and FAT is used in Symbian. Furthermore, FAT is the standard file system format for flash
 188 cards like MMC and SD.)

Deleted: ¶
¶

189 In all implementations of file systems over block devices some of the sectors are used to store the actual
 190 information content associated with files, whereas additional sectors are used to store supporting,
 191 addressing and organization information such as allocation tables, directory entries, etc. The important
 192 implication of this method for the context of this paper is that most application level storage/disk related
 193 operations (transactions) require several sector (or device block) level operations; from the device “point
 194 of view” these are unrelated and independent “atomic” operations.

195 NAND Flash storage devices exhibit several functional characteristics which are originating from the
 196 essence of this technology: information can be read and written in the form of *pages* (currently of typically
 197 512 or 2048 bytes, with a tendency to increase as the capacity of the devices increases), once a page is
 198 written, it cannot be re-written or updated, unless the *erase block* (consisting a large number of pages)
 199 containing the page in question is erased first. Obviously, erasing a block implies moving/copying of all or
 200 at least some of the pages to another (“fresh”, yet unwritten) location in order to avoid information loss.
 201 Yet another important characteristics of the flash technology is that the number of erase cycles for a given
 202 erase block is limited, thus the number of erase operations needs to be minimized in order to prolong the
 203 useful life of the device.



204

205

Figure 1: FAT and FTL Architecture

206

207 Due to the above technology specific limitations the block device emulation cannot be directly facilitated
208 by a flash device, thus the common technique is use an intermediate software layer between the Operating
209 System component (usually: the file system) “expecting” to interact with a sector based block device and
210 the physical flash memory (see figure 1). A common name for this block device or disk emulation layer is
211 *Flash Translation Layer* (FTL).

212
213 The current practice of the Hard Disk emulation paradigm is problematic in the sense that it creates
214 performance and reliability problems originating from mismatch between this block device emulation and
215 the unique characteristics of the NAND Flash technology; The following sections outline some of these
216 problems.

217 **4.1 Power Loss Robustness**

218 By Power Loss Robustness in the context of this paper we refer to the capability of a NAND Flash based
219 device with its corresponding FTL to ensure the integrity of the information stored (primarily the file
220 content “payload” information, but also the flash management related metadata) even in the presence of
221 power loss failure events.

222 The root cause of this problem is the fact that in file system implementation over any **Module** block device
223 most application level transactions (e.g. create a file, write a new record or data extent to a file, etc) require
224 more than one device/sector level atomic operations (e.g. in order to add a new data extent at least two
225 sector write operation are required: one to update the Allocation table and one to write the data content).
226 These operations cannot be “recognized” by the device as belonging to the same transaction. Consequently,
227 a power failure event encountered when only a part of the atomic operations pertaining to such a
228 transaction were actually executed by the device, can lead to an inconsistency, which may result in an
229 irrecoverable data lost.

230 While, as mentioned, this problem may exist in **any** block device, it can be more significant and
231 problematic in a NAND Flash device. The reason for this situation is that the management of the NAND
232 Flash device (by FTL) requires additional mapping mechanisms (in addition to the Allocation Tables and
233 directory structures used by the File System) which in turn imply storage and manipulation of additional
234 metadata on the flash media. Furthermore, some of the flash management operations (for example the
235 copying and relocation of sectors mentioned in the previous section, and “garbage collection” operations
236 resulting from the need to prepare for reuse “used” but no longer valid pages and blocks) are significantly
237 time consuming. This might result in increasing the time separation between the actual execution of the
238 composite atomic operations pertaining to an application level transaction, and thus increasing the
239 sensitivity of the process to an eventual power failure.

240 **4.2 Longevity**

241 NAND flash devices usually wear out as the result of repeated programming and erasing; with the
242 contemporary device technologies the typical device lifetime is limited to about 100,000 program/erase
243 cycles per erase block (even lower for MLC device).

244 To maximize the useful life of the flash, care must be taken by the Flash management algorithms to spread
245 the wear evenly over the device, and in fact the different algorithms employed by different vendors invest a
246 significant effort (and contain significant amount of related Intellectual Property) to address this issue.

247 While this longevity problem is an inherent result of the nature of Flash device technology, its importance
248 and possible severity are intensified by the block device emulation practice: the need to repeatedly update
249 the File Allocation Tables create and represent a potential “hot spot” for the extended wear of the file
250 system related sectors of the flash media.

251 **4.3 Flash Media Reliability**

252 An additional peculiarity of the NAND Flash device technology is a problem usually referred to as “bit
253 flipping”, posing a significant reliability problem if not properly handled by the Flash Management
254 algorithms.

255 While in a magnetic disk for example imperfections of the recording surface/media are usually of
256 permanent nature, thus “bad blocks” can be detected, marked and avoided a-priori while formatting a
257 magnetic storage device, in NAND Flash devices there is a phenomenon called bit flipping (manifesting
258 itself by an inconsistency between the content intended to be written and the content actually read later)
259 which can evolve, rather than being permanent. The root cause of the problem may be not only marginal
260 imperfections on the chip but may also be wear resultant. This necessitates further, dynamic measures to
261 be exercised by the Flash management algorithms, implying both additional mapping for dynamic bad
262 block detection and avoidance.

263 Once again, although the bit flipping problems are not a direct result of the block device emulation method,
264 nevertheless, in the situation where application level transactions are split into several independent atomic
265 device/sector level operations, this creates yet another problem spot as a result of the management
266 operations required to handle this problem

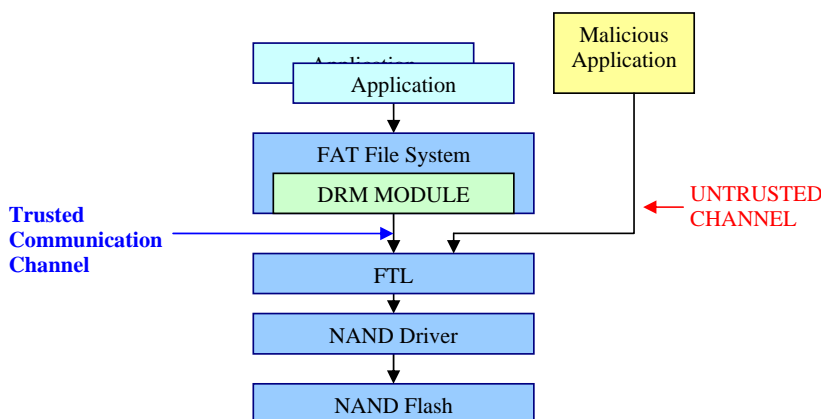
267 **5 Possible new application domain requirements**

268 The practical use of storage devices in complex applications involves the necessity to handle several
269 storage related functional requirements which are of higher level than the basic input/output or
270 read/write/delete transactions. Closer examination of several such functional requirements in the following
271 subsections reveals that the feasibility and/or efficiency of implementing these are dependent and
272 interrelated to the nature and functionality of the interface between the storage device and the host side
273 Operating System.

274 **5.1 Digital Rights Management (DRM)**

275 There is an increasing need in the industry for implementing controlled access to information content
276 stored on memory/storage devices, i.e. enforcing a DRM scheme with respect to the protected content. The
277 ideal solution for DRM would be the enforcement/implementation of the DRM access rules internally at
278 the storage device itself.

279 The nature of the protected data entity is a “content” object, e.g. a song, a picture, a document – these are
280 usually translatable to data objects, and most frequently to files. Clearly, the “sectors” as used the interface
281 with a block-device (or block-device emulation) based storage device, are well below the natural
282 granularity of the associated “rights” in any DRM scheme, thus the device implementation of **trusted and**
283 **reliable** DRM in these type of devices is practically impossible because it’s always possible to by-pass the
284 highest software module (e.g. FAT File System) and access the device through the FTL layer which has no
285 “knowledge” of DRM rights (figure 2).



286

287

Figure 2 Possible attack on FTL based DRM

288

However other, higher level abstraction interface schemes, where the content exchanged between a host and a device represents files or data objects, implementing/enforcing DRM at the device level is feasible.

289

290 **5.2 Compression**

291

In order to increase the effective capacity of a storage device, there is a requirement in some mass storage applications to implement transparent compression at the storage device level. This means that the volume of information actually stored/recorded onto the physical storage media is reduced by a compression algorithm implemented at the device on the information to be written, and a complementary decompression is performed on the read from the storage media before transferring it to the host.

292

293

294

295

296 A possible use case or application example: a handheld computer for use in a “mobile office” type
 297 application environment, where the flash disk in the device has to store a heterogeneous mix of file types
 298 and sizes, such as word processing documents, presentations, spreadsheets, messages, notes, etc.
 299 Obviously, the effective capacity of the flash drive is a factor of critical importance for the usefulness of
 300 such a device, so it is expected to be deemed advantageous by the end users if the effective capacity could
 301 be increased without paying for extra flash, or requiring increased size and/or power consumption, by
 302 employing data compression. File type dependent data compression (for example those well established
 303 compressed file types for pictures/video/music) might not be applicable or sufficient in this case, due to the
 304 heterogeneous nature of the file types as mentioned above, and the associated need to be compliant (i.e.
 305 readable, transparent, usable) with a diverse set of applications. Hence, if such a compression could be
 306 provided transparently and seamlessly by the storage device itself, this could be the optimal solution. (By
 307 transparency we mean for example that the application program is totally unaware that the 2.3 megabyte
 308 document it had just written to the flash disk is actually occupying only 0.68 megabytes of space).

309

In a block-device emulation environment, where the atomic transactions are based upon transparent sectors directly addressable by the host, there is practically very little which can be accomplished by compression decompression, since sectors cannot/must-not be concatenated, thus the efficiency of the is very limited. Another limiting factor is the mapping mechanism from file system sectors to the flash pages.

310

311

312

313 Another problem related to the compression is that usually multimedia data (e.g. MP3 files) are already
314 compressed and applying the compression algorithm don't cause a real benefit. With a block emulation
315 schema it's impossible to avoid this situation because (again) the FTL cannot know what kind of data it's
316 storing into the device. With a different software architecture that is not based on the Block Emulation
317 approach it could be possible to implement algorithm that (knowing the "nature" of the data) could
318 compress only that data that should be compressed.

319 **5.3 Encryption/security**

320 This subject is somewhat similar to the DRM issue described above: the implementation of some DRM
321 schemes involves encryption; however the topic of encryption is useful even out of the context DRM. It is
322 sometimes necessary to ensure that the content on a storage device is protected from unauthorized access
323 and use. As an extreme use case, for example, sometimes it is necessary to ensure that actual content stored
324 on the flash is encrypted, so it is incomprehensible without the possession of an encryption key which is
325 not stored on the media itself. The encryption and the corresponding keys are clearly content related: using
326 the traditional directory/file terminology, different files must be able to be encrypted different, user specific
327 keys, some files and/or directory entries needs to remain visible (non-protected, non-encrypted) while
328 others needs to be hidden (protected, encrypted).

329 In a block-device emulation environment, where the atomic transactions across the interface between the
330 storage device and the host are based upon sectors, and where the interpretation of the sectors content as
331 either net, "payload" data or file system structural/support info is only visible to the File System software
332 on the host, there is practically very little possibility to implement the necessary encryption/protection
333 provisions on the device itself, therefore the encryption/decryption must be inevitably done by the
334 application.

335 If however the block-device scheme would be replaced by some alternative scheme, where the interface
336 transactions are based on data objects or complete files (see below), then these schemes might **enable**
337 implementation of the encryption/decryption, if required, at the (future NAND-like) device level,
338 transparent to the application.

339 (The assessment whether device level versus application level implementation of encryption/decryption is
340 required/desirable/beneficial is left open for the discussion).

341

342 **5.4 Data portability**

343 One of the most important use cases for NAND Flash devices (and removable storage media and storage
344 devices in general) is their use as vehicles for data/information/content transfer among different platforms
345 (PC to PC, camera to PC, mobile phone to PDA, etc.). Obviously, in order to ensure a seamless flow of
346 content among different platforms, sometimes involving different Operating Systems, both types of the
347 software components involved need to be compatible: the File System on the host side and the data
348 representation on the media side. Within the contemporary block-device emulation scheme the device
349 stores and handles transparent sectors, some of which are data sectors (i.e. carrying the "net" content)
350 while some others are used solely by the host File System (like allocation tables, directory entries etc.). In
351 order to ensure that a computing platform, other than the one used in creating/writing the content, is
352 capable of reading and "comprehending" the stored content, the reading platform must contain a File
353 System which is binary compatible (with respect to the sectors stored on the device/media) with the writing
354 platform. This requirement is a severe and limiting one: while most contemporary platforms support the
355 FAT file system for cross platform data portability, this is obviously sub-optimal, since in this way the
356 advantages of the more advanced native file systems (e.g. NTFS or advanced UNIX/LINUX file systems)
357 are lost.

358 On the other hand, if an advanced/alternative interface technology is employed, based on data object
359 transactions rather than transparent sectors (see Section 7), then the problem of data portability is
360 automatically solved (or at least significantly diminishes).

361 While the primary focus of the MIPI NAND Software Workgroup is on embedded, i.e. non-removable,
362 applications of NAND Flash storage devices, nevertheless, this subject of data portability is relevant due
363 to the fact that the mobile/embedded hosts including the NAND Flash storage (as an integral non-
364 removable component) are frequently used as detachable/removable storage devices in other computing
365 platforms: for example a mobile phone (with embedded NAND Flash storage device) may act as a
366 removable storage device for a PC host when connected to a docking/sync station.

367 **6 Performance issues**

368 There are emerging streaming multimedia use case scenarios which may highlight performance issues with
369 NAND devices such as simultaneously recording a DVB-H stream at one point while rendering from the
370 same stream at another, advanced video trick plays, and advanced video telephony scenarios. For these use
371 cases, some concurrencies and sequences should be considered.

372 Because the relative performance overhead of the file system for NAND Flash is so small compared to
373 other activities within the system, the performance of the file system is most critical when considering time
374 extended use cases which require very little power for short periods of time. The most common use case
375 today falling into this category is the rendering of MP3/AAC music or audio books with competitive
376 playback times in excess of a hundred hours on a single charge. For use cases like this one, file system
377 performance impact to playback time is largely a function of the data rate.

378 Larger data rates with extended write scenarios might highlight file system for NAND Flash performance
379 even more. An emerging use case for consideration is that of recording a series of digital TV broadcasts.
380 Here again, the amount of time this scenario can be executed is a function of concurrencies implemented
381 between receive and write subsystems. If writes can be successfully executed concurrently with received
382 bursts, the operating frequency required to enable this concurrency may be a minor discriminating factor. If
383 the write subsystem is executed sequentially after reception, the amount of power consumed by the
384 receiver and the relative on time of the application processor in order to complete the write may become
385 discriminating factors.

386 The longevity of the recorded TV series scenario may also be compromised if a limited amount of memory
387 is also used as a circular buffer for data with a short life expectancy (such as concurrent commercial
388 offers).

389 The block emulation software solution (File System plus FTL) limits the possibility to achieve the best
390 performance possible and to guarantee a constant or at least a minima value of write and read throughput.
391 In fact the FS plus FTL solution is not a "NAND friendly" solution and it's not possible to implement
392 dedicated algorithms to manage the write or read operations. As usual FTL is not aware of the object (e.g.
393 file, directory) that the application is managing and it can only manage sectors (usually 512 byte size). In
394 this situation is almost impossible to implement (in a safe a robust way) algorithms that can use all the
395 potentiality of NAND throughput: read/write caching, multiple operations, large page architecture (e.g.
396 2KB, 4KB).

397 **7 Review of non Block Device oriented alternatives**

398 There are numerous studies conducted by various research, educational and industrial organizations
399 proposing various alternative (non block-device or non FAT oriented) storage organization and
400 management methods. The following sections include overviews of four of these alternatives which are
401 well known and/or covered by publicly available documents.

402 7.1 Journaling Flash File System (JFFS2 and JFFS3)

403 The Journaling Flash File System [JFFS], especially its advanced derivatives [JFFS2 and JFFS3] represent
404 an industry wide well known attempt to create a non block device emulation based NAND Flash File
405 System, with special attention to LINUX based embedded systems.

406 JFFS is a log-structured file system designed by Axis Communications AB in Sweden specifically for use
407 on flash devices in embedded systems, which is aware of the restrictions imposed by flash technology and
408 which operates directly on the flash devices, thereby avoiding the inefficiency of having two file systems
409 on top of each other.
410

411 The log-structured design in JFFS/JFFS2 basically means that the whole file system content, as stored on
412 the flash, may be regarded as one large log of changes: any modification to an entity of the host file
413 system (i.e., file change, directory creation, changing file's attributes, etc) is appended to the log. The log is
414 the only data structure on the flash media. Modifications are encapsulated into small data structures called
415 nodes. In principle, each node in the log contains a data object (file, directory, link) name and meta-data,
416 along with the (variable length) related data content.

417 The technical details related to the implementation and representation of the JFFS2 on the flash are beyond
418 the scope of this paper, however, the importance of this solution and its differentiating characteristics from
419 a block-device emulation based flash file system lays in the semantics difference in the interface with the
420 host: while in the case of the block-device emulation all the high level information regarding the data
421 objects is transparent to the device (i.e. from the device point of view the interaction is based on reading
422 and writing addressed sectors, all sectors are indistinguishable), in the case of the JFFS2, the atomic
423 transactions are changes related to named data objects (e.g. File-name and additional meta-data, data
424 segment offset, data segment length, data segment content).

425 7.2 SCSI Object-based Storage Device model (OSD)

426 The Technical Committee T10 of ANSI INCITS (responsible for the SCSI and related standards) had
427 created a published a working draft specification for Object-Based Storage Device Commands [OSD2].
428 The **Object-Based Storage Device (OBSD)** is defined as a device that stores **data objects** instead of
429 blocks of data. The purpose of this abstraction is to assign to the storage device more responsibility for
430 managing the location and organization of the data within the storage media.
431

432 Some important definitions, used in the following description, are:

433 **OSD object:** An ordered set of bytes within an object-based storage device that is associated with
434 a unique identifier. Data in the object is referenced by the identifier and offset information within
435 the object. Objects are allocated and placed on the media by the OSD logical unit.

436 **user object:** An OSD object that contains user data that is referenced by byte offset within the
437 OSD object.

438 **root object:** An OSD object that is always present whose attributes contain global characteristics
439 for the OSD logical unit. Each OSD logical unit has one and only one root object.

440 **collection:** An OSD object in which references to one or more user objects from a single partition
441 may be collected.

442 **partition:** An OSD object used for creating distinct management domains (e.g., for naming,
443 security, quota management).

444

445

446

447

448

449

448 OSD object abstraction as described in [OSD2]:

450 The OSD object abstraction is designed to re-divide the responsibility for managing the access to
451 data on a storage device by assigning to the storage device additional responsibilities in the area of

452
453
454

space management. Figure 3 shows the relationship between the OSD model and a traditional SBC-based (SCSI Block Commands) model for a file system.

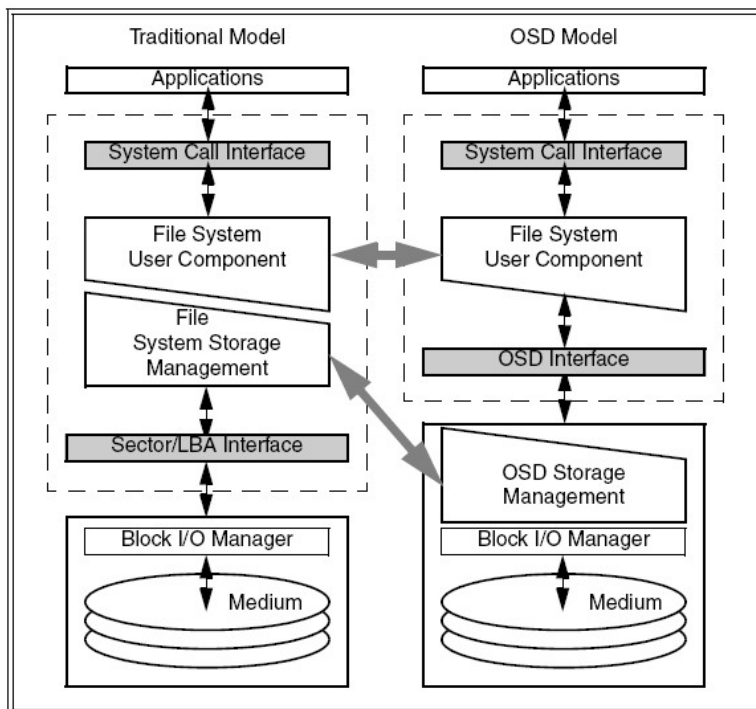


Figure 3: Comparison of traditional and OSD storage models

455
456

The user component of the file system contains such functions as:

457
458
459
460
461
462
463

- a) Hierarchy management;
- b) Naming; and
- c) User access control.

464
465
466
467
468
469
470
471
472
473
474
475
476

The storage management component is focused on mapping logical constructs (e.g., files or database entries) to the physical organization of the storage media. In the OSD model, the logical constructs are called user objects. The root object, partitions, and collections provide additional navigational aids for user objects.

In addition to mapping data, the storage management component maintains other information about the OSD objects that it stores (e.g., size, and usage quotas, and associated username) in attributes. The user component may have the ability to influence the properties of object data through the specification of attributes (e.g. directing that the location of an object to be in close proximity to another object or to have some higher performance characteristic) via mechanisms that are outside the scope of this standard.

In this model, the OSD makes the decisions as to where to allocate storage capacity for individual data entities and managing free space.

477
478

While the OSD model as defined by the T10 Committee applies to storage devices in general, it seems exceptionally well adapted to the specific case NAND Flash storage devices. The necessity and complexity

479 of managing a NAND Flash storage device usually requires the use of an intelligent device controller (the
480 “intelligence” can be software or hardware implemented) significantly more “smart” than what is required
481 to manage a “natural” block device (e.g. a rotating magnetic disk). This SW/HW controller can also be
482 used to implement the added functionality (or intelligence) required for the OSD functions.

483 Similarly to the case of the JFFS above, the interface with the host is based upon the atomic transactions
484 are related to named data objects as opposed to addressed sectors as used in the context of the traditional
485 block-devices.

486 Additional material related to initial industrial OSD projects and implementations can be found in [OSD
487 PROJ1], [OSD PROJ2] and [OSD DEMO]. [\[Check for the board approval\]](#)

488 7.3 FTP Storage

489 An intelligent storage device can in principle be configured with a functionality equivalent to an ftp (File
490 Transfer Protocol) server: the atomic interface transfer operations on such a device would be file related,
491 i.e. read/write (or get/put) and delete single or multiple files. In a sense this could be considered as a
492 private case of Object-based Storage Device with the “file” being the only object type.

493 For example a related implementation to this concept is the [PTP].

494 8 Attributes of the “ideal solution”

495 As stated above, the main objectives of the MIPI NAND Software Workgroup, in collaboration with
496 leading Operating Systems vendors and mobile/embedded computing platform vendors is to define/specify
497 a Standardized Software Interface between storage devices based upon the NAND Flash technology, and
498 the OS in embedded/mobile computing platforms.

499 This final paragraph attempts to characterize the intended end product (i.e. the Standardized NAND
500 Software Interface) in terms of a list of desirable attributes (“wish list”). Since the satisfaction of some of
501 these requirements may be mutually contradictory, this list should be considered as representing an “ideal
502 solution”, whereas the actual specification resulting, might involve certain compromises and/or options
503 and/or variations.

504 Thus, the Standardized NAND Software Interface should be:

- 505 ○ **Technologically Scalable, Extensible and Future Proof:** it should take into consideration the
506 foreseeable evolution path of both the NAND-like storage device technology and the evolution of
507 the storage related software technology;
- 508 ○ **Providing a “soft” migration path,** both for Operating Systems currently employing
509 contemporary and pre-standard block-device emulation, and for NAND Flash vendors and their
510 current devices;
- 511 ○ **Capable of continuing to support “legacy” systems and solutions,** employing pre-standard
512 devices and/or Operating Systems (e.g. retain support for the “good old FAT” and bootability
513 feature)

514 *NOTE* by bootability we mean: enable the loading of the executable Operating System image from
515 the storage device/media to the memory of the embedded application system processor main
516 memory, and facilitate secure and controlled means and methods for storing/updating/replacing
517 the OS image to the storage device/media).

- 518 ○ **Adequately address the new application domain requirements** (like DRM, compression,
519 security/encryption, data portability) as listed in above, and also considering new/additional
520 requirement resulting from the collaboration with the OS vendors;

- 521 ○ **Versatile**, i.e. compatible with, or at least adaptable to, all embedded computing platforms and
522 Operating Systems.

- 523 ○ **Robust and low-risk, by** relying on well founded and established software engineering principles
524 and practices.